Beliefs, learning, and personality in the indefinitely repeated prisoner's dilemma *

David Gill †

Yaroslav Rosokha ‡

This version: June 11, 2021

First version: July 15, 2020

Abstract

The indefinitely repeated prisoner's dilemma (IRPD) captures the trade-off between the short-term payoff from exploiting economic partners and the long-term gain from building successful relationships. We aim to understand how people form and use beliefs about others in the IRPD. To do so, we elicit beliefs about the supergame strategies chosen by others. We find that heterogeneity in beliefs and changes in beliefs with experience are central to understanding behavior and learning in the IRPD. Beliefs strongly predict cooperation, initial beliefs match behavior quite well, most subjects choose strategies that perform well given their beliefs, and beliefs respond to experience while becoming more accurate over time. Furthermore, experience affects both transitions between strategies and cooperation. Finally, we uncover a novel mechanism by which trusting subjects learn to cooperate through their interaction with experience, which helps to explain how trust underpins successful economic exchanges.

Keywords: *Infinitely repeated prisoner's dilemma; cooperation; optimism; belief elicitation; supergame strategies; experimentation; trust; experiment.*

JEL Classification: *C72; C73; C91; D91*

---

*We are grateful for helpful comments from seminar participants and in private conversations. We thank: the Blake Family Fund for Ethics, Leadership and Governance (administered by the Krannert School of Management) for funding; the Vernon Smith Experimental Economics Laboratory for hosting our experiment; and Benjamin Raymond, Peter Wagner, Lu Wang, Chen Wei and Junya Zhou for excellent research assistance.

†Purdue University, Department of Economics; dgill.econ@gmail.com

‡Purdue University, Department of Economics; yrosokha@purdue.edu

1 Introduction

Repeated interactions that persist for an uncertain length of time underpin many economic transactions and relationships.[1] In such environments, the indefinitely repeated prisoner's dilemma captures well the trade-off between the short-term payoff from exploiting economic partners and the long-term gain from building successful enduring relationships. Recent experimental work has advanced our understanding of behavior in the indefinitely repeated prisoner's dilemma (e.g., Dal Bó, 2005, Blonski et al., 2011, Dal Bó and Fréchette, 2011, Fudenberg et al., 2012, Bigoni et al., 2015, Aoyagi et al., 2019, Proto et al., 2019).[2] However, we know little about how people form and use beliefs about others when they play indefinitely repeated prisoner's dilemma games.

In this paper we want to understand the role of beliefs in the indefinitely repeated prisoner's dilemma (following the literature, we call one indefinitely repeated prisoner's dilemma game a 'supergame'). In particular, we aim to: (i) elicit initial supergame strategies and initial beliefs about the supergame strategies chosen by others; (ii) analyze the relationship between behavior and beliefs; (iii) use beliefs to provide new evidence about the mechanism that underlies learning from experience; and (iv) understand whether personality predicts behavior and beliefs. To achieve these aims, we build on recent work that: (i) directly elicits supergame strategies in the indefinitely repeated prisoner's dilemma after some experience of playing the game (Romero and Rosokha, 2018; Cason and Mui, 2019; Dal Bó and Fréchette, 2019);[3] and (ii) studies the role of personal traits in the indefinitely repeated prisoner's dilemma without eliciting strategies or beliefs about strategies (e.g., Dreber et al., 2014, Proto et al., 2019, 2020).

[1]For example: (i) firms compete in an industry until their products becomes obsolete; (ii) a labor union bargains with a firm for so long as the firm continues to exist; and (iii) countries sign free trade agreements that last until a protectionist government is elected.

[2]Recent experimental work has also studied the prisoner's dilemma with finite repetition (e.g., Embrey et al., 2018), in continuous time with a finite time horizon (Friedman and Oprea, 2012), when subjects are randomly rematched after every round of the stage game (e.g., Camera and Casari, 2009, Duffy and Fehr, 2018), and when decisions are made in teams (Cason and Mui, 2019; Cooper and Kagel, 2020).

[3]In their setting, Dal Bó and Fréchette (2019) find no evidence that eliciting supergame strategies systematically changes behavior.

We find that beliefs play a central role in explaining behavior throughout our experiment. In the first supergame, before subjects gain any experience, we find that heterogeneity in beliefs across subjects accounts for 36% of the variation in cooperation (this represents 80% of the variance accounted for together by beliefs, the return to joint cooperation, demographics, personality and cognitive ability). In the final supergame, after subjects have learned from experience, beliefs account for 23% of the variation in cooperation. Furthermore, we find that: (i) initial beliefs match behavior quite well; (ii) most subjects choose strategies that perform well given their beliefs; and (iii) beliefs respond to experience and become more accurate over time. In sum, we find that heterogeneity in beliefs and changes in beliefs with experience are crucial to understanding behavior and learning in the indefinitely repeated prisoner's dilemma. We also find that, by the end of the experiment, more trusting subjects cooperate more often and hold more optimistic beliefs, and we uncover a novel mechanism by which trusting subjects learn to cooperate through their interaction with experience.

Our 394 subjects played 25 supergames with random rematching and between-subject treatment variation in the the return to joint cooperation (cooperation is an equilibrium in all treatments). Since we wanted to elicit both supergame strategies and beliefs about the supergame strategies chosen by others, we restricted attention to ten strategies (see Table 1 in Section 2.4; we randomized the strategy presentation order across subjects). In each supergame the subject's chosen strategy and that of her opponent were played out round-by-round on the subject's screen. We elicited beliefs twice, in the first supergame and again in the final (25th) supergame. Building on Costa-Gomes and Weizsäcker (2008), we elicited incentivized beliefs about the distribution of the ten supergame strategies using the Quadratic Scoring Rule (QSR).[4] Web Appendix I provides screenshots from the experiment.

In order to elicit initial strategies and beliefs, we did not allow subjects to interact in any way with each other before eliciting strategies and beliefs in the first supergame. Nonetheless, we wanted subjects to understand the structure of the game and the nature of

[4]Costa-Gomes and Weizsäcker (2008) find that when using the QSR to elicit beliefs about a distribution over strategies, the belief elicitation has a mostly insignificant effect on behavior. Nonetheless, to avoid contamination of behavior, we elicited beliefs after subjects had chosen their strategy.

repeated game strategies before eliciting initial strategies and beliefs. To achieve this, we designed two forms of training. First, each subject played 'practice' supergames against herself using the direct-response method (that is, choosing round-by-round actions rather than supergame strategies). Second, each subject tested pairs of supergame strategies against each other in training supergames that were played out round-by-round on the subject's screen.

We also selected short directed measures of personality that we judged most likely to help explain behavior and beliefs in the indefinitely repeated prisoner's dilemma. At the beginning of the experiment we used short surveys to measure: anxiety; cautiousness; forgiveness; kindness; manipulativeness; and trust (we also measured cognitive ability and basic demographics). Using our setting where we elicited supergame strategies and beliefs over strategies, we aimed to increase the potential to detect any effects of personality on cooperation in repeated games by measuring aspects of personality that are conceptually linked to categories of strategies and to motivations in the indefinitely repeated prisoner's dilemma.

In Section 3 we analyze initial beliefs and behavior in the first supergame. We find that average beliefs in the first supergame match the distribution of chosen strategies quite well, although initial beliefs respond less strongly to the return to joint cooperation than does behavior. Furthermore, most subjects choose strategies that perform well given their beliefs, and beliefs strongly predict cooperation. Deviations from best responses to beliefs follow an interesting pattern: when the return to joint cooperation is low, optimists often fail to choose the payoff-maximizing always defect strategy; and when the return to joint cooperation is high, pessimists often fail to choose payoff-maximizing lenient strategies. We also find that expected earnings increase with the accuracy of beliefs and with the ability of subjects to choose strategies that perform well given beliefs.

In Section 4 we analyze how beliefs and behavior evolve in response to experience (of opponent cooperation and realized supergame length), and we use beliefs to provide new evidence about the mechanism that underlies learning from experience. The previous literature focuses on behavior in the first round of each supergame (Dal Bó and Fréchette, 2011, 2018), and we replicate these findings in our dataset. We then go beyond this analysis of first-round behavior in three ways. First, because we elicited strategies in consecutive supergames, we are able to show that experience affects cooperation at the level of the whole supergame strategy. We also study the factors that predict experimentation and use transitions between strategies to help understand how experience changes cooperation: for example, when a subject's opponent in the previous supergame cooperated more, the subject is much less likely to change to an "unfriendly" strategy (i.e., one which defects for sure in the first round). Second, because we elicited beliefs in the first supergame, we can show that cooperation depends on both experience and initial beliefs; thus we find that beliefs elicited in the first supergame predict cooperation throughout the experiment. Third, because we also elicited beliefs in the final supergame, we can show that beliefs themselves respond to experience and that beliefs in the final supergame predict cooperation at the end of the experiment. We also find that beliefs become more accurate over time as subjects gain experience.

In Section 5 we find that subjects who reveal themselves to be more trusting in our personality survey cooperate more on average. When we elicit beliefs in the final supergame, we also find that trust predicts optimism about others' cooperation, which suggests that beliefs mediate the relationship between trust and cooperation. Interestingly, we find no evidence that trust predicts behavior or beliefs in the first supergame. To understand how trusting subjects learn to cooperate, we study how trust interacts with experience. We find that subjects high in trust respond more strongly when they experience cooperative behavior; by contrast, subjects low in trust respond more strongly when they experience uncooperative behavior. The tendency of high trust subjects to amplify good news and discount bad news drives their cooperation up.

The main methodological contribution of our paper is to elicit beliefs over supergame strategies in the indefinitely repeated prisoner's dilemma. To the best of our knowledge, we are the first to elicit such beliefs. Davis et al. (2016) elicit round-by-round beliefs about the opponent's behavior in the indefinitely repeated prisoner's dilemma; like us, they find that beliefs correlate with behavior.[5] We complement the literature that measures beliefs in the one-shot or finitely repeated prisoner's dilemma, which consistently finds that

[5] Davis et al. (2016)'s main focus is on the relationship between personal characteristics and behavior, and so their analysis of beliefs appears only in the appendix (where they also find little correlation between beliefs and personal characteristics). Web Appendix II.1 summarizes other work that considers beliefs in indefinitely repeated prisoner's dilemma games.

beliefs correlate with cooperation.[6]

In contemporaneous work, Aoyagi et al. (2020) elicit round-by-round beliefs in definitely and indefinitely repeated prisoner's dilemmas, and develop a novel identification strategy to estimate beliefs over supergame strategies from the round-by-round beliefs.[7] Aoyagi et al. (2020)'s focus is different to ours: because they estimate beliefs using data from later supergames, they do not study initial beliefs over supergame strategies or the evolution of these beliefs in response to experience. Aoyagi et al. (2020) find that estimated beliefs over supergame strategies depend on whether the game is finitely or indefinitely repeated, that different behavioral types hold different beliefs, and that most subjects choose strategies that perform well given their beliefs.[8]

Turning now to personality, to the best of our knowledge we are the first to study the relationship between a survey measure of trust and behavior or beliefs in the indefinitely repeated prisoner's dilemma; the positive effect of trust on cooperation that we find is broadly consistent with evidence from related one-shot or finitely repeated games.[9] Nor are we aware of any previous studies that consider how trust interacts with experience

[6] In finitely repeated prisoner's dilemmas: (i) Zhang et al. (2019) elicit round-by-round beliefs, finding that beliefs and behavior correlate within rounds, and that beliefs in one round vary with behavior in the previous round; and (ii) Kagel and McGee (2016) find that over time teams form beliefs that the opposing team will defect toward the end of each supergame (where these beliefs are inferred from team chat). For literature on beliefs in one-shot prisoner's dilemmas, see Web Appendix II.2.

[7] Aoyagi et al. (2020) estimate beliefs under the identifying assumption that subjects correctly Bayes update from observed within-supergame histories. Because their approach is data intensive, they estimate beliefs at the level of behavioral types rather than at the individual level.

[8] Aoyagi et al. (2020) also study the round-by-round beliefs themselves in later supergames, finding that average beliefs track cooperation rates closely, that beliefs correlate with within-supergame experience, and that beliefs correlate with behavior. Unlike Croson (2000) in the one-shot prisoner's dilemma, they find no indication of important changes in behavior caused by the belief elicitation.

[9] Web Appendix II.3 summarizes this evidence. A few papers correlate behavior in trust games with that in the prisoner's dilemma, finding mixed results (Chaudhuri et al., 2002; Capraro et al., 2014; Haesevoets et al., 2015; Davis et al., 2016).

in the prisoner's dilemma.[10] Outside of the lab, trust matters for economic outcomes: for example, Algan and Cahuc (2010) find that trust fosters growth; Lopez-de Silanes et al. (1997) find that trust promotes performance in large organizations; and Aghion et al. (2010) find that trust lowers the demand for government regulation. Our finding that, in indefinitely repeated interactions, trust predicts cooperative behavior and beliefs about how much others cooperate provides one mechanism for the broader role of trust in underpinning successful economic exchanges. In the lab, Proto et al. (2019) find a transitory effect of agreeableness on cooperation in the indefinitely repeated prisoner's dilemma (they also find that the cautiousness facet of conscientiousness lowers cooperation, but only when subjects are matched by their level of conscientiousness).[11] Since trust is one facet of agreeableness, any effects of agreeableness on cooperation might partly be driven by trust.

Finally, Dal Bó and Fréchette (2018) survey the broader literature on the relationship between personal characteristics and cooperation in the indefinitely repeated prisoner's dilemma: noting that the current evidence is sparse, they conclude that there is as yet no robust evidence of a link between personal characteristics and cooperation when cooperation is an equilibrium, arguing instead that the main motivation behind cooperation is strategic. The role of trust that we uncover is consistent with a predominantly strategic rather than preference-based motivation for cooperation, since our evidence suggests that trusting subjects learn to cooperate, with the effect of trust on behavior mediated by beliefs.

The paper proceeds as follows: Section 2 describes the experimental design; Section 3 analyzes initial beliefs and behavior in the first supergame; Section 4 studies the evolution

[10] In a prisoner's dilemma game involving deception (subjects were told that they were playing against another subject in the room, but in fact were playing against the experimenter), and in which subjects were given no information about the number of times the game would be repeated, Parks et al. (1996) find that when subjects received a message from the experimenter stating that they planned to cooperate (defect), high (low) trust subjects responded by increasing (decreasing) cooperation. Parks et al. (1996) also find that the response to messages depends on the consistency of the message and behavior.

[11] Proto et al. (2020) also find some effect of agreeableness on cooperation in the indefinitely repeated prisoner's dilemma. Kagel and McGee (2014) find that agreeableness predicts cooperation in the finitely repeated prisoner's dilemma.

of beliefs and behavior over the course of the 25 supergames; Section 5 considers the relationship between personality and behavior and beliefs; Section 6 concludes; and the Web Appendix provides further details.

2 Experimental design

2.1 Procedures

We ran our experimental sessions at Purdue University's Vernon Smith Experimental Economics Laboratory (VSEEL) between mid-November 2018 and early February 2019. The study was reviewed by Purdue's Institutional Review Board, and all participants gave informed consent after reading the participant information sheet.

In total, 394 subjects participated, earning \$23.57 on average including a show-up fee of \$5.00 (the rate of exchange was \$1.00 for every two hundred experimental 'points'). Subjects were drawn randomly from the VSEEL student subject pool and invited to participate.[12] We excluded subjects who had participated in the related repeated prisoner's dilemma experiments reported in Romero and Rosokha (2018, 2019a,b) and Cason and Mui (2019).

The experiment was between subject, with three treatments that varied only in the structure of the stage-game payoff matrix. We ran 27 sessions, with nine sessions for each of the three treatments.[13] Sessions lasted just under one hour and included 12, 14, or 16 subjects, with 14.6 subjects on average.[14] The experiment was programmed in oTree (Chen et al., 2016). Web Appendix I provides screenshots from the experiment.

2.2 Overview

The experiment proceeded as follows:

1. Measurement of personal characteristics. We measured each subject's personality and cognitive ability. We included the following personality measures: anxiety;

[12]The VSEEL subject pool is administered using ORSEE (Greiner, 2015).

[13]We ran three sessions (one for each treatment) on each of nine separate days. Session start times were constant across days, and we balanced start times across treatments. On the ninth day, one session did not fill up, and so we ran that session exactly one week later.

[14]Dal Bó and Fréchette (2011)'s sessions included 14.8 subjects on average.

cautiousness; forgiveness; kindness; manipulativeness; and trust. We measured cognitive ability using a test of matrix reasoning.

2. Training. Each subject played ten 'practice' indefinitely repeated prisoner's dilemma games ('supergames') against herself using the direct-response method. Subjects then spent three minutes reading the description of the ten available supergame strategies. Finally, each subject spent five minutes testing pairs of strategies against each other in training supergames.

3. Supergames with strategy elicitation. Subjects played 25 supergames with random rematching. At the beginning of each supergame, each subject chose one of the ten available strategies to play the supergame, and then the subject's chosen strategy and that of her opponent were played out round-by-round on the subject's screen.

4. Belief elicitation. We elicited beliefs twice, in the first supergame and the final (25th) supergame. We elicited beliefs about the distribution of the ten strategies using the Quadratic Scoring Rule (QSR). To prevent any contamination of initial supergame strategies, we elicited beliefs in the first supergame immediately after subjects had chosen their strategy for the first supergame.

5. Demographic questionnaire. Subjects completed a short demographic questionnaire at the end of the experiment.

2.3 Supergame design

We call an indefinitely repeated prisoner's dilemma game a 'supergame'. Following Dal Bó and Fréchette (2011), our between-subject design used the stage-game payoff matrix in Figure 1 with $R \in \{32, 40, 48\}$ and a continuation probability $\delta = 0.75$. This payoff matrix determines the two players' payoffs in each round of a supergame. After each round, the supergame ends with probability $1 - \delta = 0.25$.

Each value of $R \in \{32, 40, 48\}$ corresponds to one of our three between-subject treatments. Dal Bó and Fréchette (2011) found substantial variation in cooperation rates across these values of R with $\delta = 0.75$, even though subgame-perfect Nash equilibria with full cooperation exist in all three cases.[15]

[15] Dal Bó and Fréchette (2011) also included treatments with $\delta = 0.5$.

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

Figure 1: Payoff matrix in each round of a supergame

Notes: In the experiment, we used neutral labels 'A' and 'B' to represent 'C' (cooperate) and 'D' (defect).

2.4 Choice of ten supergame strategies

Our aim was to elicit initial supergame strategies and initial beliefs about the supergame strategies chosen by others, and to study how these choices and beliefs change with experience.[16] Since we wanted to elicit both strategies and beliefs, we restricted attention to the ten supergame strategies in Table 1.[17]

Our strategy definitions follow Fudenberg et al. (2012), except for the two strategies among our ten that they do not include, namely DG and RAND.[18] Figure A.5 in Web Appendix I shows the wording of our ten strategies. The order that the ten strategies appeared on the subject's screen was randomized across subjects, and remained constant

[16]We study both initial behavior and learning. Many experimental game-theoretic studies consider only initial behavior by suppressing feedback (e.g., Costa-Gomes et al., 2001).

[17]In indefinitely repeated prisoner's dilemma games, Romero and Rosokha (2018, 2019a), Cason and Mui (2019) and Dal Bó and Fréchette (2019) elicit supergame strategies from a larger strategy space, while Bruttel and Kamecke (2012) ask subjects to choose a memory-1 rule after three rounds of the supergame.

[18]For simplicity, under our definition of G, a player does not defect unless her opponent has defected at least once; that is, the player does not respond to her own defections (the same is true for G2). In our setting without mistakes in implementation, our definitions of G and G2 are behaviorally equivalent to those in Fudenberg et al. (2012). We use the simpler definitions because: (i) they are easier to understand; and (ii) to avoid subject confusion about why a strategy specifies a response to a player's own unilateral deviation(s) that can never occur under that strategy. Similarly, Dal Bó and Fréchette (2019, p.3935) use the term 'Grim' to denote a memory-1 strategy that, in the absence of mistakes in implementation, is behaviorally equivalent to Grim as defined in Fudenberg et al. (2012).

throughout the session.[19]

Strategy	Description
AD	Always Defect
DG	Defect in first round, and then play Grim (ignoring own first-round defection)
DTFT	Defect in first round, and then play Tit-for-Tat
RAND	In every round, choose C with probability 0.5 and D with probability 0.5
G	Grim
2TFT	2-Tits-for-1-Tat
TFT	Tit-for-Tat
G2	Lenient Grim 2
TF2T	Tit-for-2-Tats
AC	Always Cooperate

Table 1: Description of our ten supergame strategies

Notes: Figure A.5 in Web Appendix I shows the wording of our ten strategies.

We selected strategies from the twenty considered by Fudenberg et al. (2012, Table 2) based on the results of the Strategy Frequency Elicitation Method (SFEM; see Dal Bó and Fréchette, 2011) applied to the data collected together by Dal Bó and Fréchette (2018) from 1,734 subjects playing indefinitely repeated prisoner's dilemma games, and then selecting the eight most popular strategies.[20] These eight strategies include the five 'key' strategies identified by Dal Bó and Fréchette (2018, Table 10).[21]

We added DG and RAND to the eight strategies from Fudenberg et al. (2012) selected by our SFEM exercise, giving the ten strategies in Table 1. DG is the the natural extension of G corresponding to DTFT as the extension of TFT.[22] RAND corresponds to

[19]We randomly created sixteen orders, one for each of the sixteen possible subjects in a session, such that every strategy appeared first in the order for at least one of the first twelve subjects (recall the minimum session size was twelve), and no strategy appeared first in the order for more than two subjects.

[20]The data include 32 combinations of paper and parameter values with $\delta > 0$. The most popular strategies were determined by running SFEM for each combination and taking the average.

[21]Fudenberg et al. (2012) call DTFT 'Exploitative Tit-for-Tat', while Dal Bó and Fréchette (2018, 2019) call it 'Suspicious Tit-for-Tat'; we use the neutral term 'DTFT' to avoid implying a motive for choosing this strategy.

[22]Our strategy DG is not equivalent to Fudenberg et al. (2012, fn.25)'s D-Grim, which responds to a player's own first-round defection and so is behaviorally equivalent to AD in our

‘RANDOM-50’ in Dal Bó and Fréchette (2019).[23] We included RAND as an option for subjects who had difficulty choosing among the other strategies (RAND is the equivalent of level-0 behavior in stage-game strategies). RAND also ensures that, despite the limited number of available strategies, subjects never perfectly learn their current opponent’s deterministic strategy,[24] which increases external validity of the learning process about the population across supergames. Finally, RAND creates more behavioral separation between G-type strategies (DG, G, G2) and TFT-type strategies (DTFT, 2TFT, TFT, TF2T) since random defection(s) under RAND induce persistent punishment in the first case but not the second.

2.5 Strategy categories

We find it useful to categorize our ten supergame strategies, as illustrated in Figure 2.

Along the horizontal axis, we categorize strategies according to when they first defect. The three ‘unfriendly’ strategies (AD, DG, DTFT) defect in the very first round. The three ‘provocable’ strategies (G, 2TFT, TFT) start by cooperating but defect immediately in response to the opponent’s first defection. The three ‘lenient’ strategies (G2, TF2T, AC) start by cooperating and do not defect immediately in response to the opponent’s first defection.[25]

Along the vertical axis, we categorize strategies according to whether they punish a rival’s defection forever or whether, after punishing a rival’s defection, they eventually relent and cooperate if the opponent cooperates. The three ‘unrelenting’ strategies are the G-type strategies (DG, G, G2). The four ‘relenting’ strategies are the TFT-type strategies (DTFT, 2TFT, TFT, TF2T).[26] The three ‘non-responsive’ strategies (AD, RAND, AC) never respond to a rival’s defection.

setting without mistakes in implementation.

[23]The round-by-round randomization for RAND was implemented in real time as play progressed.

[24]The reason is that RAND replicates every deterministic strategy with positive probability.

[25]Fudenberg et al. (2012) categorize strategies as lenient in the same way that we do here. According to Fudenberg et al. (2012)’s terminology, our unfriendly strategies are ‘fully noncooperative’, while our ‘provocable’ strategies are ‘fully cooperative’ but not lenient.

[26]We use the term ‘relenting’ rather than ‘forgiving’ because Axelrod (1980)’s concept of

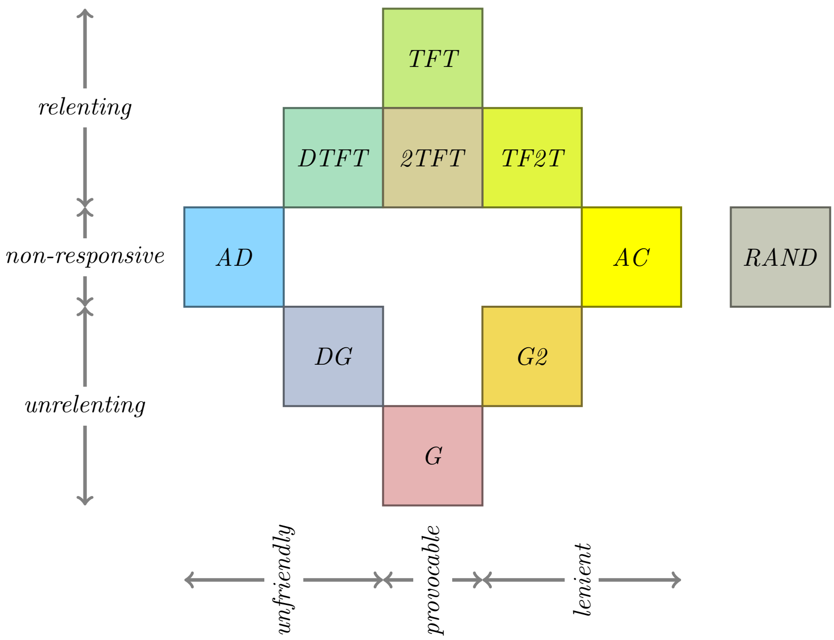

Figure 2: Strategy categories

2.6 Training

As noted in Section 2.4, one of our main aims was to elicit initial supergame strategies and initial beliefs about the supergame strategies chosen by others. Thus, we did not allow subjects to interact in any way with each other before eliciting strategies and beliefs in the first supergame.

Nonetheless, we wanted subjects to understand the structure of the game and the nature of repeated game strategies before eliciting initial strategies and beliefs. To achieve this, we designed two forms of training (without incentives). First, each subject played ten 'practice' supergames against herself using the direct-response method: in each round of each 'practice match', the subject chose an action for herself and an action for the 'other'; and at the end of each round the subject pressed a button to roll a computerized four-sided die that determined whether the supergame ended. Figures A.3 and A.4 in Web Appendix I show screenshots.

Once subjects had become familiar with the structure of the game from this direct-response training, we introduced subjects to the ten available supergame strategies. In particular, subjects spent three minutes reading the descriptions of the ten strategies,

forgiving in the prisoner's dilemma includes both relenting after a punishment and being lenient by not immediately punishing a defection.

and then each subject spent five minutes testing pairs of strategies against each other in training supergames that were played out round-by-round on the subject's screen (each round lasted a short amount of time, to mimic the feedback from direct-response play). On average subjects tested 16.2 pairs of strategies (including repeat tests of the same pair). In each 'test match', the subject chose one of the ten 'plans' for herself and one of the ten for the 'other'. Figures A.5 and A.6 in Web Appendix I show screenshots.

2.7 Supergames with strategy elicitation

Subjects played 25 supergames, with random rematching between supergames and the same supergame lengths across treatments.[27,28] At the beginning of each supergame, each subject chose one of the ten available strategies to play the supergame, and then the subject's chosen strategy and that of her opponent were played out round-by-round on the subject's screen.[29] In order to mimic the feedback from direct-response play, each round lasted two seconds, and at the end of each round the subjects observed the outcome of the four-sided die role that determined whether the supergame ended. The full supergame history (round-by-round choices, payoffs and die rolls) was displayed during the supergame and after the supergame ended (until the subject chose to continue to the screen for the next supergame). Subjects observed the round-by-round choices made by their opponent's strategy, but they did not directly observe the strategy chosen by the opponent.[30] Figures A.8 and A.11 in Web Appendix I show screenshots.

Standard theoretical analyses of repeated games (and the SFEM procedure for estimat-

[27]We randomly drew the lengths of the 25 supergames in advance. In particular, we randomly drew nine sequences of 25 supergame lengths; that is, we drew a new sequence for each of the nine sessions of a particular treatment. To keep supergame lengths the same across treatments, we used the same nine sequences for each of the three treatments.

[28]Dal Bó and Fréchette (2011, 2019) also use random rematching. Dal Bó and Fréchette (2011, fn.4) provide evidence that random rematching does not induce repeated-game effects across supergames.

[29]Previously chosen strategies did not act as defaults: subjects made an active choice of strategy at the beginning of each supergame.

[30]This choice preserves external validity: in real-world strategic interactions, people usually do not directly observe others' strategies.

ing strategies) assume that agents choose a supergame strategy, which specifies actions for the whole supergame. Eliciting such a supergame strategy is an example of the strategy method: Brandts and Charness (2011) survey the experimental evidence on the strategy method more generally. Compared to the direct-response data in Dal Bó and Fréchette (2011), Dal Bó and Fréchette (2019) find no evidence that eliciting supergame strategies systematically changes behavior in their setting (like us, Dal Bó and Fréchette, 2019, included treatments with $\delta = 0.75$ and $R = 32$ or $R = 48$).

2.8 Belief elicitation

We used the Quadratic Scoring Rule (QSR) to elicit incentivized beliefs about the distribution of strategies in the first supergame and in the final (25th) supergame. The QSR is incentive compatible (Selten, 1998), which means that money-maximizing (risk-neutral) subjects are incentivized to report their true belief.

Figure A.9 in Web Appendix I shows how we described the belief elicitation procedure to the subjects. We endeavored to keep the description concise. The text of the second and third lines is similar to that used by Costa-Gomes and Weizsäcker (2008). Following Costa-Gomes and Weizsäcker (2008): (i) we told subjects that they would make the most money if they reported their true beliefs; but (ii) we also provided a complete description of the QSR. Figure A.10 shows the belief entry screen, which included a reminder of the main points.

Costa-Gomes and Weizsäcker (2008) find that when using the QSR to elicit beliefs about a distribution over three strategies in one-shot games, the belief elicitation has a mostly insignificant effect on behavior.[31] More broadly, Schotter and Trevino (2014)'s survey concludes that eliciting beliefs either has no effect on behavior or hastens learning, and so is mostly innocuous.

Nonetheless, to prevent any contamination of initial supergame strategies, we elicited beliefs about behavior in the first supergame after subjects had chosen their strategy for the first supergame (but before the first supergame was played out), and subjects did

[31]Other papers that use the QSR to elicit beliefs about a distribution over three or more choices include Terracol and Vaksmann (2009), Danz et al. (2012), Hyndman et al. (2012) and Gee and Schreck (2018).

not know that beliefs would be elicited when they chose their strategy.[32] To minimize contamination of later supergame strategies, we elicited beliefs only twice, in the first supergame and in the final supergame.[33] Furthermore, we gave the subjects no feedback about the accuracy of their guesses.

Given that we elicit a belief about a distribution over ten strategies, we wanted to keep the belief elicitation procedure as simple as possible. In this respect, the QSR has the advantage that it is deterministic: that is, the subject's payoff depends deterministically on their reported belief and the realized state. Schlag and van der Weele (2013) show theoretically that all deterministic scoring rules impair truth-telling incentives for risk-averse subjects. However, in our setting, we judged that introducing an element of randomization would make the belief elicitation procedure too complicated.[34] Furthermore, in our setting with ten strategies, the bias toward flattening the reported distribution is unlikely to be important: Harrison et al. (2017) find that for empirically plausible levels of risk aversion, the bias is small unless the set of events over which beliefs are elicited is binary or close to binary.[35]

---

[32]Similarly, we elicited beliefs in the final supergame after subjects had chosen their strategy but before the supergame was played out. Schlag et al. (2015, p.479)'s survey finds no consensus on whether beliefs are influenced by first making a choice.

[33]After eliciting beliefs in the first supergame, subjects were told on a transition screen that they would not be asked a similar question until the end of the experiment.

[34]Furthermore, even if our subjects could understand the mechanics of a belief elicitation procedure with randomization, they might still not understand why the randomization gives the incentive to report truthfully with risk aversion. Schlag et al. (2015, p.482)'s survey discusses the contradictory evidence on whether randomized payments induce risk neutrality even in simple settings.

[35]We do not expect hedging due to risk aversion to be a significant concern in our complex setting. Schlag et al. (2015, p.481)'s survey summarizes evidence that hedging across actions and beliefs is more of a problem in simple environments. For example, Blanco et al. (2010) find hedging in a coordination game with obvious hedging incentives, but find no hedging in a more complicated prisoner's dilemma game. As noted above, Costa-Gomes and Weizsäcker (2008) use the QSR to elicit beliefs about a distribution over three strategies, and they find no evidence of hedging.

2.9 Measurement of personal characteristics and demographic questionnaire

2.9.1 Personality questionnaire

We measured personality at the beginning of the experiment. In particular, we measured: anxiety; cautiousness; forgiveness; kindness; manipulativeness; and trust. We selected these six personality measures because we judged them most likely to help explain behavior and beliefs in the indefinitely repeated prisoner's dilemma. As part of this exercise, we carefully read through the questions underlying a large number of personality measures. By design, we selected short directed measures of personality rather than longer measures that confound different concepts. For this reason, our three measures that come from the Big Five (John et al., 2008) capture specific facets of the five broader personality measures: anxiety is one of six facets that make up neuroticism; cautiousness (sometimes called 'deliberation') is one of six facets that make up conscientiousness; and trust is one of six facets that make up agreeableness.[36]

We included forgiveness, kindness and trust because we judged that these measures linked well to the strategy categories described in Section 2.5 (unfriendly, provocable, lenient, and relenting/unrelenting); indeed, the questions underlying the forgiveness measure relate to aspects of leniency and of being relenting, and thus this measure captures the spirit of Axelrod (1980)'s concept of 'forgiving' (see footnote 26 above). We included manipulativeness because the underlying questions capture a willingness to exploit others. We included anxiety because we conjectured that anxiety might affect the ability to perform in strategic interactions.[37] Finally, we included cautiousness because Proto et al. (2019) find a negative association between this facet of conscientiousness and cooperation in an indefinitely repeated prisoner's dilemma.

We gave subjects six minutes to complete the personality questionnaire (the screen

[36]Web Appendix II.4 describes the sources for our 52 personality questions.

[37]Anxiety is an important facet of neuroticism. Gill and Prowse (2016) find a negative association between neuroticism and performance in a repeated p-beauty contest game, Al-Ubaydli et al. (2016) find that neuroticism negatively predicts joint cooperation in a finitely repeated prisoner's dilemma, while DeYoung et al. (2010) find that neuroticism correlates with volume in areas of the brain associated with threat and punishment.

showed a countdown clock).[38] We randomly drew the order of the 52 questions, subject to the constraint that no two consecutive questions could come from the same personality measure (subjects all faced the same order). We told subjects that their answers would not affect the experiment in any way.

2.9.2 Personality factors

As described in Section 2.9.1, we selected short directed measures of personality from a number of different sources. As a result, some of our measures are highly correlated with each other (e.g., the highest correlation of 0.52 is between trust and forgiveness), although our personality measures are not correlated with cognitive ability (the highest correlation is 0.08). The high degree of correlation between our six measures of personality justifies the construction of a smaller number of uncorrelated factors to identify the effects of personality on behavior and beliefs. We undertake a principal factor analysis using maximum likelihood factoring and Varimax rotation, implemented to give factors that are uncorrelated with each other (see Luo et al., 2019). Before rotation, five factors have eigenvalues above one, and so these are retained in the rotation; retaining factors with eigenvalues above one is a standard criterion for choosing the number of factors due to Kaiser (1960).

The five factors are: anxiety; cautiousness; kindness; manipulativeness; and trust. Each factor's loadings are highest for the questions underlying one of the personality measures, and so we name each factor after that personality measure.[39] By construction, the factors are uncorrelated with each other, and have zero mean and unit variance.

2.9.3 Cognitive ability test: Matrix reasoning

After measuring personality but before subjects played any games, we measured cognitive ability using a matrix reasoning test that measures fluid intelligence.[40] In particular, we

[38] We followed the time per-question recommended for the Big Five Inventory (John et al., 2008, p.137). Web Appendix II.5 describes further implementation details.

[39] Web Appendix II.6 provides more information about the factor loadings.

[40] Fluid intelligence is "the ability to reason and solve problems involving new information, without relying extensively on an explicit base of declarative knowledge" (Carpenter et al., 1990). Matrix reasoning tests have been used in economics by, e.g., Burks et al. (2009), Charness et al.

used the eleven-question matrix reasoning test developed by the International Cognitive Ability Resource Team (ICAR), which is similar to the Raven Progressive Matrices test (Raven et al., 2000).[41] For each question, subjects have to identify (among six choices) the missing element that completes a visual pattern.

We gave subjects seven minutes to complete the test (the screen showed a countdown clock). We told subjects that their answers would not affect the experiment in any way. Following the convention in the psychometric literature, we did not provide monetary incentives for completing the test, and we did not tell subjects anything about their performance.

2.9.4 Demographic questionnaire

At the end of the experiment subjects completed a short demographic questionnaire. We asked subjects whether: (i) they were aged 'under 20' or '20 and over'; (ii) they were 'male' or 'female'; (iii) their major was in 'Economics or Management', 'STEM (Science, Technology, Engineering, Math)', 'Liberal Arts' or 'other'; (iv) they went to high school 'in US' or 'outside of US'. In each case, the subject could report 'prefer not to say'. Four subjects did not complete the questionnaire (answering 'prefer not to say' to one or more questions), and so we exclude those subjects from regressions that control for demographic characteristics. Those regressions also use a binary major categorization ('STEM' or 'not STEM').

3 Initial beliefs and behavior

3.1 Strategy choices, average beliefs, optimism and cooperation

We start by presenting raw data that describe the initial supergame strategies chosen by our subjects and their initial beliefs about the supergame strategies chosen by others. Figures 3(a) to 3(c) show the distribution of strategies chosen in Supergame 1 for each treatment. When the return to joint cooperation is low (i.e., $R = 32$), unfriendly strategies are popular. As the return to joint cooperation increases (i.e., as R increases from 32

(2018), Gill and Prowse (2016), Fe et al. (2019) and Proto et al. (2019).

[41] See Condon and Revelle (2014). For research purposes, the questions are available on request from ICAR (`https://icar-project.com`).

to 48), the proportion of unfriendly strategies falls substantially, while provocable and lenient strategies become more popular. In Supergame 1, our subjects choose unfriendly strategies at broadly similar rates to Dal Bó and Fréchette (2019)'s subjects in their menu elicitation treatment when, as here, $\delta = 0.75$.[42]

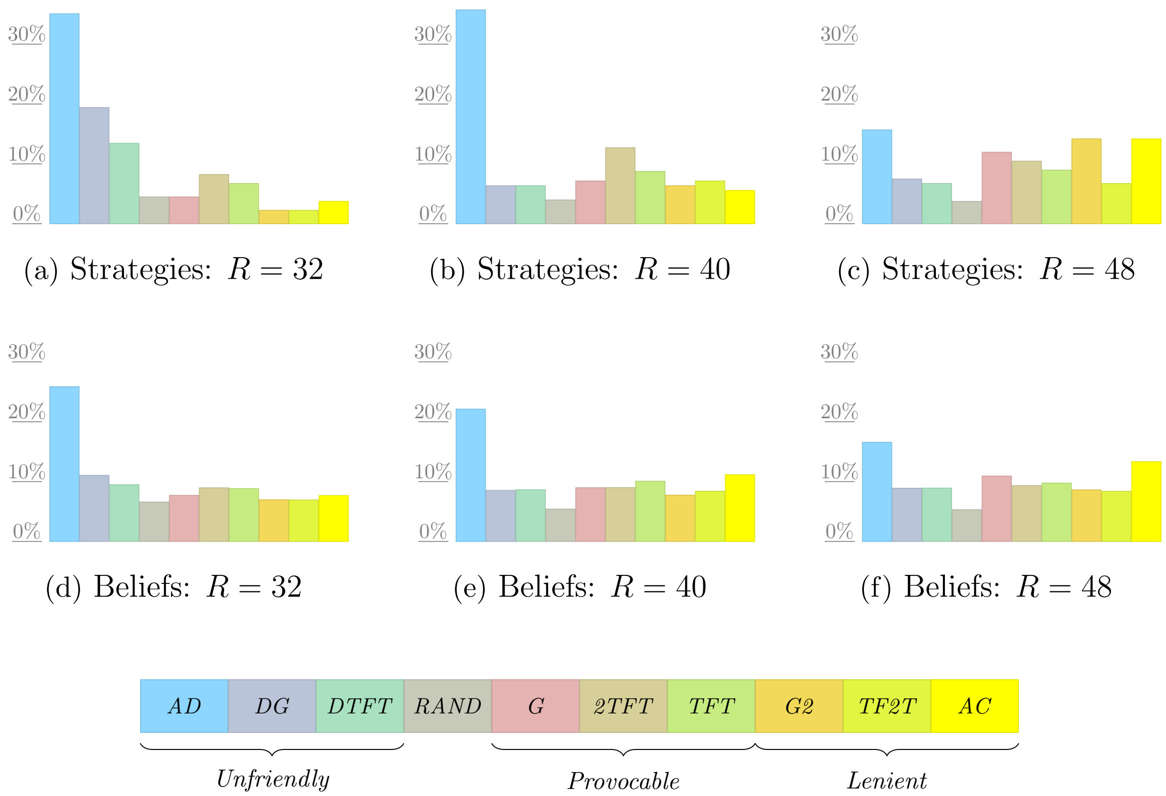

Figure 3: Strategies and average beliefs in Supergame 1

Notes: Figures 3(a) to 3(c) show the distribution of strategies chosen in Supergame 1 at the treatment level. Figures 3(d) to 3(f) show the mean probability weight placed on each strategy across each subject's belief distribution in Supergame 1 at the treatment level.

Figures 3(d) to 3(f) show average beliefs in Supergame 1 for each treatment. Broadly speaking, average beliefs match behavior quite well. This is true even though subjects did not interact in any way before eliciting strategies and beliefs in the first supergame. When the return to joint cooperation is low, subjects underestimate the proportion of unfriendly strategies, but when the return to joint cooperation is high, average beliefs track behavior closely. Generally speaking, average beliefs respond to the treatment less strongly than does behavior. When interpreting these results on average beliefs, the reader

[42]When $R = 48$ ($R = 32$), 30 (68) percent of our subjects choose unfriendly strategies, while 27 (57) percent do so in Dal Bó and Fréchette (2019) (who elicit strategies after experience of direct-response play).

should bear in mind the considerable heterogeneity in beliefs across subjects (illustrated below in Figure 4).

Our finding that average beliefs match behavior quite well could reflect projection bias, whereby subjects project their thinking onto others and so assume that others are likely to choose the same strategy that they themselves chose. Although projection bias is likely part of the story, we note that the majority of subjects (60%) report beliefs that place 20% or less weight on the subject's chosen strategy.

In order to delve more deeply into the data, we construct summary measures of initial beliefs and behavior. Briefly, Figure 4 shows that: (i) when the return to joint cooperation increases, optimism about how often others will cooperate responds more slowly than subjects' own willingness to cooperate and the actual level of cooperation in the population; (ii) when the return to joint cooperation is low, subjects are excessively optimistic, but beliefs become more accurate as the return to joint cooperation goes up; and (iii) within-treatment, subjects who are more optimistic tend also to be more cooperative.

When interpreting the results from Figure 4, note that the horizontal bars in Panels I, II, IV and V show 95 percent confidence intervals, while Panels III and VI report 95 percent confidence intervals from OLS regressions. Turning to the first row of Figure 4, 'Optimism' measures how often a subject expects others to cooperate, while 'Cooperativeness' measures how often a subject expects her chosen strategy to cooperate given her beliefs about others (and so captures willingness to cooperate). We see that when the return to joint cooperation increases (i.e., when R increases from 32 to 48), subjects become: (i) more optimistic (Panel I); and (ii) more willing to cooperate (Panel II). Interestingly, cooperativeness responds more strongly to the treatment than does optimism. Finally, within-treatment, there is a strong positive correlation between optimism and willingness to cooperate (Panel III).

Turning to the second row of Figure 4, 'OptimismRelTruth' measures optimism relative to how often others actually cooperate, while 'Cooperation' measures how often a subject's chosen strategy cooperates on average given how others actually behave. We see that when the return to joint cooperation is low ($R = 32$), subjects are too optimistic relative to the truth (Panel IV). Because subjects are overly optimistic when $R = 32$, the level of cooperation (Panel V) is lower than the willingness to cooperate (Panel II). However, when $R = 48$ optimism matches well the degree of cooperation actually present in the population (Panel IV), and so cooperation (in Panel V) and willingness to cooperate (Panel II) match up well. Finally, Panel VI shows that, within-treatment, subjects who are more optimistic relative to the truth also cooperate more.

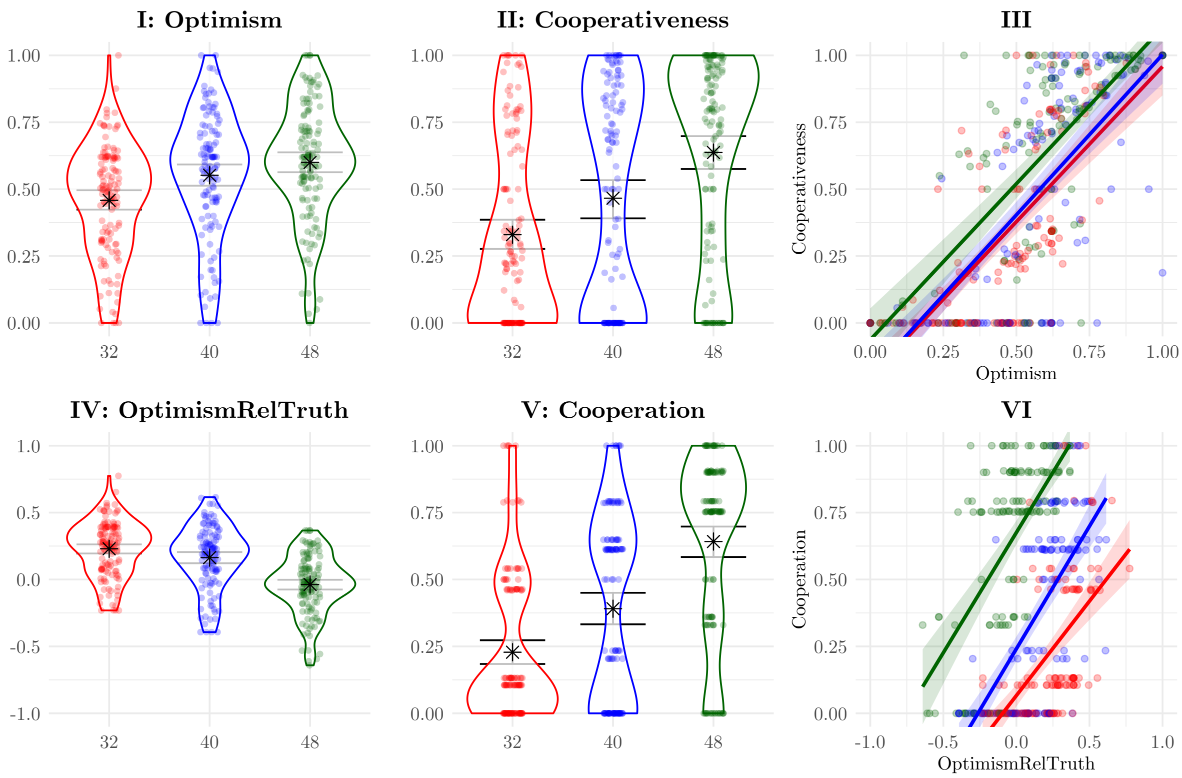

Figure 4: Beliefs and behavior in Supergame 1: Violin and scatter plots

Notes: 'Optimism' measures the expected cooperation rate of a subject's belief distribution playing against itself, while 'Cooperativeness' measures the expected cooperation rate of a subject's chosen strategy playing against the subject's belief distribution (the expected cooperation rate of a distribution of strategies playing against itself is the weighted sum of the expected cooperation rate of each strategy playing against the distribution, with the weight on each strategy given by the distribution). 'OptimismRelTruth' measures optimism minus the expected cooperation rate of the treatment-level strategy distribution (excluding the subject's own choice) playing against itself, while 'Cooperation' measures the expected cooperation rate of a subject's chosen strategy playing against the treatment-level strategy distribution (excluding the subject's own choice). In all cases, cooperation rates measure the expected number of rounds of cooperation divided by the expected number of rounds, and are based on analytical calculations of cooperation rates for every possible combination of strategies (see Table A.2 in Web Appendix V). In the violin plots, the unit of observation is an individual subject, stars are means and horizontal bars are 95 percent confidence intervals, calculated using non-parametric bootstrapping. The scatter plots report linear OLS regression lines and 95 percent confidence intervals.

Panel IV of Figure 4 suggests that the accuracy of beliefs increases as the return to joint cooperation goes up. To test this, we construct a measure of accuracy of beliefs about the within-treatment level of cooperation, and then regress this measure of accuracy on the treatment.[43] We find that accuracy does indeed increase as the return to joint cooperation goes up, with the effect statistically significant at the one-percent level.[44]

3.2 Responding to beliefs

3.2.1 Introducing good responding to beliefs

The subjects in our experiment face a difficult choice among ten strategies. Furthermore, in the first supergame, we elicit initial supergame strategies and beliefs before subjects gain any experience of play against other subjects. As a result, it is not surprising that subjects frequently fail to best respond to their initial beliefs: on average, in the first supergame 25.4 percent of subjects best respond perfectly to their beliefs.

Since best responding perfectly is difficult in our environment, we consider a less stringent, more pragmatic, definition of best response that we call 'good response'. In particular, we say that a subject good responds to her beliefs if she chooses a strategy that achieves an expected payoff within 3.15 percent of that from the best response (given the subject's beliefs). We chose this threshold so that exactly 50 percent of our subjects good respond to their beliefs in the first supergame, which is similar to the rate of best responding in simpler environments (see Alempaki et al., 2019, Costa-Gomes and Weizsäcker, 2008, and the references therein). To calculate the proportion of good responses and best responses, we use subjects' beliefs together with analytical calculations of payoffs for every combination of strategies (see Tables A.5 to A.7 in Web Appendix V).

Data from the first supergame support good responding as a useful measure. First, a payoff loss of up to 3.15 percent is small relative to the range of losses across subjects: Figure A.12 in Web Appendix IV shows the cumulative distribution function of payoffs

[43] We base our measure on the absolute value of OptimismRelTruth, which captures the deviation from the truth of the subject's expectation about how much others cooperate (the notes to Figure 4 define OptimismRelTruth). In particular, we define accuracy to be the negative of the absolute value of OptimismRelTruth; we take the negative so that accuracy increases (toward zero) as beliefs become more accurate.

[44] Web Appendix II.7 describes the regression.

relative to best responding. Second, good responding matters for outcomes: in Section 3.3 we show that good responding is a strong predictor of earnings.[45] Third, good responding changes with beliefs: in Section 3.2.2 we show that the frequency of good responding varies with the optimism of subjects' beliefs.

3.2.2 Analysis of good responding

Table 2 shows the frequency with which each strategy is a good response to subjects' beliefs, split by treatment. When the return to joint cooperation is low ($R = 32$), the unfriendly strategy AD is a good response for almost all subjects (97 percent), while DG and DTFT are good responses for around 35 percent of subjects. When $R = 40$, the lenient strategies G2 and TF2T are good responses for around 75 percent of subjects, while the provocable strategies G, 2TFT and TFT are good responses for around 55 percent, and AC is a good response for around 45 percent. When the return to joint cooperation is high ($R = 48$), the lenient strategies G2 and TF2T are good responses for almost all subjects (94 percent), and AC is a good response for around 60 percent; perhaps surprisingly, when $R = 48$ the provocable strategies G, 2TFT and TFT are good responses for only around 35 percent of subjects.[46]

R	*AD*	*DG*	*DTFT*	*RAND*	*G*	*2TFT*	*TFT*	*G2*	*TF2T*	*AC*
32	0.97	0.37	0.34	0.03	0.13	0.13	0.12	0.10	0.09	0.04
40	0.33	0.17	0.16	0.02	0.58	0.56	0.55	0.75	0.75	0.45
48	0.07	0.09	0.09	0.01	0.34	0.33	0.31	0.94	0.94	0.60

Table 2: Frequency each strategy is a good response in Supergame 1

Notes: For each strategy, the table shows the proportion of subjects for whom that strategy is in the subject's set of good responses (given the subject's beliefs), split by treatment. Good responding is defined in the second paragraph of Section 3.2.1. Table A.3 in Web Appendix V replicates the table for best responding.

[45]The relationship between good responding and earnings is not immediate, since good responding is defined relative to subjects' beliefs, while earnings depend on realized choices of others.

[46]Compared to lenient strategies, provocable strategies provide more protection against AD. However, unlike lenient strategies, provocable strategies never achieve mutual cooperation against DG and DTFT, which matters most when $R = 48$.

The frequency with which subjects good respond to their beliefs interacts with optimism in an interesting way. Recall from Section 3.1 that optimism measures how often a subject expects others to cooperate. We find that optimism is unhelpful when the return to joint cooperation is low ($R = 32$), in the sense that optimism reduces the probability that subjects good respond to their beliefs, while optimism is helpful when the return to joint cooperation is high ($R = 48$).[47] As evidenced by Table A.4 in Web Appendix V: (i) when the return to joint cooperation is low, optimists good respond less frequently because they often fail to understand that the unfriendly strategy AD is a good response to their (relatively) optimistic beliefs; and (ii) when the return to joint cooperation is high, pessimists good respond less frequently because they often fail to understand that the lenient strategies G2 and TF2T are good responses to their (relatively) pessimistic beliefs.

3.3 The determinants of earnings

In this section, we analyze the determinants of earnings. In particular, we want to understand how earnings in the first supergame depend on subjects' initial beliefs and behavior given those beliefs. Earnings in Supergame 1 are noisy, since they depend on the behavior of the specific opponent that a subject is matched with. To reduce this noise, we analyze subjects' expected earnings given their choice of strategy and how others behave within-treatment (recall that subjects had not yet interacted with each other when we elicited strategies and beliefs in the first supergame).

To put our analysis in context, Panel I of Figure 5 shows expected earnings in the first supergame by treatment, while Panel II shows expected earnings as a proportion of the maximum available (from choosing the strategy that performs best in expectation given how others actually behave). The first panel shows that, unsurprisingly, expected earnings increase with the return to joint cooperation. The second panel shows that subjects generally leave little money on the table: on average, subjects achieve expected earnings of around 95 percent of the maximum.[48]

[47]Web Appendix II.8 describes the regression.

[48]A subject who achieved expected earnings of 100 percent of the maximum would still leave money on the table relative to the best response to the specific strategy chosen by her randomly selected opponent.

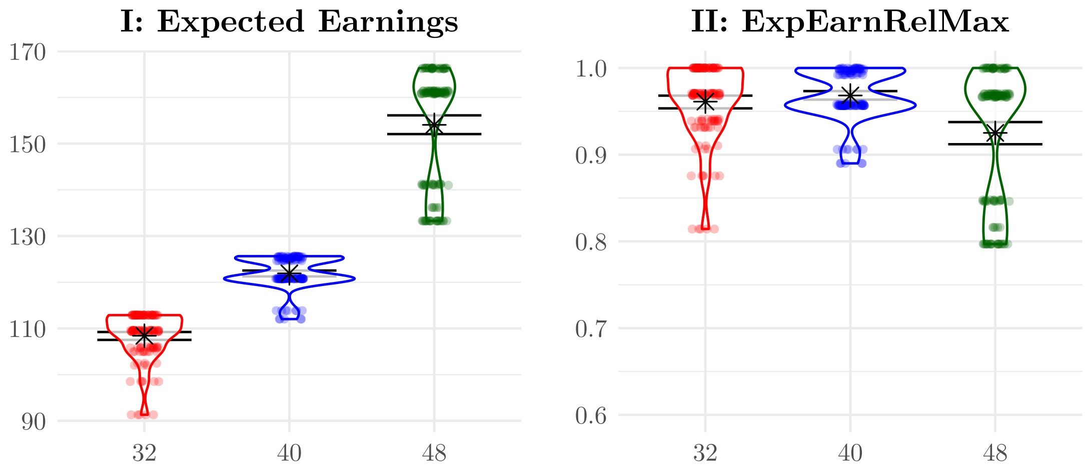

Figure 5: Earnings in Supergame 1: Violin plots

Notes: We define 'Expected earnings' to be the expected earnings in points of a subject's chosen strategy playing against the treatment-level strategy distribution (excluding the subject's own choice); we derive Expected earnings using analytical calculations of payoffs for every possible combination of strategies (see Tables A.5 to A.7 in Web Appendix V). We define 'ExpEarnRel-Max' to be Expected earnings as a proportion of the expected earnings from the best response to the treatment-level strategy distribution (excluding the subject's own choice). In the violin plots, stars are means and horizontal bars are 95 percent confidence intervals, calculated using non-parametric bootstrapping.

We now turn to our analysis of the determinants of earnings. Table 3 reports the results of regressions of expected earnings in Supergame 1 on the variables of interest; throughout, the omitted category is $R = 32$. Confirming Figure 5, the first two rows of Table 3 show that expected earnings increase with the return to joint cooperation. More interestingly, the table tells us that expected earnings depend on both the accuracy of subjects' beliefs and the ability of subjects to choose well given those beliefs.

The third row of Table 3 shows that expected earnings increase with the accuracy of beliefs about the level of cooperation (the final paragraph of Section 3.1 introduces our notion of accuracy).[49] Thus, the quality of subjects' initial beliefs helps to determine how much they earn in the first prisoner's dilemma supergame. Furthermore, the fourth row shows that expected earnings are higher for subjects who good respond to their beliefs (Section 3.2.1 introduces our notion of good responding). Thus, the ability of subjects to select strategies that perform well given their beliefs also helps to determine how much

[49]To help interpret the effect size, note that our measure of accuracy is defined from -1 to 0. In footnote 43 we define accuracy to be the negative of the absolute value of OptimismRelTruth; Panel IV of Figure 4 shows the distribution of OptimismRelTruth.

they earn; this relationship between good responding and earnings is not immediate, since good responding is defined relative to subjects' beliefs, while earnings depend on the actual choices of others. Table A.8 in Web Appendix V shows that our results are robust when we replace our binary measure of good responding with a continuous measure of the proximity of a subject's strategy to the best response to her beliefs (based on expected payoffs relative to those from best responding).

	(1)	(2)	(3)
R40	13.02***	12.82***	12.61***
	(0.55)	(0.53)	(0.53)
R48	44.57***	46.59***	45.52***
	(1.18)	(1.03)	(1.07)
Accuracy of beliefs	12.09***		10.26***
	(2.66)		(2.68)
Good responder to beliefs		7.11***	6.86***
		(0.72)	(0.72)
Mean of dependent variable	128.15	128.15	128.15
N	390	390	390

Table 3: Expected earnings in Supergame 1

Notes: Each column reports a linear OLS regression of expected earnings in Supergame 1, controlling for the five personality factors, demographic characteristics and standardized cognitive ability (see Section 2.9), and with $R = 32$ as the omitted category. Expected earnings is defined in the notes to Figure 5. Accuracy of beliefs is defined in footnote 43. Good responding is a binary variable defined in the second paragraph of Section 3.2.1. $N = 390$ because four subjects did not complete the demographic questionnaire. Heteroskedasticity-robust standard errors are shown in parentheses. ***, ** and * denote significance at the 1 percent, 5 percent and 10 percent levels (two-sided tests).

4 Evolution of beliefs and behavior

Next, we turn to the evolution of beliefs and behavior over the course of the 25 supergames.

4.1 Evolution of cooperation

Figure 6 describes the evolution of cooperation at the treatment and session levels. The figure shows that the differences across treatments in the level of cooperation that we found in Supergame 1 persist over the 25 supergames as subjects gain experience (Tables 4 and 5 below evidence strongly statistically significant effects of the treatment on cooperation).[50] Furthermore, in each treatment, the aggregate level of cooperation is fairly stable over the 25 supergames. The figure also emphasizes a considerable degree of across-session heterogeneity in cooperation. Although the average level of cooperation remains broadly stable, at the individual level we find that subjects' beliefs and behavior respond to experience; the next section describes these results.

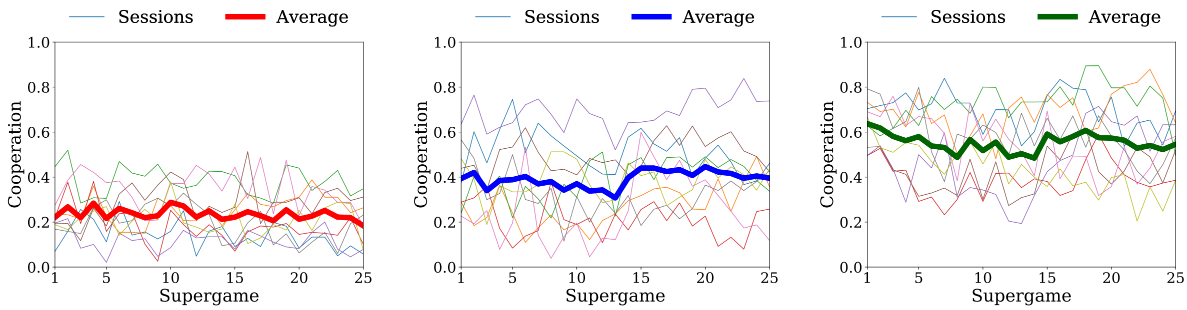

(a) $R = 32$ (b) $R = 40$ (c) $R = 48$

Figure 6: Evolution of cooperation across supergames

Notes: 'Cooperation' is defined as in the notes to Figure 4 for Supergame 1, except that we now measure the cooperation of a subject's strategy in a given supergame using the session-level strategy distribution in that supergame (again excluding the subject's own choice) instead of the treatment-level distribution.

4.2 Learning from experience

When studying the effect of experience, the previous literature focuses on behavior in the first round of each supergame, and finds that subjects are more likely to cooperate in the first round of a supergame when the previous supergame lasted longer and when their opponent in the previous supergame cooperated in the first round of that earlier supergame (Dal Bó and Fréchette, 2011, 2018). We start by showing that this previous finding replicates in our setting with strategy elicitation over consecutive supergames, and we extend the existing literature by showing that the effects of experience on first-round cooperation are robust when we control for initial beliefs.

[50] As we noted in Section 3.1, in Supergame 1 our subjects choose unfriendly strategies at broadly similar rates to Dal Bó and Fréchette (2019)'s subjects in their menu elicitation treatment (see footnote 42).

We then go beyond this analysis of first-round behavior in three ways. First, because we elicited strategies in consecutive supergames, we are able to show that experience affects cooperation at the level of the whole supergame strategy, and we also use transitions between strategies to help understand how experience changes cooperation. Second, because we elicited beliefs in the first supergame, we can show that cooperation depends on both experience and initial beliefs. Third, because we also elicited beliefs in the final supergame, we can show that beliefs themselves respond to experience and that beliefs in the final supergame predict cooperation. Thus, we use beliefs to provide new evidence about the mechanism that underlies learning from experience.

	(1)	(2)	(3)	(4)	(5)
R40	0.152** (0.060)	0.138** (0.055)	0.067 (0.048)	0.084 (0.056)	0.040 (0.043)
R48	0.304*** (0.044)	0.277*** (0.039)	0.153*** (0.043)	0.203*** (0.045)	0.113*** (0.038)
Length of Supergame $t-1$	0.007*** (0.001)				0.007*** (0.001)
Other's Round 1 coop in Supergame $t-1$		0.087*** (0.019)			0.080*** (0.016)
Own Round 1 coop in Supergame 1			0.387*** (0.036)		0.299*** (0.035)
Own optimism in Supergame 1				0.696*** (0.071)	0.345*** (0.061)
Mean of dependent variable	0.464	0.464	0.464	0.464	0.464
N	9360	9360	9360	9360	9360

Table 4: Round 1 cooperation in Supergame t

Notes: Each column reports a linear OLS regression of Round 1 cooperation in Supergame t (for $t > 1$), controlling for the five personality factors, demographic characteristics and standardized cognitive ability (see Section 2.9), and the supergame number, and with $R = 32$ as the omitted category. 'Round 1 coop in Supergame t' is a variable taking value 1 if the relevant subject cooperated in the first round of Supergame t, and taking value 0 if not, where the cooperation decision was determined by the subject's chosen strategy. 'Optimism in Supergame t' is the optimism of the relevant subject's beliefs in Supergame t (optimism is defined in the notes to Figure 4). 'Other' refers to the subject's opponent in a given supergame. 'Own' refers to the subject herself. 'Length of Supergame t' is the number of rounds that Supergame t lasted for. N is in multiples of 390 because four subjects did not complete the demographic questionnaire. Heteroskedasticity-robust standard errors with clustering at the session level are shown in parentheses. ***, ** and * denote significance at the 1 percent, 5 percent and 10 percent levels (two-sided tests).

Table 4 shows that subjects are more likely to choose a strategy that cooperates in the first round of a supergame when the previous supergame was longer and when the strategy chosen by the subject's opponent in the previous supergame cooperated in the first round of that earlier supergame. Using our data from strategy elicitation, we find effect sizes of experience that are close to those found by Dal Bó and Fréchette (2018, Table 9) using meta-data from round-by-round choices. Following Dal Bó and Fréchette (2018), our analysis of learning controls for first-round cooperation in the first supergame (shown in the panel) and the supergame number (not shown in the panel). Extending the previous literature, our analysis controls for initial beliefs by including the subject's optimism in the first supergame (recall from Section 3.1 that optimism measures how often a subject expects others to cooperate). We find that optimism in the first supergame strongly predicts first-round cooperation in later supergames, even after controlling for the subject's first-round cooperation in the first supergame. At the same time, our other parameter estimates are broadly robust to including optimism as a control for initial beliefs. For simplicity, we use linear regressions to estimate parameters; Table A.9 in Web Appendix V shows that our results are robust when instead we use Probit regressions.

Panel I of Table 5 moves beyond this analysis of first-round cooperation by measuring cooperation at the level of the whole supergame strategy, which we call 'strategy cooperation' for short. Here we want a measure of cooperation that does not depend on the behavior of others, and so we measure strategy cooperation by how much a strategy cooperates on average against a uniform distribution over the ten available strategies.[51] Once again, we control for the supergame number (not shown in the panel). Panel I shows clear evidence of learning at the level of the whole supergame strategy. In particular, the panel shows that strategy cooperation increases in both the length of the previous supergame and in the cooperation of the whole supergame strategy chosen by the opponent in the previous supergame.[52] In Section 4.3 we use transitions between strategies to help un-

[51]In particular, here we cannot measure the cooperation of a strategy in a given supergame by its realized cooperation rate; if we did, then we would create a confounding positive correlation between a subject's cooperation in Supergame t and her opponent's cooperation in Supergame $t-1$ driven by the fact that subjects' own propensity to cooperate would influence the measurement of the cooperation of their opponent's strategy. Similarly, here we do not measure the cooperation of a strategy using the session-level strategy distribution because of session-level heterogeneity in the level of cooperation exhibited in Figure 6, which would influence the measurement of subjects' own cooperation and that of their opponent.

[52]Of course, subjects do not directly observe the strategy chosen by their opponent: Table A.10 in Web Appendix V shows that the response to the opponent's cooperation is robust

I: Strategy cooperation in Supergame t

	(1)	(2)	(3)	(4)	(5)
R40	0.103**	0.096**	0.059*	0.049	0.039
	(0.041)	(0.038)	(0.032)	(0.037)	(0.029)
R48	0.209***	0.194***	0.090***	0.128***	0.068**
	(0.033)	(0.030)	(0.029)	(0.029)	(0.026)
Length of Supergame $t-1$	0.004***				0.004***
	(0.001)				(0.001)
Other's strategy coop in Supergame $t-1$		0.072***			0.066***
		(0.018)			(0.014)
Own strategy coop in Supergame 1			0.422***		0.334***
			(0.034)		(0.036)
Own optimism in Supergame 1				0.556***	0.232***
				(0.048)	(0.039)
Mean of dependent variable	0.467	0.467	0.467	0.467	0.467
N	9360	9360	9360	9360	9360

II: Optimism in Supergame 25

	(1)	(2)	(3)	(4)	(5)
R40	0.110**	0.025	0.086*	0.062	-0.016
	(0.045)	(0.025)	(0.042)	(0.043)	(0.024)
R48	0.244***	0.069**	0.178***	0.171***	0.005
	(0.043)	(0.034)	(0.045)	(0.041)	(0.035)
Length of Supergames 1 to 24	0.040				0.025*
	(0.024)				(0.013)
Others' strategy coop in Supergames 1 to 24		0.830***			0.760***
		(0.110)			(0.097)
Own strategy coop in Supergame 1			0.232***		0.069**
			(0.036)		(0.033)
Own optimism in Supergame 1				0.493***	0.411***
				(0.042)	(0.040)
Mean of dependent variable	0.430	0.430	0.430	0.430	0.430
N	390	390	390	390	390

Table 5: Effect of experience on behavior and beliefs

Notes: Each column reports a linear OLS regression of the variable in the panel title (for $t > 1$ in the case of Panel I), controlling for the five personality factors, demographic characteristics and standardized cognitive ability (see Section 2.9), and with $R = 32$ as the omitted category. The regressions in Panel I also include the supergame number as a control. 'Strategy coop in Supergame t' measures the expected cooperation rate of the relevant subject's chosen strategy in Supergame t playing against the uniform distribution over the ten strategies (the notes to Figure 4 define 'cooperation rate'), while 'Strategy coop in Supergames 1 to 24' is the mean over the first 24 supergames. 'Optimism in Supergame t' is the optimism of the relevant subject's beliefs in Supergame t (optimism is defined in the notes to Figure 4). 'Other' refers to the subject's opponent in a given supergame. 'Own' refers to the subject herself. 'Length of Supergame t' is the number of rounds that Supergame t lasted for, while 'Length of Supergames 1 to 24' is the mean over the first 24 supergames. N is in multiples of 390 because four subjects did not complete the demographic questionnaire. Heteroskedasticity-robust standard errors with clustering at the session level are shown in parentheses. ***, ** and * denote significance at the 1 percent, 5 percent and 10 percent levels (two-sided tests).

derstand how experience changes cooperation: for example, when a subject's opponent in the previous supergame cooperated more, the subject is much less likely to change to an unfriendly strategy (see Table A.1 and the associated discussion in Web Appendix III). Panel I further shows that optimism in the first supergame predicts strategy cooperation in later supergames, even after controlling for the subject's strategy cooperation in the first supergame. In summary, Panel I shows that experience and initial beliefs both predict cooperation at the level of the whole supergame strategy.

In Panel II of Table 5 we extend our analysis of learning to the evolution of beliefs themselves. In particular, we study how optimism changes with experience. Since we elicited beliefs only in the first and final supergames, we regress optimism in the final (25th) supergame on the average length of the previous 24 supergames played by the subject and the average strategy cooperation of the subject's 24 previous opponents. Panel II shows that beliefs respond to experience. In particular, the panel shows that optimism in the final supergame increases in the average cooperation of the whole supergame strategies chosen by the subject's previous opponents.[53] When we include both measures of experience in the same regression and control for the subject's strategy cooperation and optimism in the first supergame, the effect of supergame length is marginally statistically significant, but the effect of others' strategy cooperation remains large and strongly statistically significant. In summary, Panel II shows that subjects learn, in the sense that their beliefs respond to what they experienced over the course of the experiment.

Furthermore, the subjects' updated beliefs in the final supergame strongly predict cooperation: a one-unit increase in optimism in Supergame 25 is associated with an increase of 0.625 in strategy cooperation in the same supergame; and updated beliefs account for 23% of the variance in cooperation in the final supergame (this represents 70% of the variance accounted for together by updated beliefs, the return to joint cooperation,

---

to using a model in which subjects make inferences based on realized play. The model with inferences is not our preferred specification because the model needs to make an assumption about prior beliefs and because a subject's own behavior influences her inferences (although the direction of the effect is ambiguous).

[53]Just like for Panel I, Table A.10 in Web Appendix V shows that this response to the opponents' cooperation is robust to using a model in which subjects make inferences based on realized play.

demographics, personality and cognitive ability). By comparison, in the first supergame, initial beliefs account for 36% of the variance in cooperation (this represents 80% of the variance accounted for together by initial beliefs, the treatment, demographics, personality and cognitive ability).[54]

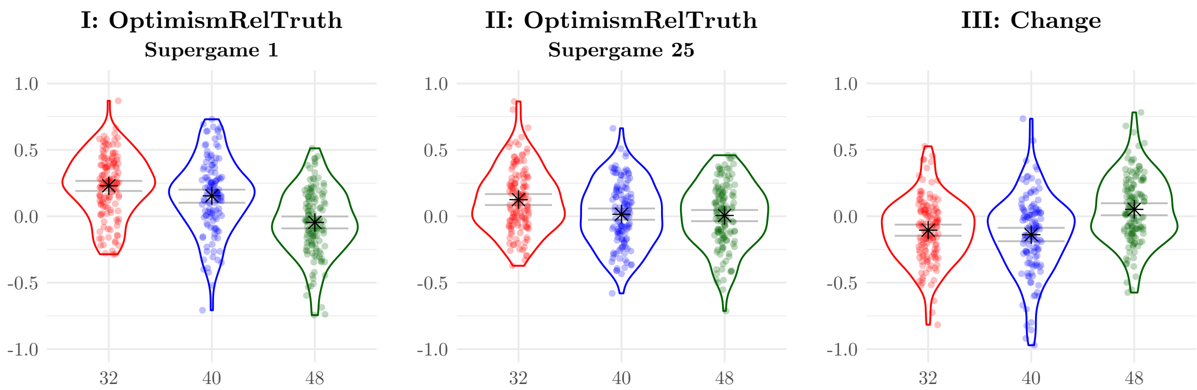

Figure 7: Evolution of OptimismRelTruth at the session level: Violin plots

Notes: 'OptimismRelTruth' is defined as in the notes to Figure 4 for Supergame 1, except that we now use the session-level strategy distribution (again excluding the subject's own choice) instead of the treatment-level distribution. In the violin plots, the unit of observation is an individual subject, stars are means and horizontal bars are 95 percent confidence intervals, calculated using non-parametric bootstrapping.

Recall from Section 3.1 that 'OptimismRelTruth' measures optimism relative to how often others actually cooperate. Our finding from Table 5 that optimism responds to experience suggests that OptimismRelTruth moves toward zero over the course of the 25 supergames as beliefs about how often others actually cooperate become more accurate. Since subjects learn within their session, in Figure 7 we measure OptimismRelTruth at the session level. Figure 7 confirms that, on average, beliefs do indeed move toward the truth in all three treatments. Confirming our finding from Section 3.1, when the return to joint cooperation is low ($R = 32$), subjects' initial beliefs are too optimistic relative to

[54]When we run an OLS regression of strategy cooperation in Supergame 25 on optimism in Supergame 25, controlling for the treatment, the five personality factors, demographic characteristics and standardized cognitive ability (see Section 2.9), the coefficient on optimism is 0.625, with the effect statistically significant at the one-percent level. We decompose the R^2 of this regression using the Shapley value (Huettner and Sunder, 2012; we include the treatment indicators in one group). We then repeat the procedure, using strategy cooperation and optimism in Supergame 1 instead of Supergame 25.

the truth; however, with experience OptimismRelTruth falls toward zero as this excess optimism declines. When the return to joint cooperation is high ($R = 48$), initial beliefs are slightly too pessimistic, and with experience OptimismRelTruth rises toward zero as this modest excess pessimism disappears on average.

4.3 Experimentation and strategy revisions

In this section we delve deeper into the evolution of behavior by studying the factors that drive experimentation and strategy revisions over the course of the 25 supergames. On average, subjects tried four of the ten available strategies at least once; furthermore, 33 percent of the time subjects changed their choice of strategy from one supergame to the next.

To help understand why subjects change their strategy from one supergame to the next, in Table 6 we regress an indicator for changing strategy on the same variables that we analyzed in Panel I of Table 5 when studying learning from experience. To those variables, we add an indicator for good responding to beliefs in the first supergame, which we interpret here as a measure of quality of thinking given the beliefs that the subject has formed (Section 3.2.1 introduces our notion of good responding). We also add the quality of the strategy chosen by the subject in the previous supergame, which we measure by how well that strategy performs in expectation given the subject's experience. In particular, 'Quality of Supergame $t-1$ strategy' is proportional to the earnings of the strategy chosen in the previous supergame when it plays against the distribution of strategies chosen by the subject's opponents up to and including the previous supergame (the notes to Table 6 provide the formal definition).

To summarize the main findings in Table 6, subjects change strategy less frequently: (i) when their opponent cooperated more in the previous supergame; (ii) when the subject exhibits a higher quality of thinking; (iii) when the subject chose a higher quality strategy in the previous supergame; and (iv) when the subject has gained experience by playing more supergames.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
R40	0.014	0.036	0.014	0.009	0.019	0.057*	0.014	0.057*
	(0.032)	(0.032)	(0.033)	(0.033)	(0.032)	(0.031)	(0.032)	(0.032)
R48	-0.026	0.018	-0.028	-0.034	-0.034	0.103**	-0.026	0.076
	(0.035)	(0.035)	(0.037)	(0.034)	(0.035)	(0.044)	(0.035)	(0.046)
Length of Supergame $t-1$	-0.000							0.000
	(0.002)							(0.002)
Other's strategy coop in Supergame $t-1$		-0.209***						-0.176***
		(0.017)						(0.014)
Own strategy coop in Supergame 1			0.007					0.002
			(0.038)					(0.038)
Own optimism in Supergame 1				0.056				0.079
				(0.063)				(0.060)
Good responder to beliefs in Supergame 1					-0.073***			-0.067***
					(0.025)			(0.023)
Quality of Supergame $t-1$ strategy						-0.323***		-0.212***
						(0.060)		(0.054)
Supergame number							-0.005***	-0.005***
							(0.001)	(0.001)
Mean of dependent variable	0.332	0.332	0.332	0.332	0.332	0.332	0.332	0.332
N	9360	9360	9360	9360	9360	9360	9360	9360

Table 6: Strategy changed from Supergame $t-1$ to Supergame t

Notes: Each column reports a linear OLS regression of a binary variable that takes value 1 if the subject changed her strategy from Supergame $t-1$ to Supergame t, and taking value 0 if not, controlling for the five personality factors, demographic characteristics and standardized cognitive ability (see Section 2.9), and with $R = 32$ as the omitted category. 'Length of Supergame $t-1$', 'Other's strategy coop in Supergame $t-1$', 'Own strategy coop in Supergame 1' and 'Own optimism in Supergame 1' appear in Panel I of Table 5 (see the table notes for definitions). Good responding is a binary variable defined in the second paragraph of Section 3.2.1. 'Quality of Supergame $t-1$ strategy' is proportional to the expected earnings of the subject's chosen strategy in Supergame $t-1$ playing against a distribution made up of the $t-1$ strategies chosen by the subject's opponents in Supergames 1 to $t-1$; each unit of quality corresponds to earnings of 100 points, and we derive this measure using analytical calculations of payoffs for every possible combination of strategies (see Tables A.5 to A.7 in Web Appendix V). N is in multiples of 390 because four subjects did not complete the demographic questionnaire. Heteroskedasticity-robust standard errors with clustering at the session level are shown in parentheses. ***, ** and * denote significance at the 1 percent, 5 percent and 10 percent levels (two-sided tests).

In more detail, the third row of Table 6 shows that the likelihood of changing strategy does not depend on the length of the previous supergame. The fourth row shows that the likelihood of changing strategy falls in the cooperation of the whole supergame strategy chosen by the subject's opponent in the previous supergame (we introduced the notion of 'strategy cooperation' in Section 4.2 when discussing the results from Panel I of Table 5). The seventh row shows that subjects who good respond to their beliefs in the first supergame, and so have a higher quality of thinking, are less likely to change strategy from one supergame to the next. The eighth row shows that subjects who chose a higher quality strategy in the previous supergame are less likely to change strategy. In the ninth row the coefficient on the supergame number is negative, and so subjects who have played more supergames tend to change strategy less frequently.[55] Just as in Panel I of Table 5, these results control for the subject's behavior and beliefs in the first supergame (fifth and sixth rows).

In Web Appendix III we further study transitions from one strategy category to another. For example, this analysis shows that subjects are much less likely to change from provocable or lenient strategies to unfriendly ones when their opponent in the previous supergame cooperated more, which sheds light on the mechanism by which the opponent's cooperation in the previous supergame reduces the likelihood of changing strategy (Table 6), while at the same time increasing the subject's own cooperation (Table 5).

5 Personality

In this section we consider the relationship between personality and behavior and beliefs. To do so, we revisit the regressions reported in Table 5 of Section 4.2, which we used to study the effect of experience on cooperation at the level of the whole supergame strategy (strategy cooperation) and on optimism (recall from from Section 3.1 that optimism measures how often a subject expects others to cooperate). Those regressions controlled for personality and demographics, and so we can study the relationship between personality and strategy cooperation and optimism by examining the coefficients on personality from those regressions. Columns 1 and 3 of Table 7 report the coefficients on personality from the regressions reported in Column 5 of Panels I and II of Table 5, which control for experience and for initial behavior and beliefs in Supergame 1, while Column 2 replicates Column 1 using only the last five supergames.

[55]Related work also finds that strategy revisions become less frequent over time (Cason and Mui, 2019, across supergames; Romero and Rosokha, 2019a, within supergames; and Dal Bó and Fréchette, 2019, for non-binding strategies). Cason and Mui (2019) also find that subjects who earn more in one supergame are less likely to change strategy from that supergame to the next.

	(1) Strategy cooperation (Supergames 2-25)	(2) Strategy cooperation (Supergames 21-25)	(3) Optimism (Supergame 25)
Anxiety	-0.004	-0.022	0.003
	(0.009)	(0.014)	(0.010)
Cautiousness	-0.009	-0.010	-0.008
	(0.008)	(0.012)	(0.012)
Kindness	0.008	0.001	0.017*
	(0.011)	(0.015)	(0.008)
Manipulativeness	0.004	0.006	-0.003
	(0.010)	(0.013)	(0.011)
Trust	0.021**	0.022**	0.017**
	(0.009)	(0.009)	(0.007)
Mean of dependent variable	0.467	0.461	0.430
N	9360	1950	390
Control for beliefs in Sup. 1	Yes	Yes	Yes
Control for behavior in Sup. 1	Yes	Yes	Yes
Controls for experience	Yes	Yes	Yes

Table 7: Effect of personality on behavior and beliefs

Notes: The regression reported in Column 1 (Column 3) is exactly the same as the one reported in Column 5 of Panel I (Column 5 of Panel II) of Table 5 (that table and its notes describe all the independent variables); here we report only the coefficients on the five personality factors (see Section 2.9). The regression reported in Column 2 is the same as that reported in Column 1, except that it uses observations of the dependent variable only from the last five supergames. 'Strategy cooperation' is defined in the notes to Table 5 and 'optimism' is defined in the notes to Figure 4.

Column 1 of Table 7 shows that trust predicts cooperation: a one-standard-deviation increase in trust is associated with a two-percentage-point increase in strategy cooperation over the course of the experiment, with the effect statistically significant at the five-percent level. Column 2 shows that the relationship between trust and cooperation persists in the last five supergames. Importantly, Column 3 shows that trust also positively predicts optimism in the final supergame, which suggests that the relationship between trust and cooperation is partly mediated by beliefs about how much others will cooperate (recall from footnote 54 that optimism in the final supergame strongly predicts strategy cooperation in the final supergame).[56] Table A.11 in Web Appendix V shows that the effects of trust on cooperation and optimism are robust when we do not control

[56]Column 3 also shows an effect of kindness on optimism, but kindness does not appear to have a robust effect on behavior and beliefs: there is no corresponding effect of kindness on cooperation in Columns 1 and 2 of Table 7; and the effect on optimism is smaller and not statistically significant at the ten-percent level when we do not control for experience or for

for experience or for initial behavior and beliefs in Supergame 1.[57]

Interestingly, we find no evidence that trust predicts behavior or beliefs in the first supergame, and so our data suggest that trusting subjects learn to cooperate. When we regress strategy cooperation and optimism in Supergame 1 on our personality factors, the coefficients on trust are small and far from statistical significance (see Table A.12 in Web Appendix V).

To understand how trusting subjects learn from experience, we study how trust interacts with good and bad news to drive cooperation. We say that a subject received 'good news' when her opponent's strategy cooperation in the previous supergame was above the treatment-level median, and we say that she received 'bad news' when her opponent's strategy cooperation in the previous supergame was below the median. Recall that Column 5 of Panel I of Table 5 reports the effect of the opponent's strategy cooperation in the previous supergame on a subject's own strategy cooperation; Figure 8 shows that this effect depends on whether the news was good or bad and on the subject's level of trust (the notes to the figure explain how the estimates come from a piece-wise linear spline regression). In particular, Figure 8 shows that subjects high in trust respond more strongly to their opponent's cooperation in the previous supergame when the news was good than when the news was bad; this tendency to amplify good news and discount bad news tends to drive cooperation up. By contrast, subjects low in trust respond more strongly to their opponent's cooperation in the previous supergame when the news was bad than when the news was good; this tendency to discount good news and amplify bad news tends to drive cooperation down.

---

initial behavior and beliefs (see Table A.11 in Web Appendix V).

[57]The effect sizes are similar, the coefficients on trust in Columns 2 and 3 remain significant at the five-percent level, while the coefficient on trust in Column 1 is significant at the ten-percent level ($p = 0.0503$).

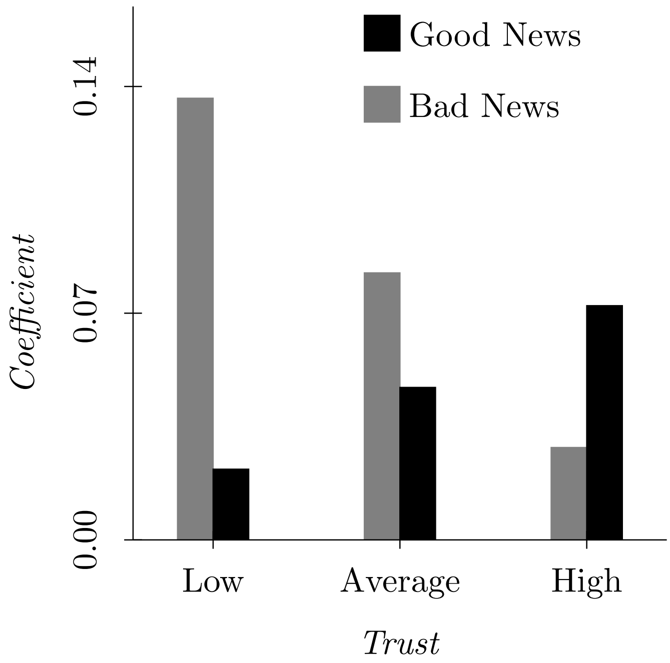

Figure 8: Effect of opponent's strategy cooperation in Supergame $t-1$ on strategy cooperation in Supergame t

Notes: We start with the regression reported in Column 5 of Panel I of Table 5 from Section 4.2, which regresses 'Strategy coop in Supergame t' on 'Other's strategy coop in Supergame $t-1$' and other variables, controlling for the five personality factors, demographic characteristics and standardized cognitive ability (see Section 2.9), and the supergame number; the notes to Table 5 define relevant terms. Let x_R for $R \in \{32, 40, 48\}$ be the treatment-specific median of 'Other's strategy coop in Supergame $t-1$' across all subjects in that treatment and $t > 1$. We say that a subject receives 'good news' ('bad news') when 'Other's strategy coop in Supergame $t-1$' $\geq$ ($<$) x_R. We run a piece-wise linear spline regression by including $\max\{0, (\text{'Other's strategy coop in Supergame } t-1\text{'} - x_R)\}$ in the regression reported in Column 5 of Panel I of Table 5; furthermore, we interact trust with 'Other's strategy coop in Supergame $t-1$' and with $\max\{0, (\text{'Other's strategy coop in Supergame } t-1\text{'} - x_R)\}$. Both interactions are statistically significant at the five-percent level; this can be seen from Column 3 of Table A.13 in Web Appendix V, which reports coefficients from the spline regression and labels $\max\{0, (\text{'Other's strategy coop in Supergame } t-1\text{'} - x_R)\}$ as ('Other's strategy coop in Sup. $t-1$' $- x_R)_+$ for conciseness. Recalling that the trust factor is standardized by construction, we define high (average) (low) trust as trust $= 1$ ($= 0$) ($= -1$), and we then calculate the coefficients in the figure based on the coefficients from Column 3 of Table A.13 in Web Appendix V. For example, when trust $= 1$, the coefficient for bad news is $0.082 + (1 \times -0.054) = 0.028$ and the coefficient for good news is $(0.082 + (1 \times -0.054)) + (-0.035 + (1 \times 0.079)) = 0.072$.

6 Conclusion

By eliciting beliefs over supergame strategies in the indefinitely repeated prisoner's dilemma, in this paper we have advanced our understanding of how people form and use beliefs in repeated interactions that last for an uncertain length of time. Our analysis suggests that beliefs matter. For example, people form initial beliefs that match actual behavior quite well on average, beliefs are correlated with behavior, initial beliefs have a persistent effect on cooperation, and beliefs change with experience and become more accurate over time. Future research should build on our findings by, for example, establishing whether it is feasible to elicit beliefs over supergame strategies using a larger strategy space, studying how the role of beliefs interacts with the continuation probability, and eliciting second-order beliefs alongside the first-order beliefs that we analyze here.

Even though we selected directed measures of personality that we judged most likely to help explain cooperation in the indefinitely repeated prisoner's dilemma, only trust predicts behavior and beliefs in our data. As we discuss in the introduction, the role of trust that we uncover is consistent with a strategic motivation for cooperation mediated by beliefs. Our results thus suggest that personal traits only predict coordination on cooperative equilibria when the personal trait interacts with strategic motives. Indeed, Proto et al. (2019, 2020) find that cognitive ability predicts cooperation in the indefinitely repeated prisoner's dilemma because more intelligent subjects make fewer errors when implementing strategies, and Dal Bó and Fréchette (2018) conclude from their literature survey that strategic motives predominate in explaining cooperation in indefinitely repeated games when cooperation can be sustained as an equilibrium. Future research could further our understanding of the role of trust by matching trusting subjects together or by designing mechanisms by which people high in trust can identify similar others and self-select to play repeated games together.

References

Aghion, P., **Algan, Y.**, **Cahuc, P.**, and **Shleifer, A.** (2010). Regulation and distrust. *Quarterly Journal of Economics*, 125(3): 1015–1049

Al-Ubaydli, O., **Jones, G.**, and **Weel, J.** (2016). Average player traits as predictors of cooperation in a repeated prisoner's dilemma. *Journal of Behavioral and Experimental Economics*, 64: 50–60

Alempaki, D., **Colman, A.**, **Kölle, F.**, **Loomes, G.**, and **Pulford, B.D.** (2019). Investigating the failure to best respond in experimental games. *CEDEX Discussion Paper 2019-13*

Algan, Y. and **Cahuc, P.** (2010). Inherited trust and growth. *American Economic Review*, 100(5): 2060–92

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

Aoyagi, M., **Fréchette, G.R.**, and **Yuksel, S.** (2020). Beliefs in repeated games. *Mimeo, NYU (dated 07/09/2020)*

Axelrod, R. (1980). Effective choice in the prisoner's dilemma. *Journal of Conflict Resolution*, 24(1): 3–25

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

Blanco, M., **Engelmann, D.**, **Koch, A.K.**, and **Normann, H.T.** (2010). Belief elicitation in experiments: is there a hedging problem? *Experimental Economics*, 13(4): 412–438

Blonski, M., **Ockenfels, P.**, and **Spagnolo, G.** (2011). Equilibrium selection in the repeated prisoner's dilemma: Axiomatic approach and experimental evidence. *American Economic Journal: Microeconomics*, 3(3): 164–92

Brandts, J. and **Charness, G.** (2011). The strategy versus the direct-response method: a first survey of experimental comparisons. *Experimental Economics*, 14(3): 375–398

Bruttel, L. and **Kamecke, U.** (2012). Infinity in the lab. How do people play repeated games? *Theory and Decision*, 72(2): 205–219

Burks, S., **Carpenter, J.**, **Goette, L.**, and **Rustichini, A.** (2009). Cognitive skills affect economic preferences, strategic behavior, and job attachment. *Proceedings of the National Academy of Sciences*, 106(19): 7745–7750

Camera, G. and **Casari, M.** (2009). Cooperation among strangers under the shadow of the future. *American Economic Review*, 99(3): 979–1005

Capraro, V., **Smyth, C.**, **Mylona, K.**, and **Niblo, G.A.** (2014). Benevolent characteristics promote cooperative behaviour among humans. *PLoS ONE*, 9(8): e102881

Carpenter, P., **Just, M.A.**, and **Shell, P.** (1990). What one intelligence test measures: A theoretical account of the processing in the Raven Progressive Matrices test. *Psychological Review*, 97(3): 404–413

Cason, T.N. and **Mui, V.L.** (2019). Individual versus group choices of repeated game strategies: A strategy method approach. *Games and Economic Behavior*, 114: 128–145

Charness, G., **Rustichini, A.**, and **Van de Ven, J.** (2018). Self-confidence and strategic behavior. *Experimental Economics*, 21(1): 72–98

Chaudhuri, A., **Sopher, B.**, and **Strand, P.** (2002). Cooperation in social dilemmas, trust and reciprocity. *Journal of Economic Psychology*, 23(2): 231–249

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

Condon, D.M. and **Revelle, W.** (2014). The International Cognitive Ability Resource: Development and initial validation of a public-domain measure. *Intelligence*, 43: 52–64

Cooper, D.J. and **Kagel, J.H.** (2020). Infinitely repeated prisoner dilemma games: Comparing teams with individuals. *Mimeo, Florida State University*

Costa-Gomes, M., **Crawford, V.P.**, and **Broseta, B.** (2001). Cognition and behavior in normal-form games: An experimental study. *Econometrica*, 69(5): 1193–1235

Costa-Gomes, M.A. and **Weizsäcker, G.** (2008). Stated beliefs and play in normal-form games. *Review of Economic Studies*, 75(3): 729–762

Croson, R.T. (2000). Thinking like a game theorist: factors affecting the frequency of equilibrium play. *Journal of Economic Behavior & Organization*, 41(3): 299–314

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

Dal Bó, P. and **Fréchette, G.R.** (2011). The evolution of cooperation in infinitely repeated games: Experimental evidence. *American Economic Review*, 101(1): 411–29

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

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

Davis, D., **Ivanov, A.**, and **Korenok, O.** (2016). Individual characteristics and behavior in repeated games: An experimental study. *Experimental Economics*, 19(1): 67–99

DeYoung, C.G., **Hirsh, J.B.**, **Shane, M.S.**, **Papademetris, X.**, **Rajeevan, N.**, and **Gray, J.R.** (2010). Testing predictions from personality neuroscience: Brain structure and the Big Five. *Psychological Science*, 21(6): 820–828

Dreber, A., **Fudenberg, D.**, and **Rand, D.G.** (2014). Who cooperates in repeated games: The role of altruism, inequity aversion, and demographics. *Journal of Economic Behavior & Organization*, 98: 41–55

Duffy, J. and **Fehr, D.** (2018). Equilibrium selection in similar repeated games: experimental evidence on the role of precedents. *Experimental Economics*, 21(3): 573–600

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

Fe, E., **Gill, D.**, and **Prowse, V.** (2019). Cognitive skills, strategic sophistication, and life outcomes. *CAGE Working Paper 448*

Friedman, D. and **Oprea, R.** (2012). A continuous dilemma. *American Economic Review*, 102(1): 337–363

Fudenberg, D., **Rand, D.G.**, and **Dreber, A.** (2012). Slow to anger and fast to forgive: Cooperation in an uncertain world. *American Economic Review*, 102(2): 720–49

Gee, L.K. and **Schreck, M.J.** (2018). Do beliefs about peers matter for donation matching? Experiments in the field and laboratory. *Games and Economic Behavior*, 107: 282–297

Gill, D. and **Prowse, V.** (2016). Cognitive ability, character skills, and learning to play equilibrium: A level-k analysis. *Journal of Political Economy*, 126(4): 1619–1676

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

Haesevoets, T., **Reinders Folmer, C.**, and **Van Hiel, A.** (2015). Cooperation in mixed-motive games: The role of individual differences in selfish and social orientation. *European Journal of Personality*, 29(4): 445–458

Harrison, G.W., **Martínez-Correa, J.**, **Swarthout, J.T.**, and **Ulm, E.R.** (2017). Scoring rules for subjective probability distributions. *Journal of Economic Behavior & Organization*, 134: 430–448

Huettner, F. and **Sunder, M.** (2012). Axiomatic arguments for decomposing goodness of fit according to Shapley and Owen values. *Electronic Journal of Statistics*, 6: 1239–1250

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

John, O.P., **Naumann, L.P.**, and **Soto, C.J.** (2008). Paradigm shift to the integrative Big Five trait taxonomy. *Handbook of Personality: Theory and Research*, 3(2): 114–158

Kagel, J. and **McGee, P.** (2014). Personality and cooperation in finitely repeated prisoner's dilemma games. *Economics Letters*, 124(2): 274–277

Kagel, J.H. and **McGee, P.** (2016). Team versus individual play in finitely repeated prisoner dilemma games. *American Economic Journal: Microeconomics*, 8(2): 253–76

Kaiser, H.F. (1960). The application of electronic computers to factor analysis. *Educational and Psychological Measurement*, 20(1): 141–151

Lopez-de Silanes, F., **La Porta, R.**, **Shleifer, A.**, and **Vishny, R.** (1997). Trust in large organizations. *American Economic Review: Papers & Proceedings*, 87(2): 333–338

Luo, L., **Arizmendi, C.**, and **Gates, K.M.** (2019). Exploratory factor analysis (EFA) programs in R. *Structural Equation Modeling: A Multidisciplinary Journal*, 26(5): 819–826

Parks, C.D., **Henager, R.F.**, and **Scamahorn, S.D.** (1996). Trust and reactions to messages of intent in social dilemmas. *Journal of Conflict Resolution*, 40(1): 134–151

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

Proto, E., **Rustichini, A.**, and **Sofianos, A.** (2020). Intelligence, errors and strategic choices in the repeated prisoners' dilemma. *IZA Discussion Paper 12925*

Raven, J., **Raven, J.C.**, and **Court, J.H.** (2000). *Manual for Raven's Progressive Matrices and Vocabulary Scales.* San Antonio: Pearson

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

Romero, J. and **Rosokha, Y.** (2019a). The evolution of cooperation: The role of costly strategy adjustments. *American Economic Journal: Microeconomics*, 11(1): 299–328

Romero, J. and **Rosokha, Y.** (2019b). Mixed strategies in the indefinitely repeated prisoner's dilemma. *Mimeo, Purdue University*

Schlag, K.H. and **van der Weele, J.J.** (2013). Eliciting probabilities, means, medians, variances and covariances without assuming risk neutrality. *Theoretical Economics Letters*, 3(1): 38–42

Schlag, K.H., **Tremewan, J.**, and **Van der Weele, J.J.** (2015). A penny for your thoughts: A survey of methods for eliciting beliefs. *Experimental Economics*, 18(3): 457–490

Schotter, A. and **Trevino, I.** (2014). Belief elicitation in the laboratory. *Annual Review of Economics*, 6(1): 103–128

Selten, R. (1998). Axiomatic characterization of the quadratic scoring rule. *Experimental Economics*, 1(1): 43–61

Terracol, A. and **Vaksmann, J.** (2009). Dumbing down rational players: Learning and teaching in an experimental game. *Journal of Economic Behavior & Organization*, 70(1-2): 54–71

Zhang, D., **Lin, Y.**, **Jing, Y.**, **Feng, C.**, and **Gu, R.** (2019). The dynamics of belief updating in human cooperation: Findings from inter-brain ERP hyperscanning. *NeuroImage*, 198: 1–12