I Z A Institute of Labor Economics
Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 15492

Beliefs, Learning, and Personality in the Indefinitely Repeated Prisoner's Dilemma

David Gill
Yaroslav Rosokha

AUGUST 2022

IZA Institute of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 15492

Beliefs, Learning, and Personality in the Indefinitely Repeated Prisoner's Dilemma

David Gill
Purdue University and IZA

Yaroslav Rosokha
Purdue University

AUGUST 2022

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Beliefs, Learning, and Personality in the Indefinitely Repeated Prisoner's Dilemma*

We aim to understand the role and evolution of beliefs in the indefinitely repeated prisoner's dilemma (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. Finally, we uncover a novel mechanism whereby trusting subjects learn to cooperate through their interaction with experience.

JEL Classification: C72, C73, C91, D91

Keywords: infinitely repeated prisoner's dilemma, cooperation, optimism, belief elicitation, supergame strategies, experimentation, trust, experiment

Corresponding author:
David Gill
Department of Economics
Purdue University
Rawls 4023
Krannert School of Management
403 W. State Street
West Lafayette, IN 47907
USA
E-mail: dgill.econ@gmail.com

* 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.

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 the role and evolution of beliefs about others when people 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, forthcoming, Kölle et al., 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.

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). 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]

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 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.

---

[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.

We find that heterogeneity in beliefs and changes in beliefs with experience are central to understanding behavior and learning in the indefinitely repeated prisoner's dilemma. In the first supergame, before subjects gain any experience, we find that: (i) average beliefs match the distribution of chosen strategies quite well; (ii) most subjects choose strategies that perform well given their beliefs; and (iii) beliefs strongly predict cooperation. Turning to learning and the evolution of beliefs and behavior, we find that: (i) cooperation depends on both experience and initial beliefs; (ii) beliefs respond to experience and become more accurate over time; and (iii) updated beliefs continue to predict cooperation after subjects have learned from experience. Thus, the elicited beliefs allow us to shed new light on the mechanisms that underlie learning from experience.

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. In particular, subjects high in trust respond more strongly when they experience cooperative behavior, while subjects low in trust respond more strongly when they experience uncooperative behavior. This tendency of high trust subjects to amplify cooperative evidence and discount uncooperative evidence 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 beliefs correlate with cooperation.[6]

[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.

[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

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 in the prisoner's dilemma.[10] Outside of the lab, trust matters for economic outcomes:

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).

[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

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. Our results do not conflict with this view, in the sense that the more optimistic beliefs of more trusting subjects can help to explain why they cooperate more, even absent any social preferences (although reciprocity can also help to explain the link between optimism and cooperation; see Rabin, 1993).

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 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.

(decreasing) cooperation. Parks et al. (1996) also find that the response to messages depends on the consistency of the message and behavior.

[11] Web Appendix II.4 discusses Proto et al. (2019)'s results on personality in more detail. Proto et al. (forthcoming) 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.

2 Experimental design

In this section, we describe our experimental design. Web Appendix III provides further details, while Web Appendix I provides screenshots from the experiment.

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. 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.

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. Sessions lasted just under one hour and included 12, 14, or 16 subjects, with 14.6 subjects on average.[12]

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; 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

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

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.[13]

	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).

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

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.[14] Since we wanted to elicit both strategies and beliefs, we restricted attention to the ten supergame strategies in Table 1.[15]

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.

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. The order that the ten strategies appeared on the subject's screen was randomized across subjects, and remained constant throughout the session.

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.[16] These eight strategies include the five 'key' strategies identified by Dal Bó and Fréchette (2018, Table 10). We added DG and RAND

[14]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).

[15]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.

[16]The data include 32 combinations of paper and parameter values with $\delta > 0$. The most

to the eight strategies from Fudenberg et al. (2012) selected by our SFEM exercise, giving the ten strategies in Table 1. DG is the natural extension of G corresponding to DTFT as the extension of TFT. RAND corresponds to 'RANDOM-50' in Dal Bó and Fréchette (2019).

2.5 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.

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, 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'.

popular strategies were determined by running SFEM for each combination and taking the average.

2.6 Supergames with strategy elicitation

Subjects played 25 supergames, with random rematching between supergames and the same supergame lengths across treatments. 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. 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.[17] Figure 2 provides a screenshot of the experimental interface.

Standard theoretical analyses of repeated games (and the SFEM procedure for estimating 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. Web Appendix III.6 describes how full defection and full cooperation are equilibrium outcomes in our setting with strategy elicitation.

[17]This choice preserves external validity: in real-world strategic interactions, people usually do not directly observe others' strategies. In our implementation, when learning from experience, subjects are uncertain about the strategy chosen by their opponents in previous rounds (because the same within-supergame history of play can arise from multiple strategies). If we removed this uncertainty by allowing subjects to directly observe strategies, we conjecture that subjects would respond more strongly to experience. As a result, the comparisons that we make in Section 4.2 to the effects of learning from experience in Dal Bó and Fréchette (2018)'s meta-data from round-by-round choices would become less informative.

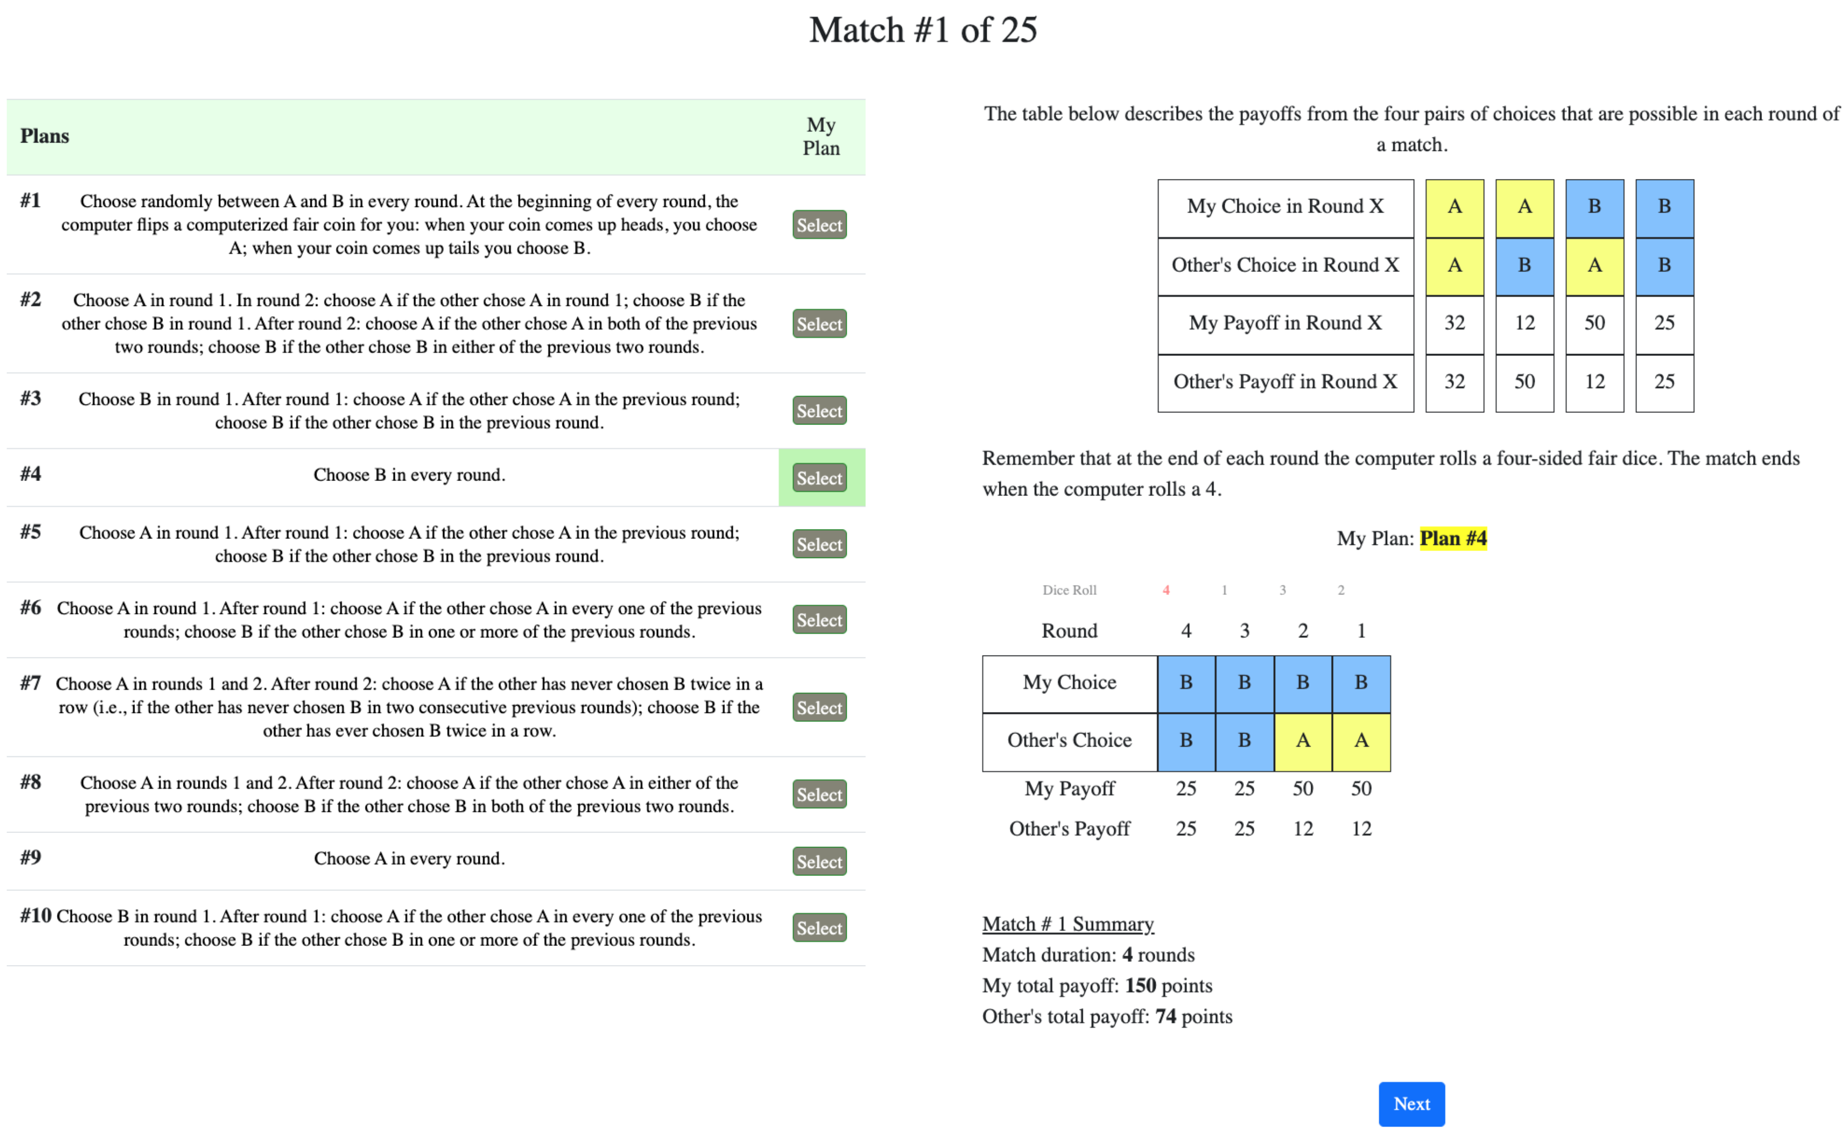

Figure 2: Summary screen (Supergame 1)

Notes: The order that the ten strategies appeared on the subject's screen was randomized across subjects (see Web Appendix III.2).

2.7 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. 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 not know that beliefs would be elicited when they chose their strategy.[18] To minimize contamination of later supergame strategies, we elicited beliefs only twice, in the first supergame and in the final supergame.[19] Furthermore, we gave the subjects no feedback about the accuracy of their guesses.

2.8 Measurement of personal characteristics

2.8.1 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. We gave subjects six minutes to complete the 52-question personality questionnaire.[20] Web Appendix III.8 describes how we selected the personality measures and the sources for the questions, together with further implementation details.[21]

Some of our personality 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

[18]Similarly, we elicited beliefs in the final supergame after subjects had chosen their strategy but before the supergame was played out.

[19]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.

[20]The screen showed a countdown clock. We followed the time per-question recommended for the Big Five Inventory (John et al., 2008, p.137).

[21]We included the personality questionnaire at the beginning of the experiment because we were concerned that experience and earnings in the prisoner's dilemma could affect subjects' answers to the personality questions. Web Appendix III.8 discusses this choice in more detail.

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. The five factors are: anxiety; cautiousness; kindness; manipulativeness; and trust. By construction, the factors are uncorrelated with each other, and have zero mean and unit variance. Web Appendix III.9 describes how we construct the factors.

2.8.2 Cognitive ability test and demographic questionnaire

After measuring personality but before subjects played any games, we measured cognitive ability using a short matrix reasoning test, and at the end of the experiment subjects completed a short demographic questionnaire. Web Appendix III.10 and Web Appendix III.11 provide details. We discuss our results relating to cognitive ability and demographics in Web Appendix VIII (and we discuss there how our results on cognitive ability relate to those of Proto et al., 2019, forthcoming).

3 Initial beliefs and behavior

In this section, we study initial beliefs and behavior in the first supergame. In the main text we focus on our key results; Web Appendix IV provides further analysis, with an outline in Web Appendix IV.1. Table A.21 in Web Appendix X summarizes the main definitions used in the paper.

3.1 Strategy choices and average beliefs

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 that defect in the very first round are popular. As the return to joint cooperation increases (i.e., as R increases from 32 to 48), the proportion of unfriendly strategies falls substantially, while provocable and lenient strategies become more popular.[22] In Supergame 1,

[22]Provocable strategies defect immediately in response to the opponent's first defection, while lenient strategies do not defect immediately in response to defection; see Web Appendix III.3.

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, $\delta = 0.75$.[23]

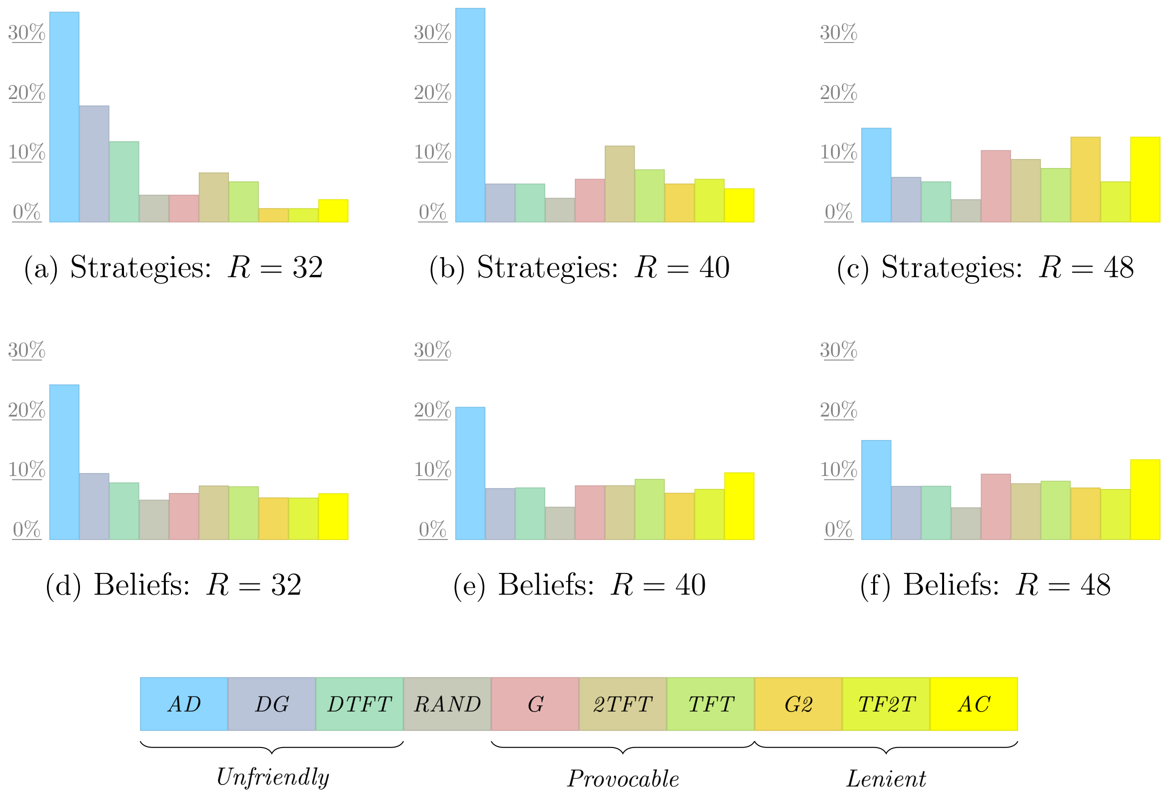

(a) Strategies: $R = 32$ (b) Strategies: $R = 40$ (c) Strategies: $R = 48$

(d) Beliefs: $R = 32$ (e) Beliefs: $R = 40$ (f) Beliefs: $R = 48$

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

[23] We elicit initial beliefs and strategies after training, while Dal Bó and Fréchette (2019) elicit strategies after experience of direct-response play. Even though are settings are not directly comparable, but consistent with evidence that our subjects learn during the training phase (see the third paragraph of Web Appendix VII), in Supergame 1 we find unfriendly strategies at broadly similar rates to Dal Bó and Fréchette (2019): 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).

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 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.

3.2 Optimism and cooperation

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).

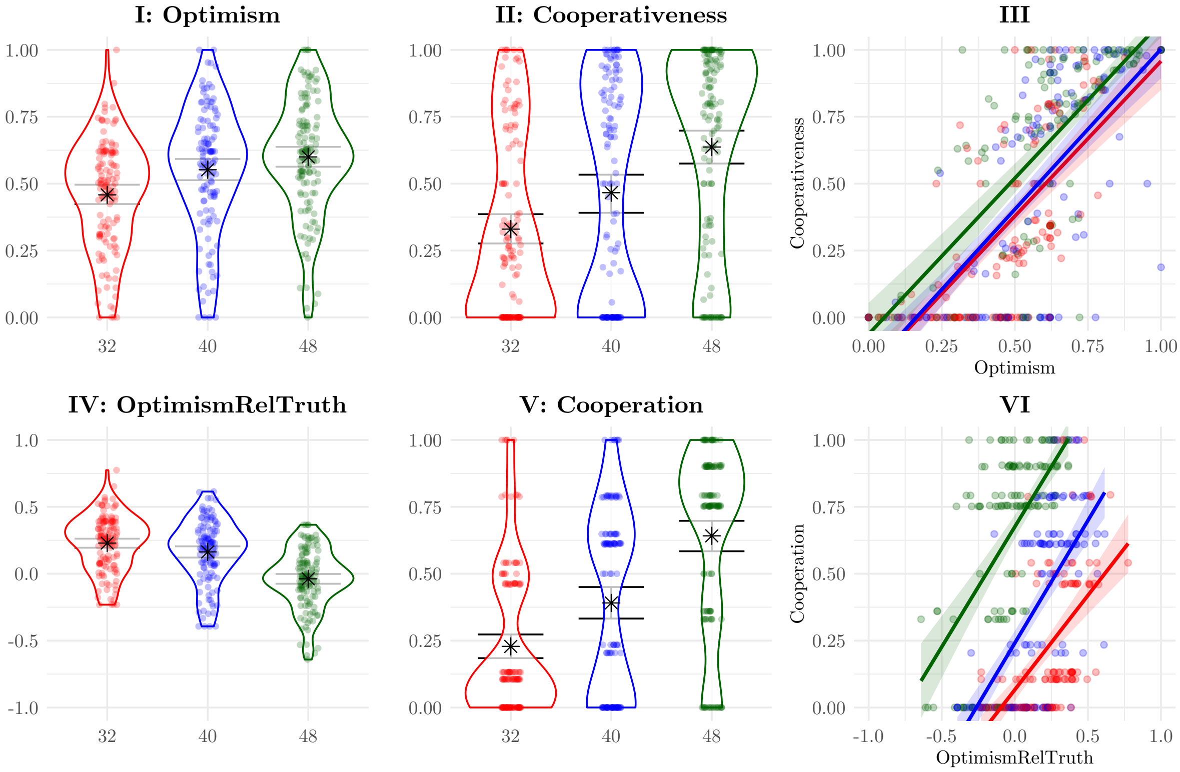

Figure 4: Beliefs and behavior in Supergame 1: Violin and scatter plots

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.11 in Web Appendix X). 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.

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.

Our definition of optimism measures how often a subject expects the population of subjects (excluding herself) to cooperate when they play against each other; we find this intuitive, and furthermore this definition allows a direct comparison of beliefs to the level of cooperation in the population: Web Appendix IV.2 provides further discussion.

3.3 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.

4 Evolution of beliefs and behavior

Next, we turn to the evolution of beliefs and behavior over the course of the 25 supergames. In the main text we focus on our key results; Web Appendix V provides further analysis, with an outline in Web Appendix V.1. For brevity, the regression tables in the main text do not report every coefficient: Table A.9 in Web Appendix VIII reports the full set of estimates.

4.1 Evolution of cooperation

Figure 5 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 2 and 3 below evidence strongly statistically significant effects of the treatment on cooperation).[24] Furthermore, in each treatment, the aggregate level of cooperation is fairly stable over the 25 supergames. Two features of our design potentially explain why the average level of cooperation is broadly stable in our data: first, subjects could learn about the structure of the game and the properties of the supergame strategies during the two forms of training; and second, subjects were unable to experiment within supergame given that we elicited supergame strategies (see Web Appendix VII for details). Although the average level of cooperation remains broadly stable, at the individual level we find that subjects' beliefs and behavior respond to experience, with the effects of experience on behavior consistent with prior work without strategy elicitation; the next section describes these results. Figure 5 also emphasizes a considerable degree of across-session heterogeneity in cooperation.

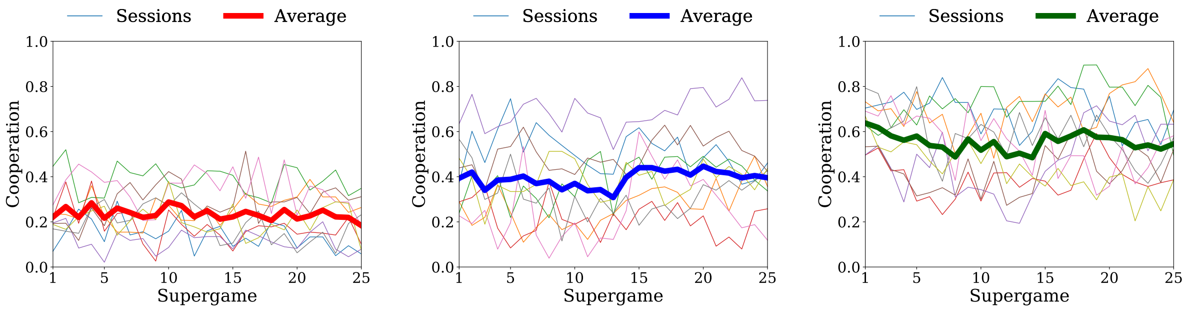

(a) $R = 32$ (b) $R = 40$ (c) $R = 48$

Figure 5: 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.

[24]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 23).

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.

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. 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 become more accurate over time, and that updated beliefs in the final supergame continue to predict cooperation. Thus, we use beliefs to provide new evidence about the mechanism that underlies learning from experience.

Table 2 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 3.2 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.

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

Table 2: 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 2.8), and the supergame number, and with $R = 32$ as the omitted category. For simplicity, we use linear regressions to estimate parameters; Table A.18 in Web Appendix X shows that our results are robust when instead we use Probit regressions. '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).

Panel I of Table 3 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.[25] Once again, we control for the supergame number (not shown in the panel). Panel I shows

[25] 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

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.[26] 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 3 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.[27] 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.

---

the cooperation of a strategy using the session-level strategy distribution because of session-level heterogeneity in the level of cooperation exhibited in Figure 5, which would influence the measurement of subjects' own cooperation and that of their opponent.

[26]Of course, subjects do not directly observe the strategy chosen by their opponent: Table A.19 in Web Appendix X shows that the response to the opponent's cooperation is robust 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).

[27]Just like for Panel I, Table A.19 in Web Appendix X shows that this response to the opponents' cooperation is robust to using a model in which subjects make inferences based on realized play.

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 3: 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 2.8), and with $R = 32$ as the omitted category (Table A.20 in Web Appendix X shows that the results are robust when we do not control for personal traits or demographics). 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).

As subjects learn over the course of the 25 supergames, their beliefs become more accurate (see Web Appendix V.2). 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, 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).[28]

5 Personality

In this section we consider the relationship between personality and behavior and beliefs. To summarize, we find that more trusting subjects cooperate more on average and hold more optimistic beliefs by the end of the experiment. Interestingly, we find no evidence that trust predicts behavior or beliefs in the first supergame. Instead, trusting subjects learn to cooperate through their interaction with experience. Finally, we find no systematic evidence that our other personality measures predict behavior or beliefs, whether in the first or later supergames.

To study the relationship between personality and behavior and beliefs, we revisit the regressions reported in Table 3 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 3.2 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 strat-

[28]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 2.8), 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.

egy cooperation and optimism by examining the coefficients on personality from those regressions. Columns 1 and 3 of Table 4 report the coefficients on personality from the regressions reported in Column 5 of Panels I and II of Table 3, 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.

	(1) Strategy cooperation (Supergames 2-25)	(2) Strategy cooperation (Supergames 21-25)	(3) Optimism (Supergame 25)
Anxiety	-0.004 (0.009)	-0.022 (0.014)	0.003 (0.010)
Cautiousness	-0.009 (0.008)	-0.010 (0.012)	-0.008 (0.012)
Kindness	0.008 (0.011)	0.001 (0.015)	0.017* (0.008)
Manipulativeness	0.004 (0.010)	0.006 (0.013)	-0.003 (0.011)
Trust	0.021** (0.009)	0.022** (0.009)	0.017** (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 4: 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 3 (that table and its notes describe all the independent variables); here we report only the coefficients on the five personality factors (see Section 2.8). 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 3 and 'optimism' is defined in the notes to Figure 4.

Column 1 of Table 4 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 28 that optimism in the final supergame strongly predicts strategy cooperation in the final supergame).[29]

Table A.5 in Web Appendix VI.1 shows that the effects of trust on cooperation and optimism reported in Table 4 are robust when we do not control for experience or for initial behavior and beliefs in Supergame 1. 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 ($p = 0.0503$).

Interestingly, and consistent with the robustness discussed in the previous paragraph, 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.6 in Web Appendix VI.1). When we regress strategy cooperation on our personality factors using the data from all supergames, and interact trust with an indicator for Supergame 1, we find that this interaction of trust with the Supergame 1 indicator is negative and statistically significant at the 5% level, which provides evidence that the difference in the effect of trust on cooperation in the first supergame compared to later supergames is statistically meaningful.[30]

To understand how trusting subjects learn from experience, we study how trust interacts with cooperative and uncooperative evidence to drive cooperation. We say that a subject received 'cooperative evidence' when her opponent's strategy cooperation in the previous supergame was above the treatment-level median, and we say that she received 'uncooperative evidence' when her opponent's strategy cooperation in the previous supergame was below the median. Recall that Column 5 of Panel I of Table 3 reports the

[29]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 4; 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 initial behavior and beliefs (see Table A.5 in Web Appendix VI.1).

[30]Specifically, we take the regressions reported in Columns 1 and 2 of Table A.6 in Web Appendix VI.1, and we run these regressions using the data from all supergames, further including an indicator for Supergame 1, the interaction of this indicator with trust, and a control for the supergame number. In both cases, $p < 0.05$ for the interaction of trust with the Supergame 1 indicator (based on heteroskedasticity-robust standard errors with clustering at the session level and two-sided tests of significance).

effect of the opponent's strategy cooperation in the previous supergame on a subject's own strategy cooperation; Figure 6 shows that this effect depends on whether the evidence was cooperative or uncooperative 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 6 shows that subjects high in trust respond more strongly to their opponent's cooperation in the previous supergame when the evidence was cooperative than when the evidence was uncooperative; this tendency to amplify cooperative evidence and discount uncooperative evidence 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 evidence was uncooperative than when the evidence was cooperative; this tendency to discount cooperative evidence and amplify uncooperative evidence tends to drive cooperation down.

To understand better the coefficients reported in Figure 6, Figure 7 describes visually the meaning of the 'uncooperative evidence' and 'cooperative evidence' coefficients. As the opponent's strategy cooperation in the previous supergame falls below the treatment-level median x_R, a subject's own cooperation falls at a rate given by the slope of the function to the left of the kink at x_R: this slope is the 'uncooperative evidence' coefficient. As the opponent's strategy cooperation in the previous supergame rises above the treatment-level median x_R, a subject's own cooperation rises at a rate given by the slope of the function to the right of the kink at x_R: this slope is the 'cooperative evidence' coefficient. The coefficients from Figure 6 therefore tell us that low trust subjects respond more strongly to uncooperative evidence than to cooperative evidence, since those coefficients mean that for low trust subjects the function is steeper to the left of the kink than it is to the right of the kink. By contrast, the coefficients from Figure 6 tell us that high trust subjects respond more strongly to cooperative evidence than to uncooperative evidence, since those coefficients mean that for high trust subjects the function is steeper to the right of the kink than it is to the left of the kink. The notes to Figure 6 describe the formal details of the piece-wise linear spline regression with a kink at x_R that produces the 'cooperative evidence' and 'uncooperative evidence' coefficients.[31]

[31] Table A.7 in Web Appendix VI.1 reports the results of the spline regression. In that regression, the slope of the function to the left of the kink is given by the coefficient on 'Other's strategy coop in Supergame $t-1$', while the slope of the function to the right of the kink is given by the sum of the coefficients on 'Other's strategy coop in Supergame $t-1$' and

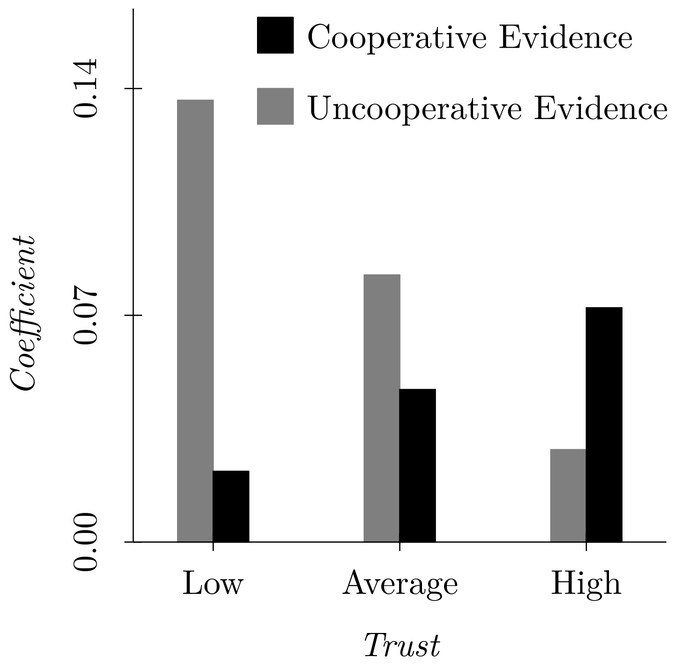

Figure 6: Effect of opponent's strategy cooperation in Supergame $t-1$ on strategy cooperation in Supergame t

Notes: We start with the regression reported in Column 5 of Panel I of Table 3 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 2.8), and the supergame number; the notes to Table 3 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 'cooperative evidence' ('uncooperative evidence') 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 3; 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.7 in Web Appendix VI.1, 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.7 in Web Appendix VI.1. For example, when trust $= 1$, the coefficient for uncooperative evidence is $0.082 + (1 \times -0.054) = 0.028$ and the coefficient for cooperative evidence is $(0.082 + (1 \times -0.054)) + (-0.035 + (1 \times 0.079)) = 0.072$.

('Other's strategy coop in Sup. $t-1$' $- x_R)_+$.

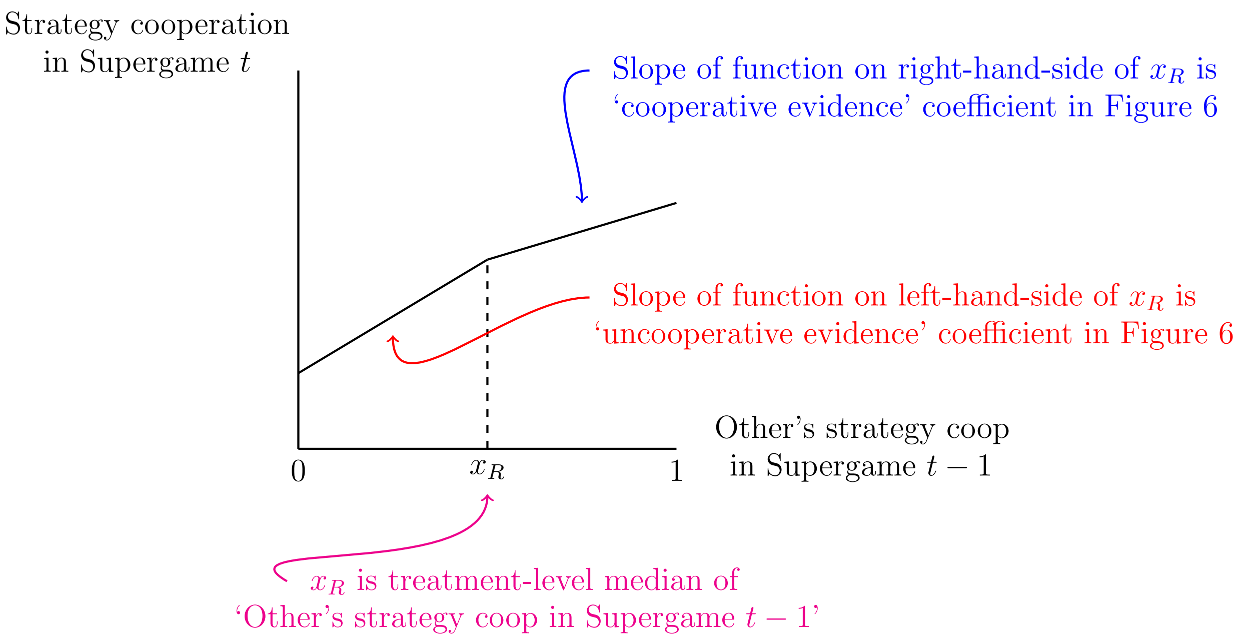

Figure 7: Understanding the coefficients from Figure 6

In the web appendix we consider two extensions. In Web Appendix VI.2: (i) we show that the pattern in Figure 6 is robust when we simplify the underlying regression by excluding the control variables; and (ii) we find no statistically significant differences in the pattern across treatments. In Web Appendix VI.3, we carry out counter-factual simulations investigating the impact of being able to split subjects across sessions based on their standardized measure of trust. In the simulations, we find that the asymmetric reaction to cooperative and uncooperative evidence by trusting and non-trusting agents yields substantially different cooperation dynamics – with sessions populated by trusting agents achieving high cooperation and sessions populated by non-trusting agents achieving low cooperation. Our experimental subjects meet a mix of high and low trust opponents, and so our experimental data hide these interesting dynamics that arise from the interaction of agents with similar trust levels.

6 Conclusion

By eliciting beliefs over supergame strategies in the indefinitely repeated prisoner's dilemma, in this paper we have advanced our understanding of the role and evolution of beliefs in repeated interactions that last for an uncertain length of time. Our analysis suggests that beliefs matter. For example, people's initial beliefs 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. Proto et al. (2019, forthcoming) 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. As we note in the introduction, our results do not conflict with Dal Bó and Fréchette (2018)'s conclusion, in the sense that the more optimistic beliefs of higher trust subjects can help to explain why they cooperate more, even in the absence of any social preferences (although reciprocity can also help to explain the link between optimism and cooperation; see Rabin, 1993). 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 (for main text and Web Appendix)

Acedo-Carmona, C. and **Gomila, A.** (2014). Personal trust increases cooperation beyond general trust. *PloS ONE*, 9(8): e105559

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

Ahn, T.K., **Ostrom, E.**, **Schmidt, D.**, and **Walker, J.** (2003). Trust in two-person games: Game structures and linkages. In E. Ostrom and J. Walker, editors, *Trust and Reciprocity: Interdisciplinary Lessons from Experimental Research*, 323–51. Russell Sage

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

Anderson, L.R., **Mellor, J.M.**, and **Milyo, J.** (2004). Social capital and contributions in a public-goods experiment. *American Economic Review: Papers & Proceedings*, 94(2): 373–376

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., **Rigotti, L.**, and **Rustichini, A.** (2016). Social surplus determines cooperation rates in the one-shot prisoner's dilemma. *Games and Economic Behavior*, 100: 113–124

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

Duffy, J. and **Ochs, J.** (2009). Cooperative behavior and the frequency of social interaction. *Games and Economic Behavior*, 66(2): 785–812

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

Emonds, G., **Declerck, C.H.**, **Boone, C.**, **Seurinck, R.**, and **Achten, R.** (2014). Establishing cooperation in a mixed-motive social dilemma. an fMRI study investigating the role of social value orientation and dispositional trust. *Social Neuroscience*, 9(1): 10–22

Engel, C. and **Zhurakhovska, L.** (2016). When is the risk of cooperation worth taking? The prisoner's dilemma as a game of multiple motives. *Applied Economics Letters*, 23(16): 1157–1161

Fe, E., **Gill, D.**, and **Prowse, V.** (forthcoming). Cognitive skills, strategic sophistication, and life outcomes. *Journal of Political Economy*

Fréchette, G.R., **Schotter, A.**, and **Trevino, I.** (2017). Personality, information acquisition, and choice under uncertainty: An experimental study. *Economic Inquiry*, 55(3): 1468–1488

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

Gächter, S., **Herrmann, B.**, and **Thöni, C.** (2004). Trust, voluntary cooperation, and socio-economic background: survey and experimental evidence. *Journal of Economic Behavior & Organization*, 55(4): 505–531

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

Goldberg, L.R. (1999). A broad-bandwidth, public domain, personality inventory measuring the lower-level facets of several five-factor models. In I. Mervielde, I. Deary, F. De Fruyt, and F. Ostendorf, editors, *Personality Psychology in Europe*, volume 7(1), 7–28. Tilburg University Press

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.**, **Bostyn, D.H.**, and **Van Hiel, A.** (2018). Behavioural consistency within the prisoner's dilemma game: The role of personality and situation. *European Journal of Personality*, 32(4): 405–426

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

Heuer, L. and **Orland, A.** (2019). Cooperation in the Prisoner's Dilemma: an experimental comparison between pure and mixed strategies. *Royal Society Open Science*, 6: 182142

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

Johnson, J.A. (2014). Measuring thirty facets of the Five Factor Model with a 120-item public domain inventory: Development of the IPIP-NEO-120. *Journal of Research in Personality*, 51: 78–89

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

Kölle, F., **Quercia, S.**, and **Tripodi, E.** (2020). Social preferences under the shadow of the future. *SSRN Working Paper 3622125*

Li, S.X. and **Liu, T.X.** (2017). Group identity and cooperation in infinitely repeated games. *Mimeo, Texas A&M (dated 01/15/2017)*

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

McGrath, R. (2017). The VIA assessment suite for adults: Development and initial evaluation. *Technical Report, VIA Institute on Character*

Messé, L.A. and **Sivacek, J.M.** (1979). Predictions of others' responses in a mixed-motive game: Self-justification or false consensus? *Journal of Personality and Social Psychology*, 37(4): 602–607

Miettinen, T. and **Suetens, S.** (2008). Communication and guilt in a prisoner's dilemma. *Journal of Conflict Resolution*, 52(6): 945–960

Mulder, L.B., **Van Dijk, E.**, **De Cremer, D.**, and **Wilke, H.A.** (2006). Undermining trust and cooperation: The paradox of sanctioning systems in social dilemmas. *Journal of Experimental Social Psychology*, 42(2): 147–162

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

Peeters, R. and **Vorsatz, M.** (2018). Simple guilt and cooperation. *University of Otago Economics Discussion Paper 1801*

Peterson, C. and **Seligman, M.E.** (2004). *Character Strengths and Virtues: A Handbook and Classification.* Oxford University Press

Peysakhovich, A., **Nowak, M.A.**, and **Rand, D.G.** (2014). Humans display a 'cooperative phenotype' that is domain general and temporally stable. *Nature Communications*, 5(1): 1–8

Proto, E. and **Rustichini, A.** (2014). Cooperation and personality. *Warwick Economic Research Paper 1045*

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.** (forthcoming). Intelligence, errors and cooperation in repeated interactions. *Review of Economic Studies*

Rabin, M. (1993). Incorporating fairness into game theory and economics. *American Economic Review*, 83: 1281–1302

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

Ridinger, G. and **McBride, M.** (2017). Theory-of-mind ability and cooperation. *Mimeo, UC Irvine (dated 09/2017)*

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*

Sato, K. (1988). Trust and group size in a social dilemma. *Japanese Psychological Research*, 30(2): 88–93

Sato, K. (1989). Trust and feedback in a social dilemma. *Japanese Journal of Experimental Social Psychology*, 29(2): 123–128

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

Simms, L.J., **Goldberg, L.R.**, **Roberts, J.E.**, **Watson, D.**, **Welte, J.**, and **Rotterman, J.H.** (2011). Computerized adaptive assessment of personality disorder: Introducing the CAT–PD project. *Journal of Personality Assessment*, 93(4): 380–389

Sutter, M. and **Untertrifaller, A.** (2020). Children's heterogeneity in cooperation and parental background: An experimental study. *Journal of Economic Behavior & Organization*, 171: 286–296

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

Thöni, C., **Tyran, J.R.**, and **Wengström, E.** (2012). Microfoundations of social capital. *Journal of Public Economics*, 96(7-8): 635–643

Wright, A.G. and **Simms, L.J.** (2014). On the structure of personality disorder traits: Conjoint analyses of the CAT-PD, PID-5, and NEO-PI-3 trait models. *Personality Disorders: Theory, Research, and Treatment*, 5(1): 43–54

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

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

Web Appendix
(Intended for Online Publication)

David Gill, Department of Economics, Purdue University

Yaroslav Rosokha, Department of Economics, Purdue University

This version: July 8, 2022

First version: July 15, 2020

Web Appendix I Screenshots from the experiment

Welcome!

Today's experiment will last about 60 minutes.

You will be paid a show-up fee of $5 together with any money you accumulate during this experiment. The amount of money you accumulate will depend partly on your actions and partly on the actions of other participants. This money will be paid at the end of the experiment in private and in cash.

It is important that during the experiment you remain **silent**. If you have a question or need assistance of any kind, please **raise your hand, but do not speak** - and an experiment administrator will come to you, and you may then whisper your question.

In addition, please **turn off your cell phones and put them away now**. Please do not use or place on your desk any personal items, including pens, paper, phones etc. Please do not look into anyone else's booth at any time.

Anybody that breaks these rules will be asked to leave.

Agenda:

- Questionnaire and test
- Experimental Instructions
- Experiment

Next

Figure A.1: Introductory screen

How Matches Work

The experiment is made up of **25 matches**.

At the start of each match you will be randomly paired with another participant in this room.

You will then play a number of rounds with that participant (this is what we call a "match").

Each match will last for a random number of **rounds**:

- At the end of each round the computer will roll a four-sided fair dice.
- If the computer rolls a 1, 2, or 3, then the match continues for at least one more round (75% probability).
- If the computer rolls a 4, then the match ends (25% probability).

When the match ends, you will again be randomly paired with another participant for the next match.

Choices and Payoffs

In each round of a match, you will choose A or B. The participant you are matched with will also choose A or B.

The table below describes the payoffs from the four pairs of choices that are possible in each round of a match.

These payoffs are in **points**.

My Choice in Round X	A	A	B	B
Other's Choice in Round X	A	B	A	B
My Payoff in Round X	32	12	50	25
Other's Payoff in Round X	32	50	12	25

In words:

- If you choose A, and the participant you are matched with chooses A, then your payoff for the round will be 32 points, and the other's payoff will be 32 points.
- If you choose A, and the participant you are matched with chooses B, then your payoff for the round will be 12 points, and the other's payoff will be 50 points.
- If you choose B, and the participant you are matched with chooses A, then your payoff for the round will be 50 points, and the other's payoff will be 12 points.
- If you choose B, and the participant you are matched with chooses B, then your payoff for the round will be 25 points, and the other's payoff will be 25 points.

At the end of the experiment, your total points will be converted into cash at the exchange rate of 250 points = $1. In addition, you will be paid your show-up fee of $5.

Next

Figure A.2: Description of matches, choices and payoffs

Practice Matches

To get some practice with the choices, you will now play 10 practice matches.

In these practice matches you are not yet matched with another participant and you will not earn any money.

In these practice matches you will make choices both for yourself and for the other participant.

Practice Match #2 of 10

The table below describes the payoffs from the four pairs of choices that are possible in each round of a match.

My Choice in Round X	A	A	B	B
Other's Choice in Round X	A	B	A	B
My Payoff in Round X	32	12	50	25
Other's Payoff in Round X	32	50	12	25

Remember that at the end of each round the computer rolls a four-sided fair dice. The match ends when the computer rolls a 4.

Dice Roll		2	1
Round	3	2	1
My Choice		B	A
Other's Choice		A	A
My Payoff		50	32
Other's Payoff		12	32

My Choice in Round **3**

A B

Other's Choice in Round **3**

A B

Figure A.3: Practice matches: Choice screen

Practice Matches

To get some practice with the choices, you will now play 10 practice matches.

In these practice matches you are not yet matched with another participant and you will not earn any money.

In these practice matches you will make choices both for yourself and for the other participant.

Practice Match #2 of 10

The table below describes the payoffs from the four pairs of choices that are possible in each round of a match.

My Choice in Round X	A	A	B	B
Other's Choice in Round X	A	B	A	B
My Payoff in Round X	32	12	50	25
Other's Payoff in Round X	32	50	12	25

Remember that at the end of each round the computer rolls a four-sided fair dice. The match ends when the computer rolls a 4.

Dice Roll	4	2	3	2	1
Round	5	4	3	2	1
My Choice	A	B	A	B	A
Other's Choice	B	B	A	A	A
My Payoff	12	25	32	50	32
Other's Payoff	50	25	32	12	32

Practice Match #2 Summary
Match duration: **5** rounds
My total payoff: **151** points
Other's total payoff: **151** points

Next

Figure A.4: Practice matches: Summary screen

Plans

When you are paired with another participant for a match, you will select a plan that will make your choices for you in every round. The other participant will also select a plan.

There are 10 possible plans.

Plans	
#1	Choose randomly between A and B in every round. At the beginning of every round, the computer flips a computerized fair coin for you: when your coin comes up heads, you choose A; when your coin comes up tails you choose B.
#2	Choose A in round 1. In round 2: choose A if the other chose A in round 1; choose B if the other chose B in round 1. After round 2: choose A if the other chose A in both of the previous two rounds; choose B if the other chose B in either of the previous two rounds.
#3	Choose B in round 1. After round 1: choose A if the other chose A in the previous round; choose B if the other chose B in the previous round.
#4	Choose B in every round.
#5	Choose A in round 1. After round 1: choose A if the other chose A in the previous round; choose B if the other chose B in the previous round.
#6	Choose A in round 1. After round 1: choose A if the other chose A in every one of the previous rounds; choose B if the other chose B in one or more of the previous rounds.
#7	Choose A in rounds 1 and 2. After round 2: choose A if the other has never chosen B twice in a row (i.e., if the other has never chosen B in two consecutive previous rounds); choose B if the other has ever chosen B twice in a row.
#8	Choose A in rounds 1 and 2. After round 2: choose A if the other chose A in either of the previous two rounds; choose B if the other chose B in both of the previous two rounds.
#9	Choose A in every round.
#10	Choose B in round 1. After round 1: choose A if the other chose A in every one of the previous rounds; choose B if the other chose B in one or more of the previous rounds.

Review Plans

Time Left: **2:42**

You have 3 minutes to review the table presenting the 10 plans.

You will have access to this table throughout the experiment.

Figure A.5: Review of strategies

Notes: The order that the ten strategies appeared on the subject's screen was randomized across subjects (see Web Appendix III.2). In this example, the order of the strategies is: RAND; 2TFT; DTFT; AD; TFT; G; G2; TF2T; AC; DG.

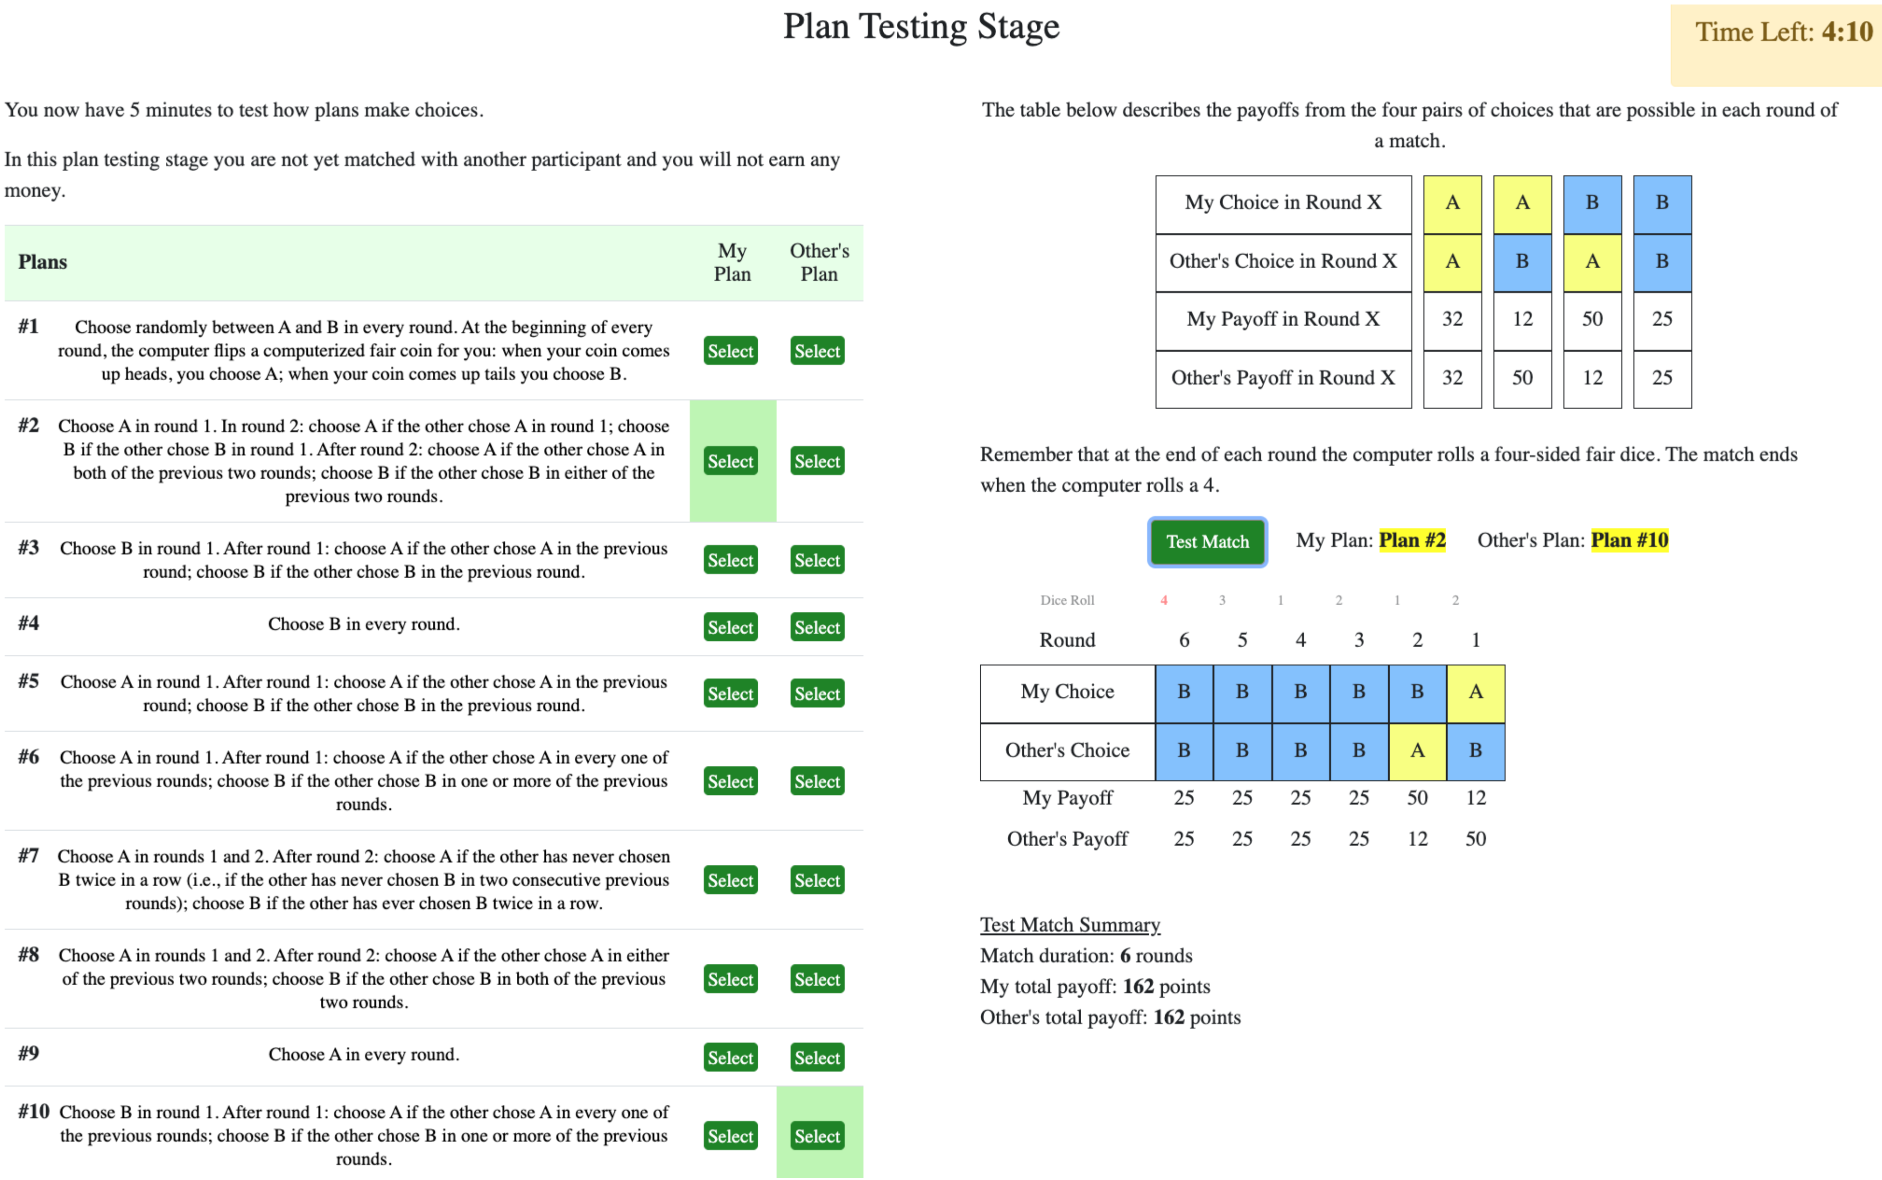

Figure A.6: Testing strategies: Summary screen

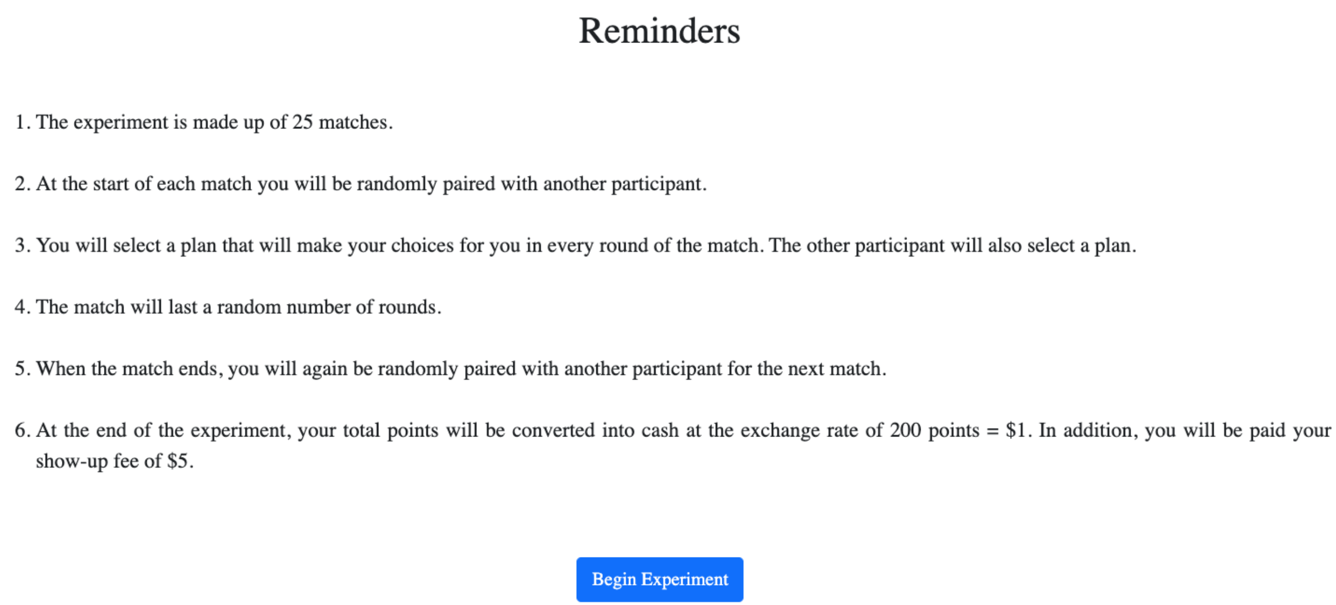

Reminders

1. The experiment is made up of 25 matches.

2. At the start of each match you will be randomly paired with another participant.

3. You will select a plan that will make your choices for you in every round of the match. The other participant will also select a plan.

4. The match will last a random number of rounds.

5. When the match ends, you will again be randomly paired with another participant for the next match.

6. At the end of the experiment, your total points will be converted into cash at the exchange rate of 200 points = $1. In addition, you will be paid your show-up fee of $5.

Begin Experiment

Figure A.7: Reminder screen

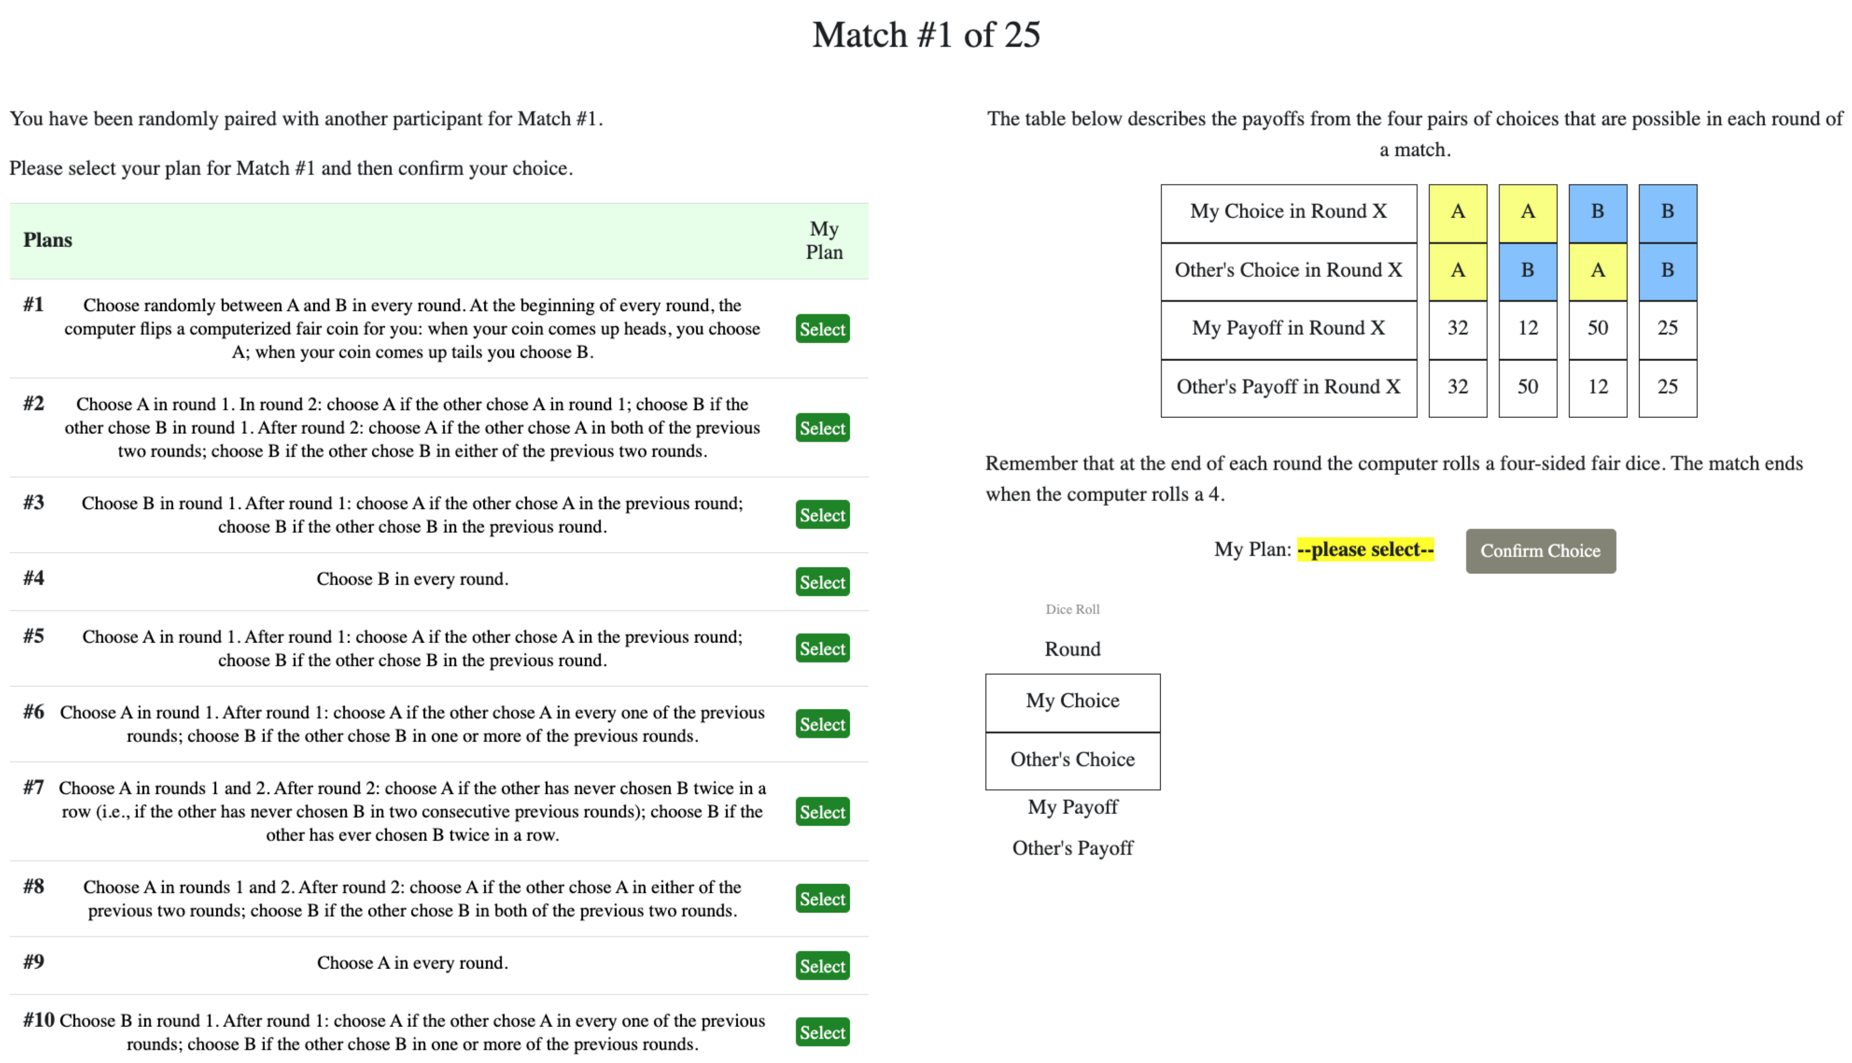

Figure A.8: Supergame 1: Choice screen

Question Before We Implement Your Plan in Match # 1

Before we implement your plan in Match #1, **on the next screen we will ask you what you think the chances are that the other participant chooses each of the ten plans.**

In other words, we will ask you how many times out of 100 you think the other participant in Match #1 would choose each of the ten plans.

Of course, the other participant chooses his/her plan for Match #1 only once (just like you), not 100 times. But you can think of the question as a way of asking how likely the other participant is to choose each of the plans.

You will be paid according to the accuracy of your answer. You will make the most money on average if you answer truthfully what you think the chances are that the other participant chooses each of the plans.

You will not be paid for your answer until the end of the experiment. Your answer will not be shown to any other participant. Your answer will not affect the experiment in any way.

You do not need to understand the details of how the payment works. If you are interested, the details follow below. If you are not interested, you can stop reading this screen now.

Details of how the payment works:

> *For this task, you start with 400 points and you will then be penalized for inaccuracy. The total penalty can never be more than 400 points. Your payment for this task will be 400 points minus the total penalty.*
>
> *For each of the ten possible plans, the penalty will be calculated as follows.*
>
> 1. *If that plan is chosen by the other participant, then the penalty will be smaller the higher your answer about the chances that the other participant chooses that plan. In that case, the penalty is calculated as follows:*
>
>   *(100 − Your Answer) × (100 − Your Answer) × (0.02 points)*
>
> 2. *If that plan is not chosen by the other participant, then the penalty will be smaller the lower your answer about the chances that the other participant chooses that plan. In that case, the penalty is calculated as follows:*
>
>   *(Your Answer) × (Your Answer) × (0.02 points)*

Next

Figure A.9: Supergame 1: Belief elicitation information

Question Before We Implement Your Plan in Match # 1

Plans		Other's Plan (Chance, %)
#1	Choose randomly between A and B in every round. At the beginning of every round, the computer flips a computerized fair coin for you: when your coin comes up heads, you choose A; when your coin comes up tails you choose B.	
#2	Choose A in round 1. In round 2: choose A if the other chose A in round 1; choose B if the other chose B in round 1. After round 2: choose A if the other chose A in both of the previous two rounds; choose B if the other chose B in either of the previous two rounds.	
#3	Choose B in round 1. After round 1: choose A if the other chose A in the previous round; choose B if the other chose B in the previous round.	
#4	Choose B in every round.	
#5	Choose A in round 1. After round 1: choose A if the other chose A in the previous round; choose B if the other chose B in the previous round.	
#6	Choose A in round 1. After round 1: choose A if the other chose A in every one of the previous rounds; choose B if the other chose B in one or more of the previous rounds.	
#7	Choose A in rounds 1 and 2. After round 2: choose A if the other has never chosen B twice in a row (i.e., if the other has never chosen B in two consecutive previous rounds); choose B if the other has ever chosen B twice in a row.	
#8	Choose A in rounds 1 and 2. After round 2: choose A if the other chose A in either of the previous two rounds; choose B if the other chose B in both of the previous two rounds.	
#9	Choose A in every round.	
#10	Choose B in round 1. After round 1: choose A if the other chose A in every one of the previous rounds; choose B if the other chose B in one or more of the previous rounds.	
	The sum must add up to 100: Check Sum	

What do you think the chances are that the other participant chooses each of the ten plans?

In other words, how many times out of 100 do you think the other participant in Match #1 would choose each of the ten plans?

- Please use only whole numbers for your answer.
- You will be paid according to the accuracy of your answer. You will make the most money on average if you answer truthfully what you think the chances are that the other participant chooses each of the plans.
- You will not be paid for your answer until the end of the experiment.
- Your answer will not be shown to any other participant.
- Your answer will not affect the experiment in any way.

Figure A.10: Supergame 1: Belief elicitation

Match #1 of 25

Plans		My Plan
#1	Choose randomly between A and B in every round. At the beginning of every round, the computer flips a computerized fair coin for you: when your coin comes up heads, you choose A; when your coin comes up tails you choose B.	Select
#2	Choose A in round 1. In round 2: choose A if the other chose A in round 1; choose B if the other chose B in round 1. After round 2: choose A if the other chose A in both of the previous two rounds; choose B if the other chose B in either of the previous two rounds.	Select
#3	Choose B in round 1. After round 1: choose A if the other chose A in the previous round; choose B if the other chose B in the previous round.	Select
#4	Choose B in every round.	Select
#5	Choose A in round 1. After round 1: choose A if the other chose A in the previous round; choose B if the other chose B in the previous round.	Select
#6	Choose A in round 1. After round 1: choose A if the other chose A in every one of the previous rounds; choose B if the other chose B in one or more of the previous rounds.	Select
#7	Choose A in rounds 1 and 2. After round 2: choose A if the other has never chosen B twice in a row (i.e., if the other has never chosen B in two consecutive previous rounds); choose B if the other has ever chosen B twice in a row.	Select
#8	Choose A in rounds 1 and 2. After round 2: choose A if the other chose A in either of the previous two rounds; choose B if the other chose B in both of the previous two rounds.	Select
#9	Choose A in every round.	Select
#10	Choose B in round 1. After round 1: choose A if the other chose A in every one of the previous rounds; choose B if the other chose B in one or more of the previous rounds.	Select

The table below describes the payoffs from the four pairs of choices that are possible in each round of a match.

My Choice in Round X	A	A	B	B
Other's Choice in Round X	A	B	A	B
My Payoff in Round X	32	12	50	25
Other's Payoff in Round X	32	50	12	25

Remember that at the end of each round the computer rolls a four-sided fair dice. The match ends when the computer rolls a 4.

My Plan: **Plan #4**

Dice Roll	4	1	3	2
Round	4	3	2	1
My Choice	B	B	B	B
Other's Choice	B	B	A	A
My Payoff	25	25	50	50
Other's Payoff	25	25	12	12

Match # 1 Summary
Match duration: **4** rounds
My total payoff: **150** points
Other's total payoff: **74** points

Next

Figure A.11: Supergame 1: Summary screen

Web Appendix II Details referred to in footnotes

Web Appendix II.1 Details referred to in footnote 5

Li and Liu (2017) find a weak relationship between group identity and beliefs in the first round of their first indefinitely repeated prisoner's dilemma game. Dreber et al. (2014) find that cooperation in indefinitely repeated prisoner's dilemma games with mistakes in implementation is correlated with beliefs (measured after the games) about the likelihood that defections were due to mistakes. When subjects play a sequence of indefinitely repeated prisoner's dilemma games where they are randomly rematched after every round within a supergame: (i) Duffy and Ochs (2009) find that cooperation is low and can be predicted well using threshold strategies elicited in an earlier one-shot prisoner's dilemma together with subjects' round-by-round forecasts of how many others in the matching group of fourteen will cooperate in that round; and (ii) Duffy and Fehr (2018) elicit beliefs about how many others in the matching group of ten will cooperate in the first round of each supergame, finding that beliefs and behavior are correlated and that beliefs respond to experience in an earlier stag hunt game.

Web Appendix II.2 Details referred to in footnote 6

In one-shot prisoner's dilemmas, see: Messé and Sivacek (1979); Croson (2000); Miettinen and Suetens (2008); Charness et al. (2016); Engel and Zhurakhovska (2016); Ridinger and McBride (2017); Peeters and Vorsatz (2018); Heuer and Orland (2019); and Sutter and Untertrifaller (2020). Croson (2000) also finds that eliciting beliefs lowers cooperation (but see footnote 8 in the main text), while Charness et al. (2016) also find that beliefs change with the payoff from joint cooperation. In Miettinen and Suetens (2008), beliefs are more cooperative when both players send a message expressing a desire for mutual cooperation.

Web Appendix II.3 Details referred to in footnote 9

In one-shot or finitely repeated prisoner's dilemmas, Haesevoets et al. (2018) find that a survey measure of trust predicts cooperation, Emonds et al. (2014) find that a survey measure of trust predicts cooperation among prosocial subjects, while Acedo-Carmona and Gomila (2014) find that subjects cooperate more when they are matched with people they know and trust personally, although Ahn et al. (2003) find no effect of a survey measure of trust. Papers that correlate survey measures of trust with contributions in one-shot or finitely repeated linear public goods games mostly find a positive effect of trust (Sato, 1988, 1989; Anderson et al., 2004; Gächter et al., 2004; Thöni et al., 2012; Peysakhovich et al., 2014), while Mulder et al. (2006) find that experiencing a treatment with sanctions undermines contributions of high trust subjects. Finally, in a real-effort game, Proto and Rustichini (2014) find that trust predicts whether subjects choose effort consistent with believing that others are cooperative.

Web Appendix II.4 Details referred to in footnote 11

Proto et al. (2019) find that when subjects are matched according to their level of conscientiousness, high conscientiousness subjects cooperate less in the indefinitely repeated prisoner's dilemma. Furthermore, this reduction in cooperation is driven by the cautiousness facet of conscientiousness. However, when subjects are not matched according to conscientiousness, Proto et al. (2019) find that conscientiousness does not have a statistically significant effect on cooperation (and cautiousness was not measured in this treatment): as they state: "the presence of two highly conscientious players – rather than one individual – seems a necessary condition for the trait to have a measurable impact on outcomes." Thus, in our setting in which subjects are not matched by traits, our result that cautiousness does not predict cooperation is consistent with Proto et al. (2019) (recall that we measured the cautiousness facet of conscientiousness but not conscientiousness itself).

Proto et al. (2019) also find that agreeableness has a transitory effect on cooperation: Proto et al. (2019) mention that trust and altruism are facets of agreeableness, but they do not study the effects of these facets on cooperation. Finally, Proto et al. (2019) do not find any systematic effect of the other Big Five personality measures (extraversion, openness and neuroticism) on cooperation in the indefinitely repeated prisoner's dilemma.

When studying cooperation in the indefinitely repeated prisoner's dilemma, Proto et al. (forthcoming) include the Big Five personality measures as controls, without matching subjects by personality or discussing any results on personality. Table 3 in Proto et al. (forthcoming) reports the effects of personality on cooperation in the very first round of the experiment, with an effect of agreeableness at the ten-percent level and no statistically significant effects of the other measures.

Web Appendix III Details on experimental design

Web Appendix III.1 Further details on procedures

All participants gave informed consent after reading the participant information sheet. The VSEEL subject pool is administered using ORSEE (Greiner, 2015). 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). 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. The experiment was programmed in oTree (Chen et al., 2016).

Web Appendix III.2 Further details on choice of strategies

We randomly created sixteen orders of the ten strategies, 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.

The round-by-round randomization for RAND was implemented in real time as play progressed. 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 (the reason is that RAND replicates every deterministic strategy with positive probability): this 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.

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.

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 setting without mistakes in implementation.

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).

Web Appendix III.3 Strategy categories

We find it useful to categorize our ten supergame strategies, as illustrated in Figure A.12.

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.[32]

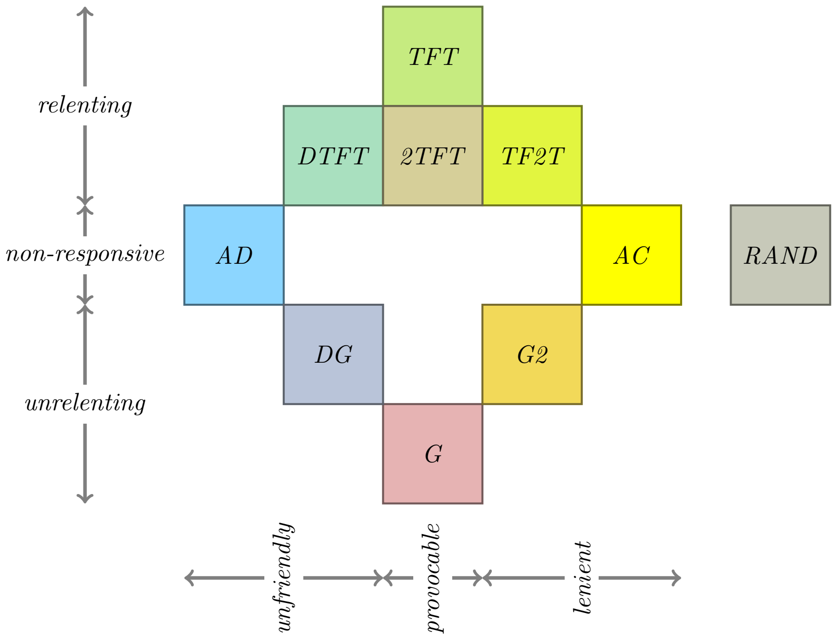

Figure A.12: Strategy categories

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,

[32]Fudenberg et al. (2012) categorize strategies as lenient in the same way that we do in Figure A.12. According to Fudenberg et al. (2012)'s terminology, our unfriendly strategies are 'fully noncooperative', while our 'provocable' strategies are 'fully cooperative' but not lenient.

TF2T).[33] The three 'non-responsive' strategies (AD, RAND, AC) never respond to a rival's defection.

Web Appendix III.4 Further details on supergames with strategy elicitation

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.

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

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.

Web Appendix III.5 Strategy elicitation and one-shot games

As we describe in more detail in Section 2.6, in our implementation of the indefinitely repeated prisoner's dilemma, at the beginning of each supergame each subject chooses one strategy to play the supergame, and then the subject's chosen supergame strategy and that of her opponent are played out round-by-round on the subject's screen. Thus, subjects in our experiment choose a strategy to which they are committed for the duration of the supergame: related work that also directly elicits supergame strategies in the indefinitely repeated prisoner's dilemma by making subjects choose a strategy to which they are committed includes Romero and Rosokha (2018), Cason and Mui (2019) and Dal Bó and Fréchette (2019). Eliciting supergame strategies in the indefinitely repeated prisoner's dilemma is an example of the strategy method, under which subjects have no opportunity to deviate from their strategy as the game proceeds (see the survey by Brandts and Charness, 2011).

This implementation of the indefinitely repeated prisoner's dilemma is related to one-shot games. However, there are differences:

- Subjects are not directly given the ten-by-ten supergame strategy expected payoff matrix. Instead, they are given the payoffs for the stage game of the indefinitely repeated prisoner's dilemma, together with a description of each supergame strategy. We also note that the payoffs that arise from any pair of strategies are not deterministic because: (i) the game

[33] We use the term 'relenting' rather than 'forgiving' because Axelrod (1980)'s concept of forgiving in the prisoner's dilemma includes both relenting after a punishment and being lenient by not immediately punishing a defection.

lasts an uncertain number of rounds; and (ii) the RAND strategy randomizes round-by-round.

- Relatedly, subjects see round-by-round feedback (choices made by the strategies, payoffs, and die rolls that determine whether the supergame continues to the next round), and each round lasts two seconds in order to mimic the feedback from direct-response play (where subjects choose their actions round-by-round).

- Unlike one-shot games with feedback, subjects do not directly observe the strategy chosen by their opponent, which preserves external validity because in real-world strategic interactions people usually do not directly observe others' strategies.

If we made our implementation more like one-shot games with feedback by allowing subjects to directly observe the strategy chosen by their opponent, we conjecture that this would change how subjects learn from experience. In our implementation, when learning from experience, subjects are uncertain about the strategy chosen by their opponents in previous rounds (because the same within-supergame history of play can arise from multiple strategies). If we removed this uncertainty by allowing subjects to directly observe strategies, we conjecture that subjects would respond more strongly to experience. As a result, the comparisons that we make in Section 4.2 to the effects of learning from experience in Dal Bó and Fréchette (2018)'s meta-data from round-by-round choices would become less informative.

Web Appendix III.6 Strategy elicitation and equilibrium

When $\delta = 0.75$ and $R \in \{32, 40, 48\}$, Dal Bó and Fréchette (2011) show that full defection and full cooperation are both equilibrium outcomes of subgame-perfect Nash equilibria; e.g., (AD,AD) and (G,G) are both equilibria. As we show below, in our setting with strategy elicitation, full defection and full cooperation remain equilibrium outcomes.

We build on recent work that directly elicits supergame strategies in the indefinitely repeated prisoner's dilemma (Romero and Rosokha, 2018; Cason and Mui, 2019; Dal Bó and Fréchette, 2019). As in that literature, subjects in our experiment choose a strategy to which they are committed for the duration of the supergame. Thus, equilibrium analysis checks whether strategies are best-responses to each other, ignoring the possibility of deviation once the supergame has started. In this sense, the equilibrium analysis is like that for one-shot games. Inspecting the expected payoff matrices (Tables A.14, A.15 and A.16 in Web Appendix X) gives the following:

- When $R = 32$, the set of symmetric pure-strategy Nash equilibria is:
  {(AD,AD), (G,G), (2TFT,2TFT)}.

- When $R \in \{40, 48\}$, the set of symmetric pure-strategy Nash equilibria is:
  {(AD,AD), (G,G), (2TFT,2TFT), (TFT,TFT)}.

- When $R = 32$, the set of asymmetric pure-strategy Nash equilibria is:

  {(DG,G2), (DTFT,TFT), (G,2TFT)}.

- When $R \in \{40, 48\}$, the set of asymmetric pure-strategy Nash equilibria is:

  {(DG,G2), (DG,TF2T), (DTFT,G2), (DTFT,TF2T), (G,2TFT), (G,TFT), (2TFT,TFT)}.

All ten strategies are rationalizable, since each strategy is a best-response to some belief.[34]

Web Appendix III.7 Further details on belief elicitation

We endeavored to keep the description of the QSR 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.

Eliciting beliefs can potentially change behavior. For example, eliciting incentivized beliefs might change beliefs or make the importance of beliefs more salient, which in turn might affect behavior. Reassuringly, 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.[35] 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. Since we are specifically interested in learning, we designed our experiment to minimize contamination of strategy choices by eliciting beliefs only twice, in the first supergame and in the final supergame (and by eliciting beliefs in these supergames after subjects chose their strategy). By eliciting beliefs in the first and final supergames, we are able to study both initial beliefs and the change in beliefs from the beginning to the end of the experiment, while minimizing concerns that eliciting beliefs could change behavior. Eliciting beliefs only twice also reduces the cognitive complexity and length of our experiment.

The QSR is incentive compatible (Selten, 1998), which means that money-maximizing (risk-neutral) subjects are incentivized to report their true belief. 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

[34]Inspecting the payoff tables (Tables A.14, A.15 and A.16 in Web Appendix X), all ten strategies except RAND are always a best-response to at least one pure strategy, and when $R = 32$, RAND is also a best-response to a pure strategy. When $R = 40$, RAND is a best-respone to a 50-50 mix over TF2T and AC. When $R = 48$, RAND is a best-response to a 11-89 mix over TF2T and AC.

[35]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).

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.[36] 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.

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.

Finally, Schlag et al. (2015, p.479)'s survey finds no consensus on whether beliefs are influenced by first making a choice.

Web Appendix III.8 Further details on personality questionnaire

We included forgiveness, kindness and trust because we judged that these measures linked well to the strategy categories described in Web Appendix III.3 (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 Web Appendix III.3). 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 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

[36]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.

[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.

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. The anxiety, cautiousness and trust measures include ten questions each; the questions come from the 300-item IPIP-NEO (see Goldberg, 1999, and Johnson, 2014) and are at: https://ipip.ori.org/newNEOFacetsKey.htm. The manipulativeness measure includes six questions and is one of thirty-three scales from the Computerized Adaptive Test of Personality Disorder (CAT–PD); the questions come from the 212-item CAT–PD–SF (see Simms et al., 2011, and Wright and Simms, 2014) and are at: https://ipip.ori.org/newCAT-PD-SFv1.1Keys.htm. The forgiveness and kindness measures include eight questions each and are two of the twenty-four character strengths from the Values in Action Inventory of Strengths (VIA-IS); the questions come from the 192-item VIA-IS-R (see Peterson and Seligman, 2004, and McGrath, 2017) and for research purposes are available on request from the VIA Institute on Character (https://www.viacharacter.org).

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. All 52 questions use a five-point Likert scale. For consistency, we presented all questions in the form 'I ...', and we used the introductory wording and response categories recommended by IPIP at: https://ipip.ori.org/New_IPIP-50-item-scale.htm. The 52 questions were split into five screens of ten questions and a final screen of two questions. Subjects could change their answers on a particular screen until they submitted their answers for that screen. Subjects were allowed to submit incomplete sets of answers, but were asked to confirm that they wanted to do so. We replaced missing responses by the sample average of nonmissing responses to that particular question.

We included the personality questionnaire at the beginning of the experiment because we were concerned that experience and earnings in the prisoner's dilemma could affect subjects' answers to the personality questions. We were less concerned that answering a personality questionnaire would affect behavior: as explained above, we randomized the order of the 52 questions, and our personality questionnaire is neutral in the sense that some questions are framed positively (e.g., "I trust what people say") while others are framed negatively (e.g., "I am wary of others"). Placing the personality questionnaire at the beginning of the experiment before subjects play games also follows recent practice in, e.g., Gill and Prowse (2016) and Proto et al. (2019, forthcoming), while Fréchette et al. (2017) measure personality before studying choice under risk and uncertainty.

Web Appendix III.9 Further details on personality factors

We undertook 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 were retained in the rotation; retaining factors with eigenvalues above one is a standard criterion for choosing the number of factors due to Kaiser (1960).

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. The ten highest loading questions underlying the trust factor are the ten questions that measure trust. The nine highest loading questions underlying the anxiety (cautiousness) factor come from the ten questions that measure anxiety (cautiousness). The six highest loading questions underlying the kindness factor come from the eight questions that measure kindness. The four highest loading questions underlying the manipulativeness factor come from the six questions that measure manipulativeness (and the other two questions also appear among the ten highest loading questions); the raw factor loads negatively on manipulativeness, and so to create the manipulativeness factor we changed the sign of all the factor loadings.

Web Appendix III.10 Further details on cognitive ability test

We 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) and measures fluid intelligence.[38] For each question, subjects have to identify (among six choices) the missing element that completes a visual pattern. For research purposes, the questions are available on request from ICAR (`https://icar-project.com`; see Condon and Revelle, 2014, for more about ICAR).

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.

Web Appendix III.11 Further details on demographics

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

[38]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. (2018), Gill and Prowse (2016), Fe et al. (forthcoming) and Proto et al. (2019).

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').

Web Appendix IV Further analysis of initial beliefs and behavior

Web Appendix IV.1 Introduction

Web Appendix IV.2 discusses our measure of optimism. Web Appendix IV.3 shows that the accuracy of beliefs increases as the return to joint cooperation goes up. Web Appendix IV.4 provides support for good responding as a useful measure and shows that the frequency with which subjects good respond interacts with their optimism. Web Appendix IV.5 finds that earnings increase with the accuracy of beliefs and with the ability of subjects to good respond to their beliefs.

Web Appendix IV.2 Discussion of our measure of optimism

As described in Section 3.2, our measure of optimism measures how often a subject expects others to cooperate: specifically, 'Optimism' measures the expected cooperation rate of a subject's belief distribution playing against itself (see the notes to Figure 4 for more details). We used this definition of optimism because it measures how often a subject expects the population of subjects (excluding herself) to cooperate when they play against each other; we find this intuitive, and furthermore this definition allows a direct comparison of beliefs to the level of cooperation in the population.

We prefer this definition to an alternative measure of optimism based on the level of cooperation that the subject's own strategy achieves against her belief distribution. The reason is that we want our measure of optimism to be independent of the subject's behavior, so that we can study cleanly the relationship between optimism and behavior.

We could have used a simpler measure of optimism that sums up the belief weights on cooperative strategies (where cooperative strategies are defined to be those that always cooperate when played against themselves). We call this simpler measure 'OptimismSimple'. This measure is cruder than ours since it weights beliefs on strategies like AC and G the same, even though such strategies cooperate differently against AD. It turns out that OptimismSimple is highly correlated with Optimism: see Figure A.13. Furthermore, when we replace Optimism in Figure 4 with OptimismSimple we get very similar results: compare Panels I and III of Figure 4 with the equivalent panels in Figure A.14 here.

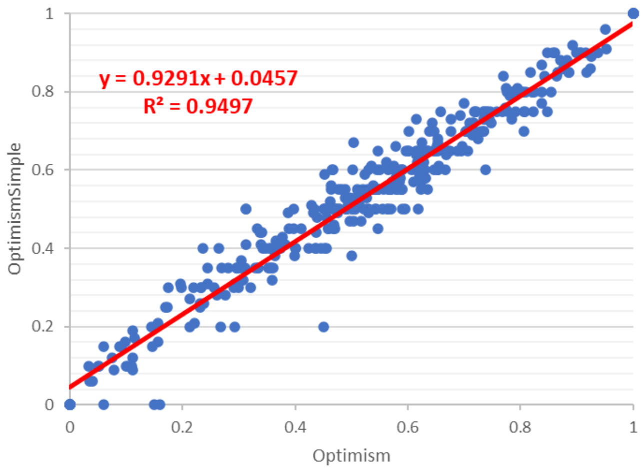

Figure A.13: Scatterplot of OptimismSimple vs. Optimism in Supergame 1

Notes: See the preceding paragraph for the definition of OptimismSimple.

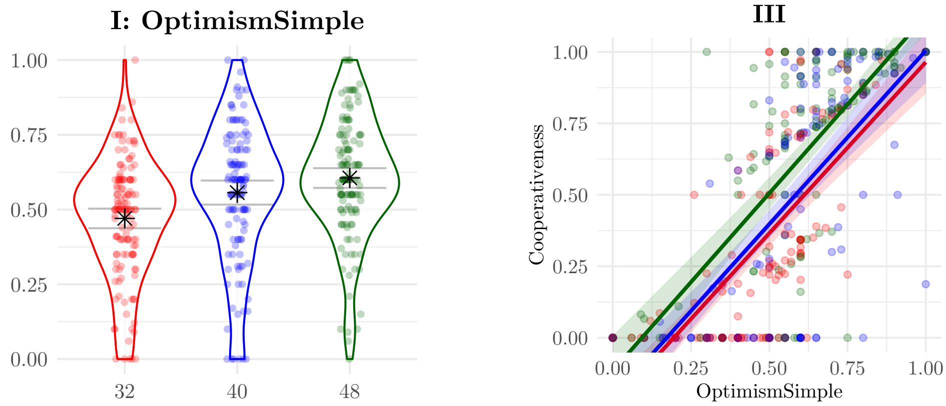

Figure A.14: Panels I and III of Figure 4, using OptimismSimple instead of Optimism

Web Appendix IV.3 Accuracy

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. 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.

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. In particular, we ran a linear OLS regression of accuracy in Supergame 1 on the treatment, controlling for the five personality factors, demographic characteristics and standardized cognitive ability (see Section 2.8), and using heteroskedasticity-robust standard errors and two-sided tests of significance. The positive effect of $R = 48$ relative to $R = 32$ is significant at the one-percent level. Column 1 of Table A.10 in Web Appendix VIII presents the full set of estimates from this regression.

Web Appendix IV.4 Analysis of good responding

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.25 in Web Appendix IX shows the cumulative distribution function of payoffs relative to best responding. Second, good responding matters for outcomes: in Web Appendix IV.5 we show that good responding is a strong predictor of earnings.[39,40] Third, good responding changes with beliefs: as we show below in this section, the frequency of good responding varies with the optimism of subjects' beliefs.

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 A.1: 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 3.3. Table A.12 in Web Appendix X replicates the table for best responding.

Table A.1 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

[39]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.

[40]If subjects' beliefs were completely wrong, then we would not expect good responding to predict earnings. In that case, our measure of good responding would remain valid while becoming less useful.

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.[41]

The frequency with which subjects good respond to their beliefs interacts with optimism in an interesting way. Recall from Section 3.2 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 ($R = 48$).[42] As evidenced by Table A.13 in Web Appendix X: (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.

Web Appendix IV.5 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 A.15 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.[43]

[41] 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 $R = 48$.

[42] We ran the regression described in Web Appendix IV.3, replacing accuracy as the dependent variable with an indicator taking value 1 if a subject chose a good response to her beliefs in Supergame 1, and further including optimism (defined in the notes to Figure 4) and the interaction of optimism with the treatment. Setting $R = 32$ ($R = 48$) as the omitted category, we find a negative (positive) effect of optimism on the probability of good responding, statistically significant at the one-percent (one-percent) level. We also find that the effect of optimism when $R = 48$ is significantly different to that when $R = 32$, again at the one-percent level.

[43] 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

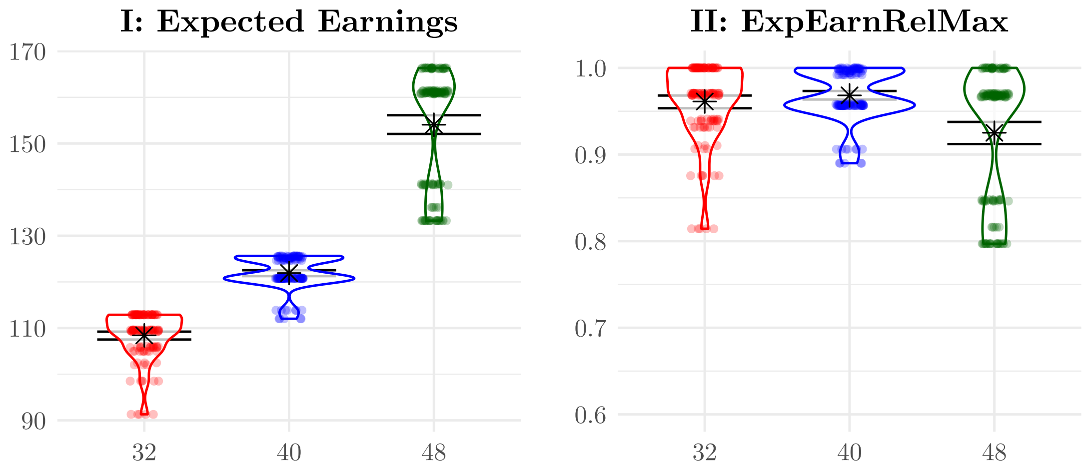

Figure A.15: Earnings in Supergame 1: Violin plots

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.14 to A.16 in Web Appendix X). We define 'ExpEarnRelMax' 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 A.2 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 A.15, the first two rows of Table A.2 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 A.2 shows that expected earnings increase with the accuracy of beliefs about the level of cooperation (Web Appendix IV.3 introduces our notion of accuracy).[44] 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 3.3 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 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.17 in Web Appendix X shows

selected opponent.

[44]To help interpret the effect size, note that our measure of accuracy is defined from -1 to 0. In Web Appendix IV.3 we define accuracy to be the negative of the absolute value of OptimismRelTruth; Panel IV of Figure 4 shows the distribution of OptimismRelTruth.

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 A.2: 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 2.8), and with $R = 32$ as the omitted category. Expected earnings is defined in the notes to Figure A.15. Accuracy of beliefs is defined in Web Appendix IV.3. Good responding is a binary variable defined in the second paragraph of Section 3.3. $N = 390$ 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).

Web Appendix V Further analysis of the evolution of beliefs and behavior

Web Appendix V.1 Introduction

Web Appendix V.2 shows that as subjects learn over the course of the 25 supergames, their beliefs about how often others cooperate become more accurate. To help understand how experience changes cooperation, Web Appendix V.3 studies the factors that drive experimentation and strategy revisions, and Web Appendix V.4 studies transitions between strategies: for example, we find that when a subject's opponent in the previous supergame cooperated more, the subject is much less likely to change to an "unfriendly" strategy that defects for sure in the first round. Web Appendix V.5 describes the evolution of equilibrium behavior and beliefs, while Web Appendix V.6 describes the evolution of strategy choices and average beliefs.

Web Appendix V.2 Evolution of OptimismRelTruth

Recall from Section 3.2 that 'OptimismRelTruth' measures optimism relative to how often others actually cooperate. Our finding from Table 3 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 A.16 we measure OptimismRelTruth at the session level. Figure A.16 confirms that, on average, beliefs do indeed move toward the truth in all three treatments. Confirming our finding from Section 3.2, when the return to joint cooperation is low ($R = 32$), subjects' initial beliefs are too optimistic relative to 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.

Importantly, OptimismRelTruth captures accuracy in terms of beliefs about behavior. This avoids two disadvantages of an alternative measure of accuracy based on the distance between a belief distribution and the average strategy distribution in a session or a treatment. First, this alternative measure penalizes mistakes that have small implications for behavior as much as mistakes that have much larger implications: for example, G and 2TFT cooperate similarly (see Table A.11 in Web Appendix X), and so believing that others choose G when they actually choose 2TFT is a minimal mistake in terms of behavior, but the alternative measure penalizes this mistake as much believing that others choose G when they actually choose AD, which implies a much bigger mistake in terms of behavior. Secondly, and relatedly, this alternative measure does not tell us whether beliefs on average are too optimistic or pessimistic about the rate of cooperation, and therefore, for example, we could not use the alternative measure to analyze how excess optimism declines with experience when $R = 32$ (see the paragraph above).

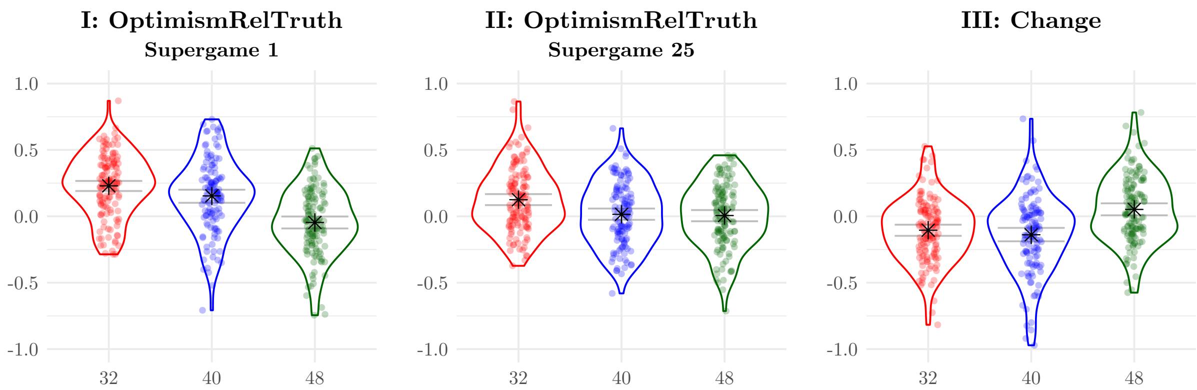

Figure A.16: Evolution of OptimismRelTruth at the session level: Violin plots

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.

Alongside our finding that beliefs become more accurate with experience, we also find that updated beliefs account for 23% of the variance in cooperation in the final supergame, while initial beliefs account for 36% of the variance in cooperation in the first supergame (final paragraph of Section 4.2). When interpreting these data, we note that considerable heterogeneity in beliefs and cooperation remains in the final supergame, both within and across treatment (Figure A.16 here and Figure A.21 in Web Appendix V.6). Table A.3 in Web Appendix V.3 shows that strategy revisions become less common as subjects gain experience: this implies less within-subject variation in cooperation across supergames but not less across-subject variation in cooperation in the final supergame.

Web Appendix V.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 A.3 we regress an indicator for changing strategy on the same variables that we analyzed in Panel I of Table 3 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 3.3 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 A.3 provide the formal definition).

	(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 A.3: 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 2.8), 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 3 (see the table notes for definitions). Good responding is a binary variable defined in the second paragraph of Section 3.3. '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.14 to A.16 in Web Appendix X). 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).

To summarize the main findings in Table A.3, 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.

In more detail, the third row of Table A.3 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 3). 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.[45] Just as in Panel I of Table 3, these results control for the subject's behavior and beliefs in the first supergame (fifth and sixth rows).

In Web Appendix V.4 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 A.3), while at the same time increasing the subject's own cooperation (Table 3).

Web Appendix V.4 Strategy transitions

In Table A.4 we study transitions from one strategy category to another (Web Appendix III.3 defines the strategy categories). Panel I uses the observations where the subject chose a strategy from the unfriendly category in the previous supergame, and we consider the factors that drive the subject to continue choosing an unfriendly strategy in the next supergame, to change to a provocable strategy, or to change to a lenient strategy. Panel II (III) repeats the exercise, using observations where the subject chose a strategy from the provocable (lenient) category in the previous supergame. The regressions reported in Table A.4 include the same independent variables as the regressions reported in Table A.3 (to save space we do not report the treatment coefficients).

[45]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.

I: Unfriendly strategy in Supergame $t-1$

	(1) Stay unfriendly	(2) Change to provocable	(3) Change to lenient
Length of Sup. $t-1$	-0.006***	0.005***	0.002*
	(0.001)	(0.002)	(0.001)
Other's strategy coop in Sup. $t-1$	0.068***	-0.025*	-0.025**
	(0.017)	(0.014)	(0.011)
Own strategy coop in Sup. 1	-0.155***	0.109***	0.050*
	(0.040)	(0.026)	(0.027)
Own optimism in Sup. 1	-0.161***	0.048	0.100***
	(0.054)	(0.037)	(0.027)
Good responder to beliefs in Sup. 1	-0.004	-0.007	0.020
	(0.022)	(0.017)	(0.013)
Quality of Sup. $t-1$ strategy	-0.010	-0.025	0.006
	(0.084)	(0.070)	(0.037)
Supergame number	0.004***	-0.002**	-0.001***
	(0.001)	(0.001)	(0.000)
Mean of DV	0.827	0.102	0.045
N	4820	4820	4820

II: Provocable strategy in Supergame $t-1$

	(1) Change to unfriendly	(2) Stay provocable	(3) Change to lenient
Length of Sup. $t-1$	-0.009***	0.009***	-0.001
	(0.002)	(0.003)	(0.001)
Other's strategy coop in Sup. $t-1$	-0.236***	0.250***	-0.006
	(0.031)	(0.031)	(0.013)
Own strategy coop in Sup. 1	-0.175***	0.171***	0.023
	(0.040)	(0.058)	(0.022)
Own optimism in Sup. 1	-0.111*	0.079	0.029
	(0.057)	(0.073)	(0.029)
Good responder to beliefs in Sup. 1	-0.023	0.027	0.001
	(0.019)	(0.024)	(0.009)
Quality of Sup. $t-1$ strategy	-0.124	0.128	-0.020
	(0.076)	(0.087)	(0.026)
Supergame number	-0.003***	0.005***	-0.002***
	(0.001)	(0.001)	(0.001)
Mean of DV	0.185	0.746	0.052
N	2676	2676	2676

III: Lenient strategy in Supergame $t-1$

	(1) Change to unfriendly	(2) Change to provocable	(3) Stay lenient
Length of Sup. $t-1$	-0.002	-0.003	0.006*
	(0.003)	(0.002)	(0.003)
Other's strategy coop in Sup. $t-1$	-0.202***	-0.047	0.260***
	(0.046)	(0.032)	(0.042)
Own strategy coop in Sup. 1	-0.190***	-0.047	0.250***
	(0.049)	(0.049)	(0.076)
Own optimism in Sup. 1	-0.175	-0.004	0.183
	(0.114)	(0.092)	(0.151)
Good responder to beliefs in Sup. 1	-0.034	0.007	0.029
	(0.027)	(0.022)	(0.035)
Quality of Sup. $t-1$ strategy	-0.010	-0.117**	0.135
	(0.063)	(0.048)	(0.084)
Supergame number	-0.003*	-0.004***	0.006***
	(0.001)	(0.001)	(0.002)
Mean of DV	0.167	0.090	0.720
N	1495	1495	1495

Table A.4: Strategy category transitions from Supergame $t-1$ to Supergame t

Notes: Panel I uses all observations from Supergames 1 to 24 where the subject chose a strategy from the unfriendly category (Web Appendix III.3 defines the strategy categories). The first column of Panel I reports a linear OLS regression of a binary variable that takes value 1 if the subject continued to choose a strategy from the unfriendly category in the next supergame, and taking value 0 otherwise; note that the variable takes value 1 even if the subject changed strategy within the unfriendly category. The 2nd (3rd) column of Panel I reports a linear OLS regression of a binary variable that takes value 1 if the subject changed her strategy to one from the provocable (lenient) category in the next supergame, and taking value 0 otherwise. Panels II and III are constructed similarly. Transition probabilities do not sum to one because subjects can change to RAND, which is not included in the unfriendly, provocable or lenient categories. All regressions control for the five personality factors, demographic characteristics and standardized cognitive ability (see Section 2.8), and the treatment (treatment coefficients are omitted to save space). The independent variables are the same as those in Table A.3; see the notes to that table. Heteroskedasticity-robust standard errors with clustering at the session level are shown in parentheses. ***, ** and * denote significance at the 1%, 5% and 10% levels (two-sided tests).

The first row of Panels I-III of Table A.4 show that when the previous supergame was longer, subjects are less likely to stay unfriendly (first column of Panel I), and more likely to stay provocable and lenient (second column of Panel II and third column of Panel III). Furthermore, the increased likelihood of staying provocable comes at the expense of changes from provocable to unfriendly strategies (first and second columns of Panel II). These findings help us to understand why previous supergame length has no effect on the likelihood of changing strategy (Table A.3), while at the same time increasing cooperation (Table 3).

The second row of Panels I-III of Table A.4 show that when a subject's opponent cooperated more in the previous supergame, the subject is more likely to stay within the unfriendly, provocable and lenient categories (first column of Panel I, second column of Panel II, and third column of Panel III). Even though subjects are somewhat more likely to stay unfriendly, they are also much less likely to change from provocable or lenient strategies to unfriendly ones when their opponent in the previous supergame cooperated more (first columns of Panels II and III). These findings shed light on the mechanism by which the opponent's cooperation in the previous supergame reduces the likelihood of changing strategy (Table A.3), while at the same time increasing the subject's own cooperation (Table 3).

Although the effects are not individually statistically significant, the fifth and six rows of Panels I-III of Table A.4 show that quality of thinking and the quality of the strategy chosen in the previous supergame reduce the likelihood of staying unfriendly (first column of Panel I), but increase the likelihood of staying provocable and lenient by more (second column of Panel II and third column of Panel III). These findings help to explain how quality of thinking and the quality of the strategy chosen in the previous supergame reduce the likelihood of changing strategy (Table A.3). Finally, the seventh row shows that as subjects gain experience by playing more supergames, they are more likely to stay unfriendly, provocable and lenient; thus, the finding that subjects who have played more supergames change strategy less frequently (Table A.3) holds at the level of all three strategy categories.

Web Appendix V.5 Evolution of equilibrium behavior

Web Appendix III.6 describes the set of equilibrium strategies for each R. Figures A.17 and A.18 show that the proportion of chosen strategies that are equilibrium strategies and the weight that beliefs place on equilibrium strategies both show a modest tendency to increase from the first supergame to the final supergame, with the effect slightly larger in magnitude for beliefs.

For completeness, this analysis considers symmetric equilibrium strategies, asymmetric equilibrium strategies and their combination, although random rematching and the absence of feedback meant that subjects had no obvious way to coordinate their roles in an asymmetric equilibrium.

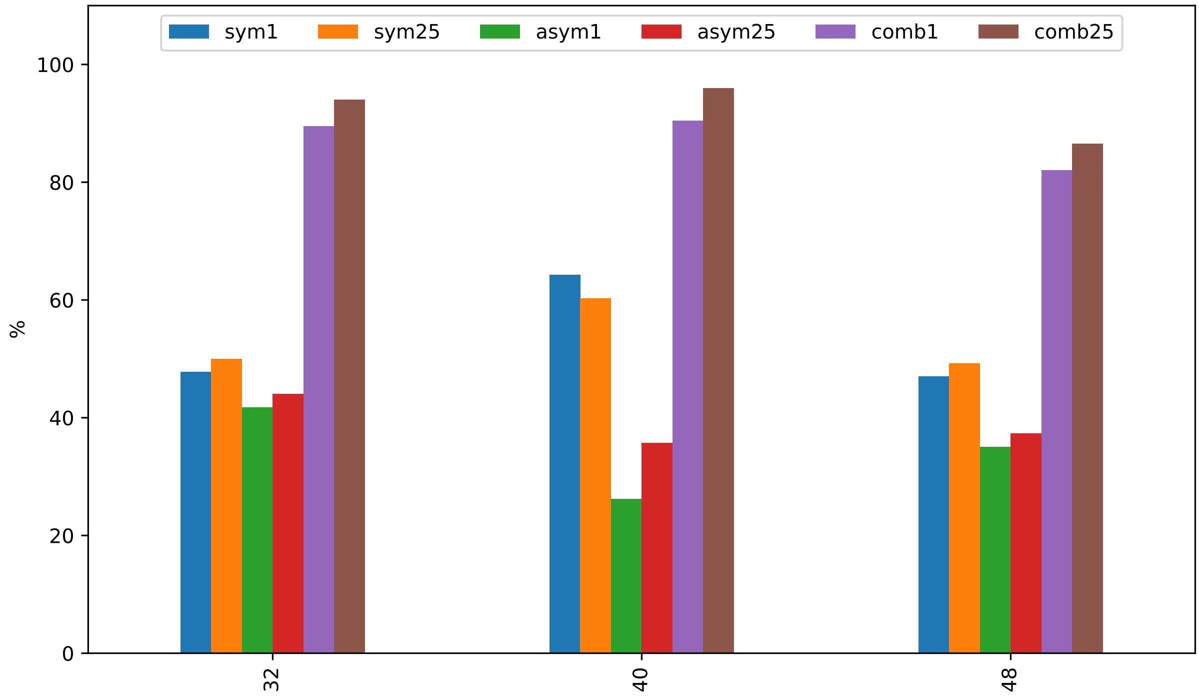

Figure A.17: Equilibrium strategies as proportion of chosen strategies in Supergame 1 vs. Supergame 25 (split by R)

Notes: 'sym1' ('sym25') is the proportion of symmetric equilibrium strategies among the strategies chosen in Supergame 1 (Supergame 25), 'asym1' ('asym25') is the proportion of asymmetric equilibrium strategies among the strategies chosen in Supergame 1 (Supergame 25), and 'comb1' ('comb25') combine the two. Web Appendix III.6 describes the set of equilibrium strategies. We say that a strategy is a 'symmetric equilibrium strategy' if that strategy is included in at least one symmetric pure-strategy Nash equilibrium, and we say that a strategy is an 'asymmetric equilibrium strategy' if it is included in at least one asymmetric pure-strategy Nash equilibrium, but is not included in any symmetric pure-strategy Nash equilibrium.

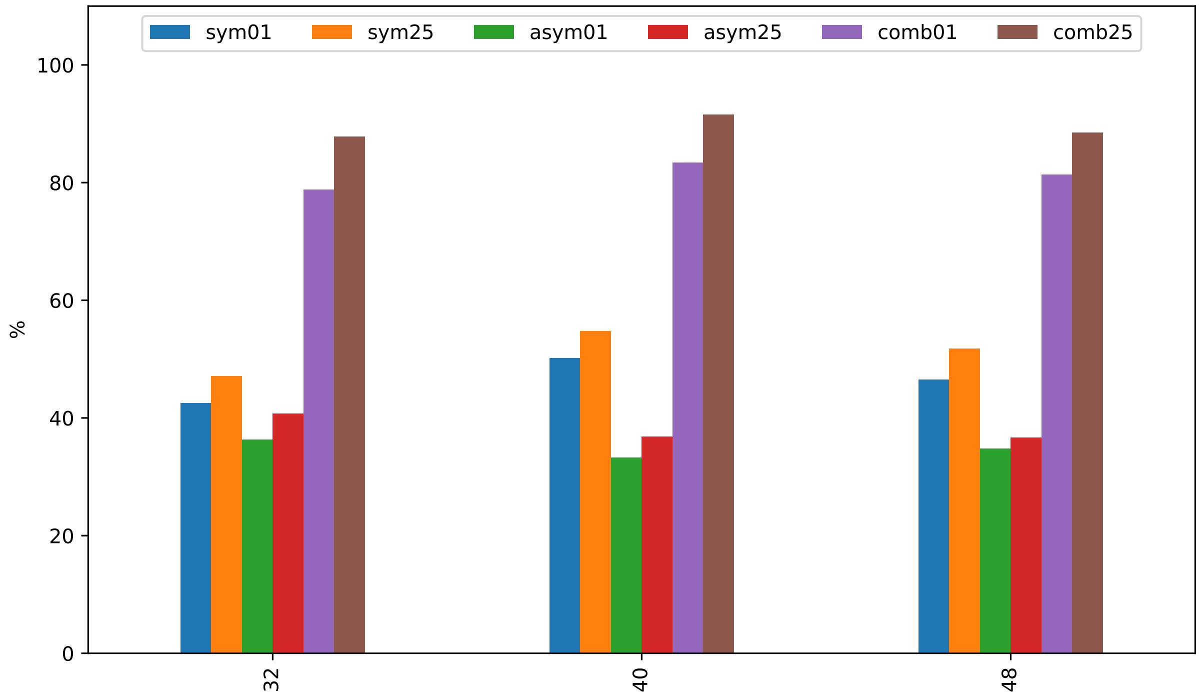

Figure A.18: Weight of beliefs on equilibrium strategies
in Supergame 1 vs. Supergame 25 (split by R)

Notes: See the notes to Figure A.17. Here, we report the mean probability weight placed on equilibrium strategies across each subject's belief distribution.

We also find that strategy revisions depend on how well subjects are responding to experience. In Web Appendix V.3 we find that subjects whose chosen strategy in supergame $t-1$ performs well in expectation given the subject's experience up to supergame $t-1$ are less likely to change strategy in supergame t (this effect of 'quality of supergame $t-1$ strategy' is statistically significant at the one-percent level). Furthermore, the mechanism links to the provocable strategies (G, 2TFT, TFT) that immediately punish a defection, which from Web Appendix III.6 make up all of the symmetric equilibrium strategies except for AD. In particular, when a subject has chosen a provocable strategy in supergame $t-1$, she is more likely to continue to choose a provocable strategy in supergame t the better her chosen strategy in supergame $t-1$ performs in expectation given the subject's experience up to supergame $t-1$ (see Web Appendix V.4, although the effect is not quite statistically significant).

Finally, Figure A.19 shows that, conditional on a subject changing her strategy from supergame $t-1$ to supergame t, her strategy choice moves in the direction of the best-response to her opponent's $t-1$ strategy (with the pattern holding for all three values of R). Panel (a) uses the opponent's actual strategy at $t-1$; since this choice is not directly observed by the subject, Panel (b) instead uses the distribution of the opponent's $t-1$ strategies inferred by the subject from the round-by-round choices in supergame $t-1$ (assuming Bayesian updating from a uniform prior).

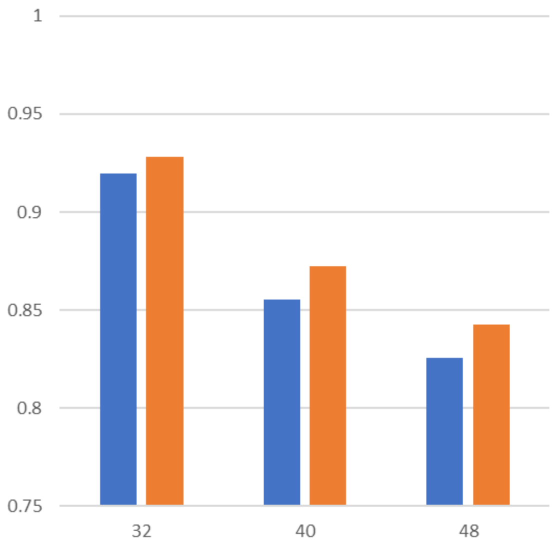

(a) vs. opponent's actual $t-1$ strategy

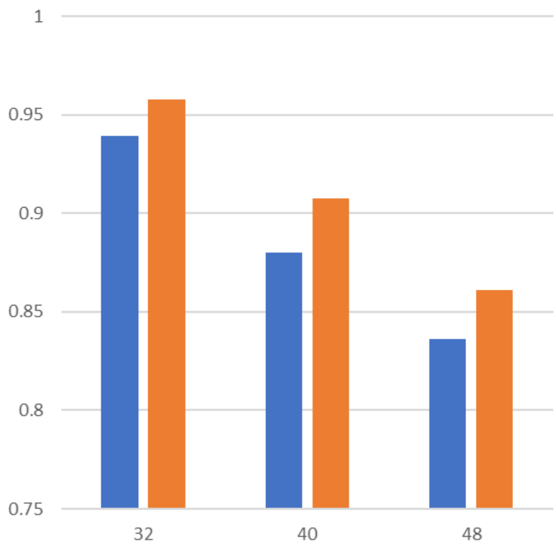

(b) vs. opponent's inferred $t-1$ strategies

Figure A.19: Payoff of revised strategy at t (orange) and original strategy at $t-1$ (blue), playing against opponent's $t-1$ strategy (as proportion of payoff from best-response)

Notes: See the preceding paragraph for details.

Web Appendix V.6 Evolution of strategy choices and beliefs

Figure 3 in Section 3.1 shows strategy choices and average beliefs in the first supergame; for ease of reference, Figure A.20 here replicates that figure. Figure A.21 presents the same data for the final supergame. To make the comparison between Figures A.20 and A.21 clearer, Figure A.21 further includes the data from the first supergame as black horizontal bars.

When comparing strategy choices and beliefs in the final supergame to those in the first supergame, a few observations stand out:

- In most of the thirty cases, the change in beliefs from the first to the final supergame is in the same direction as the change in the frequency of the corresponding strategy choice.
- In Section 3.2 we found that, in the first supergame, subjects are too optimistic about the level of cooperation when the return to joint cooperation is low ($R = 32$), and in Web Appendix V.2 we found that this excess optimism declines over time. Consistent with these findings, here we see that when $R = 32$: (i) in the first supergame, subjects' beliefs underestimate the proportion of unfriendly strategies and overestimate the proportion of lenient strategies; and (ii) comparing the final supergame to the first supergame, the beliefs about the proportions of unfriendly and lenient strategies become more accurate (unfriendly strategies are chosen more frequently, but the weight that beliefs place on unfriendly strategies increases by more; lenient strategies are chosen at around the same rate, while the weight that beliefs place on lenient strategies declines).
- In the final supergame, compared to the first supergame, the frequency of the most common strategy choice AD varies more in the return to joint cooperation R, as does the weight that beliefs place on AD.
- For all R, the frequency of DTFT (AC) and the weight that beliefs place on DTFT (AC) increase (decline) from the first to the final supergame.
- By the final supergame, average beliefs match the distribution of strategy choices quite closely.

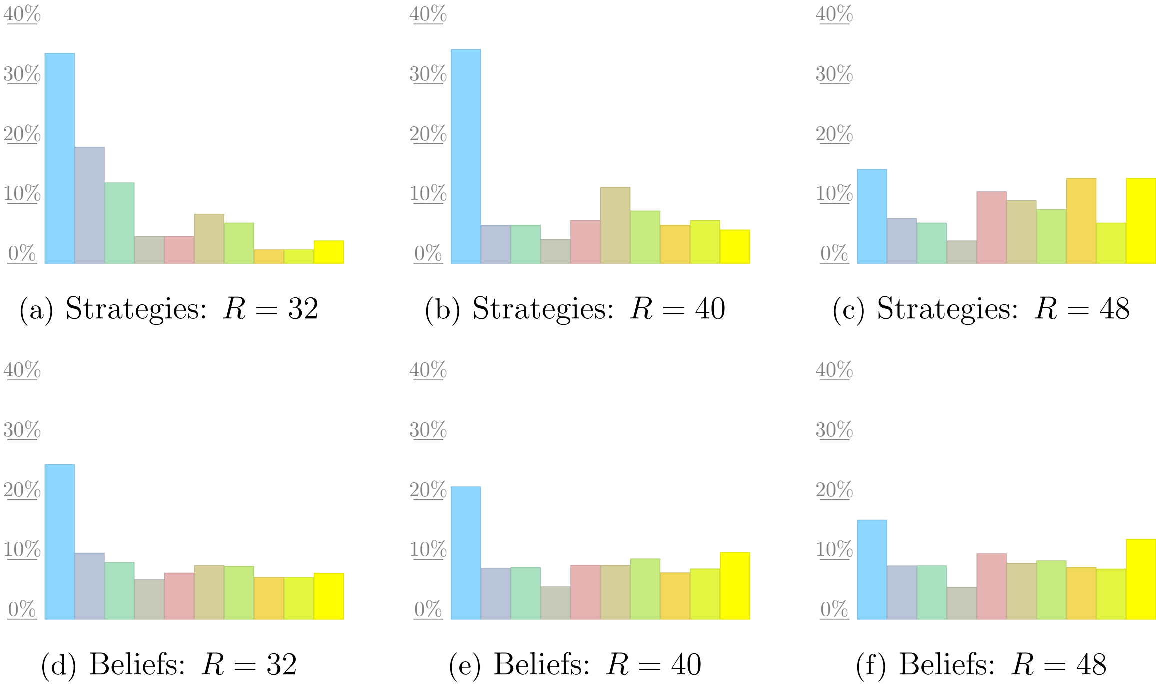

(a) Strategies: $R = 32$ (b) Strategies: $R = 40$ (c) Strategies: $R = 48$

(d) Beliefs: $R = 32$ (e) Beliefs: $R = 40$ (f) Beliefs: $R = 48$

Figure A.20: Strategies & average beliefs in Supergame 1 (replicates Figure 3)

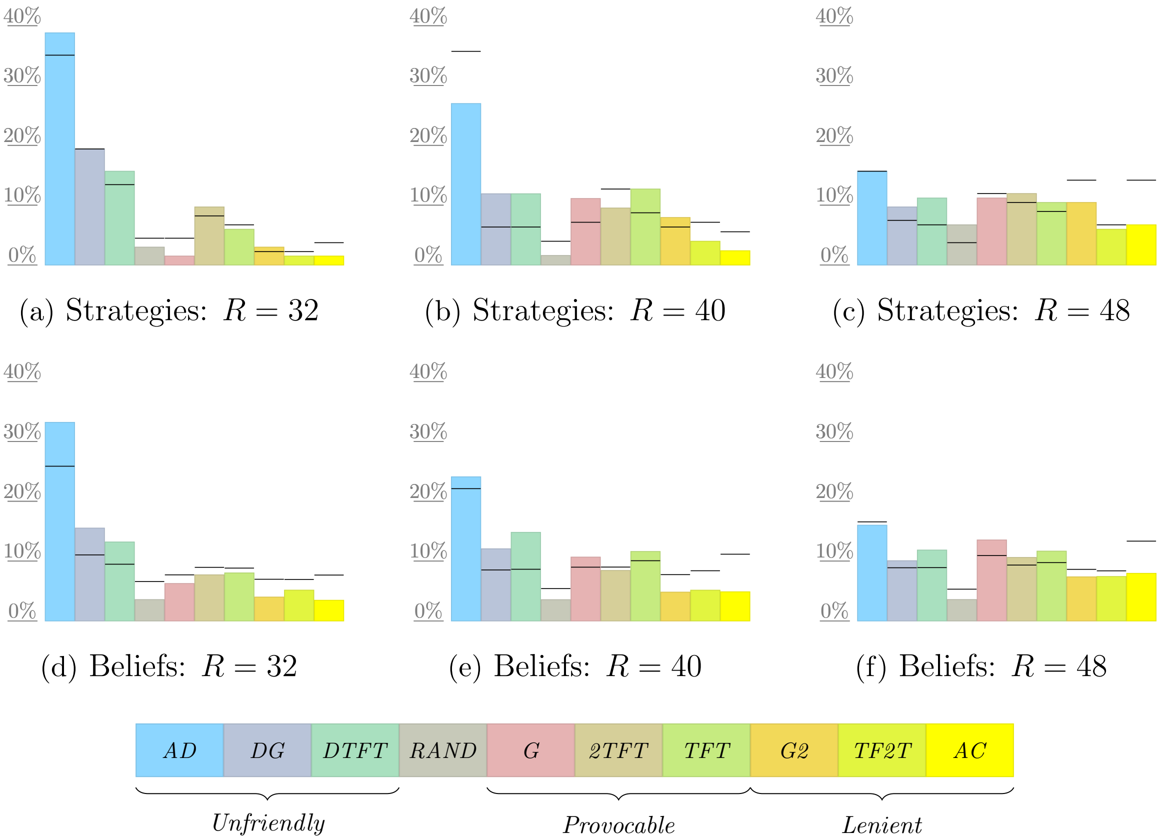

(a) Strategies: $R = 32$ (b) Strategies: $R = 40$ (c) Strategies: $R = 48$

(d) Beliefs: $R = 32$ (e) Beliefs: $R = 40$ (f) Beliefs: $R = 48$

Figure A.21: Strategies & average beliefs in Supergame 25 (black bars are Supergame 1)

Web Appendix VI Further details on personality

Web Appendix VI.1 Further tables and figures on personality

	(1) Strategy cooperation (Supergames 2-25)	(2) Strategy cooperation (Supergames 21-25)	(3) Optimism (Supergame 25)
Anxiety	-0.004	-0.023*	-0.005
	(0.009)	(0.013)	(0.012)
Cautiousness	-0.009	-0.011	-0.003
	(0.011)	(0.011)	(0.015)
Kindness	-0.004	-0.010	0.009
	(0.014)	(0.017)	(0.009)
Manipulativeness	-0.009	-0.005	-0.013
	(0.013)	(0.015)	(0.013)
Trust	0.021*	0.022**	0.024**
	(0.010)	(0.008)	(0.010)
Mean of dependent variable	0.467	0.461	0.430
N	9360	1950	390
Control for beliefs in Sup. 1	No	No	No
Control for behavior in Sup. 1	No	No	No
Controls for experience	No	No	No

Table A.5: Effect of personality on behavior and beliefs (robustness)

Notes: The regressions reported here are the same as those reported in Table 4, except that they exclude the controls for experience and for Supergame 1 behavior and beliefs presented in rows three to six of Panel I of Table 3 (in relation to Columns 1 and 2 here) and of Panel II of Table 3 (in relation to Column 3 here); Columns 1 and 2 also exclude the supergame number control (not relevant to Column 3).

	(1) Strategy cooperation (Supergame 1)	(2) Strategy cooperation (Supergame 1)	(3) Optimism (Supergame 1)
Anxiety	0.009	0.017	-0.008
	(0.019)	(0.016)	(0.011)
Cautiousness	-0.002	0.000	-0.002
	(0.017)	(0.013)	(0.012)
Kindness	-0.022	-0.006	-0.017
	(0.017)	(0.014)	(0.011)
Manipulativeness	-0.027	-0.008	-0.020*
	(0.019)	(0.015)	(0.012)
Trust	-0.005	-0.010	0.005
	(0.018)	(0.013)	(0.012)
Mean of dependent variable	0.478	0.478	0.536
N	390	390	390
Control for beliefs in Sup. 1	No	Yes	—
Control for behavior in Sup. 1	—	—	No
Controls for experience	—	—	—

Table A.6: Effect of personality on behavior and beliefs in Supergame 1

Column 1 (Column 3) reports a linear OLS regression of strategy cooperation (optimism) in Supergame 1 on the five personality factors, controlling for demographic characteristics and standardized cognitive ability (see Section 2.8), and the treatment. The regression reported in Column 2 is the same as the one reported in Column 1, except that it further includes optimism in Supergame 1 as a control. 'Strategy cooperation' is defined in the notes to Table 3 and 'optimism' is defined in the notes to Figure 4. $N = 390$ 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).

	(1)	(2)	(3)
R40	0.039	0.039	0.039
	(0.029)	(0.030)	(0.029)
R48	0.068**	0.060**	0.060**
	(0.026)	(0.025)	(0.025)
Length of Supergame $t-1$	0.004***	0.004***	0.004***
	(0.001)	(0.001)	(0.001)
Own strategy coop in Supergame 1	0.334***	0.334***	0.334***
	(0.036)	(0.036)	(0.036)
Own optimism in Supergame 1	0.232***	0.231***	0.231***
	(0.039)	(0.039)	(0.039)
Other's strategy coop in Supergame $t-1$	0.066***	0.085***	0.082***
	(0.014)	(0.023)	(0.024)
('Other's strategy coop in Sup. $t-1$' $- x_R)_+$		-0.037	-0.035
		(0.039)	(0.040)
Trust × Other's strategy coop in Supergame $t-1$			-0.054**
			(0.020)
Trust × ('Other's strategy coop in Sup. $t-1$' $- x_R)_+$			0.079**
			(0.037)
Mean of dependent variable	0.467	0.467	0.467
N	9360	9360	9360

Table A.7: Strategy cooperation in Supergame t

Notes: The regression reported in Column 1 is exactly the same as the one reported in Column 5 of Panel I of Table 3. The notes to Figure 6 describe how we run a piece-wise linear spline regression by further including $\max\{0, (\text{'Other's strategy coop in Supergame } t-1\text{'} - x_R)\}$, which for conciseness we label here as ('Other's strategy coop in Sup. $t-1$' $- x_R)_+$. Column 2 reports coefficients from this spline regression without interactions with trust, while Column 3 reports coefficients with interactions with trust. All regressions control for the five personality factors (including trust), demographic characteristics and standardized cognitive ability (see Section 2.8), and the supergame number, with $R = 32$ as the omitted category. 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).

Web Appendix VI.2 Robustness of the pattern in Figure 6

Figure A.22 and Table A.8 show that the pattern in Figure 6 in Section 5 is robust when we simplify the underlying regression by excluding the control variables.

Furthermore, we find no statistically significant differences in the pattern across treatments. In particular, when we take the regression reported in Column 3 of Table A.7 in Web Appendix VI.1 and further include interactions of the treatment indicators (R40 and R48) with the interaction of trust and uncooperative evidence (penultimate row of Table A.7) and with the interaction of trust and cooperative evidence (final row of Table A.7), the coefficients on all four of these triple interactions are far from statistical significance (all $p > 0.5$).[46] Although we find no evidence that the pattern varies across treatments, we interpret these high p-values with caution because we are not well powered to identify how trust interacts with cooperative or uncooperative evidence within a particular treatment.

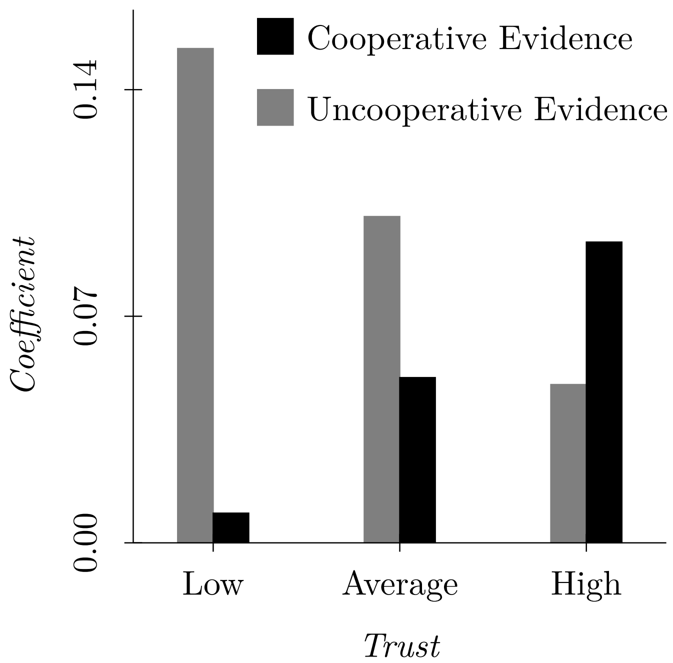

Figure A.22: Effect of opponent's strategy cooperation in Supergame $t-1$ on strategy cooperation in Supergame t (robustness)

Notes: See the notes to Figure 6. Here we use the coefficients from Column 3 of Table A.8 below (instead of Column 3 of Table A.7 in Web Appendix VI.1) that excludes the control variables from the underlying regression.

[46] This regression also includes interactions of the treatment indicators with: (i) trust; (ii) 'Other's strategy coop in Supergame $t-1$'; and (iii) ('Other's strategy coop in Sup. $t-1$' $- x_R)_+$.

	(1)	(2)	(3)
R40	0.100**	0.101**	0.100**
	(0.041)	(0.041)	(0.041)
R48	0.195***	0.185***	0.184***
	(0.033)	(0.034)	(0.034)
Trust	0.022**	0.023**	0.034**
	(0.010)	(0.010)	(0.013)
Other's strategy coop in Supergame $t-1$	0.078***	0.102***	0.101***
	(0.018)	(0.026)	(0.026)
('Other's strategy coop in Sup. $t-1$' $- x_R)_+$		-0.050	-0.050
		(0.045)	(0.047)
Trust × Other's strategy coop in Supergame $t-1$			-0.052**
			(0.024)
Trust × ('Other's strategy coop in Sup. $t-1$' $- x_R)_+$			0.094**
			(0.045)
Mean of dependent variable	0.467	0.467	0.467
N	9360	9360	9360

Table A.8: Strategy cooperation in Supergame t (robustness)

Notes: The regressions reported here are the same as those reported in Table A.7 in Web Appendix VI.1, except that they exclude the control variables (and therefore include only the variables listed above in the table). When we exclude the control variables, we continue to include the treatment indicators because the treatment is strongly correlated with 'Other's strategy coop in Supergame $t-1$'.

Web Appendix VI.3 Simulated dynamics of cooperation

In this appendix, we explore the impact of asymmetric responses to cooperative and uncooperative evidence by more and less trusting subjects (documented in Figure 6). In particular, we carry out two counter-factual simulations. The first simulation consists of sessions populated with agents whose standardized measure of trust is above zero, while the second consists of sessions populated with agents whose standardized measure of trust is below zero. Although both simulations started with the identical distribution of strategy cooperation in Supergame 1, we find that the cooperation trends diverged over the course of the 25 simulated supergames (with sessions populated by trusting subjects achieving high cooperation and sessions populated by non-trusting subjects achieving low cooperation). Figure A.23 presents the evolution of cooperation for the two simulations, and the notes to the figure describe the details of the simulations.

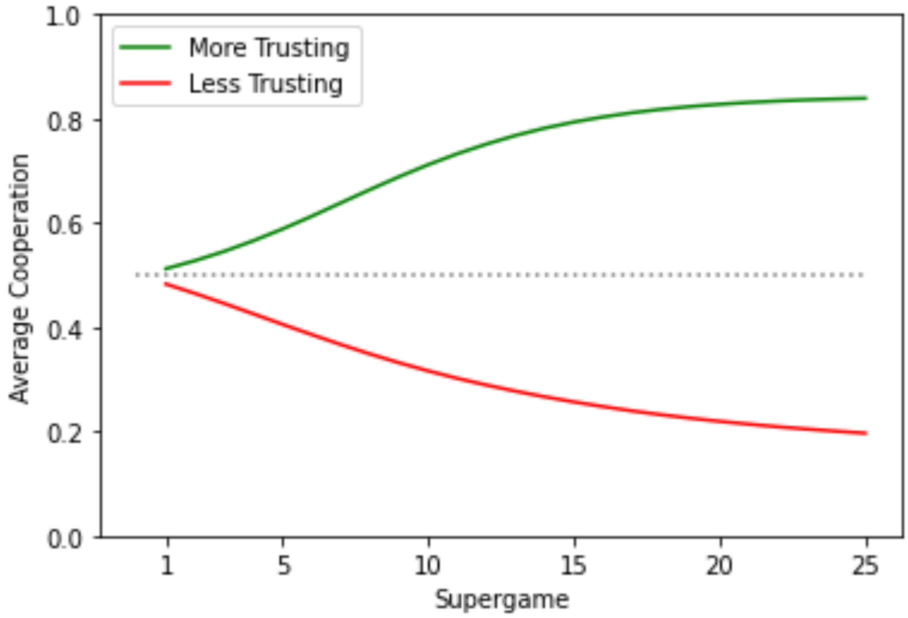

Figure A.23: Simulated dynamics of cooperation in more trusting and less trusting sessions

Notes: Each simulated session included 12 more trusting agents or 12 less trusting agents, and we average over 2,000 simulated sessions. Sessions were populated by drawing independently each agent's trust level from a standard normal distribution (recall from Section 2.8 that the trust factor is standardized by construction), and then randomly allocating agents with trust above (below) zero to a more (less) trusting session. Within session, agents were randomly rematched between supergames. Each agent's Supergame 1 cooperation level was drawn independently from a normal distribution with mean 0.50 and standard deviation 0.15. An agent's cooperation level changed from Supergame $t-1$ to Supergame t as a function of the cooperation level of the agent's opponent in Supergame $t-1$ according to the coefficients in the final four rows of Column 3 of Table A.7, setting $x_R = 0.5$.

Web Appendix VII Comparison to direct-response play

As we describe in the introduction, our experimental design with elicitation of supergame strategies allows us to study initial beliefs about the strategies chosen by others and the evolution of these beliefs with experience. Using this methodology, we replicate important features of the data from indefinitely repeated prisoner's dilemma games with direct-response play (where subjects choose their actions round-by-round instead of choosing a supergame strategy): the next paragraph summarizes this evidence. Unlike the data from direct-response play (e.g., Dal Bó and Fréchette, 2011), in our data cooperation is broadly stable over supergames when the return to joint cooperation is high. As we note in Section 4.1, two features of our design help to explain this finding: (i) subjects could learn about the game during the two forms of training; and (ii) subjects were unable to experiment within supergame given that we elicited supergame strategies. The third and fourth paragraphs of this section discuss these two explanations in detail.

We replicate important features of the data from indefinitely repeated prisoner's dilemma games with direct-response play. For example:

- Dal Bó and Fréchette (2018)'s meta-study finds that cooperation varies in the parameters of the payoff matrix (Table 4). Similarly, we find that cooperation varies in the return to joint cooperation R (Column 1 of Table 2 and Column 1 of Panel I of Table 3, with the effects of $R = 48$ relative to $R = 32$ statistically significant at the one-percent level).

- Dal Bó and Fréchette (2018)'s meta-study (Table 9) finds that cooperation in Round 1 of a supergame depends on the length of the previous supergame (coefficient of 0.006), the Round 1 cooperation of the subject's opponent in the previous supergame (coefficient of 0.12) and the subject's own Round 1 cooperation in the first supergame (coefficient of 0.29). In Table 2, we find effect sizes that are close to those from Dal Bó and Fréchette (2018)'s meta-data, with all three effects statistically significant at the one-percent level.

Learning about the structure of the game and the properties of the supergame strategies during our two forms of training (described in Section 2.5) can help to explain why, in our data, cooperation is broadly stable over supergames when the return to joint cooperation is high. Features of our data suggest that subjects learned during the training phase to vary their cooperation level with the return to joint cooperation R:

- Figure A.24 below shows that subjects varied the frequency with which they tested particular strategies with the return to joint cooperation R.

- In our data, even in the first supergame subjects cooperate more when $R = 40$ than when $R = 32$ (Panel V of Figure 4, Panels (a) and (b) of Figure 5, and Column 1 of Table A.9 in Web Appendix VIII). By contrast, when $\delta = 0.75$ as in our data, with direct-response play

Dal Bó and Fréchette (2011) find a substantially smaller difference that is not statistically significant (top panel of Table 3), while the difference becomes larger and statistically significant after subjects learn from playing the game (bottom panel of Table 3).

- Relatedly, when R changes from 32 to 48, in the first supergame of our experiment subjects vary the frequency with which they choose each of the five key strategies in the same direction as found by Dal Bó and Fréchette (2011, 2019) after subjects learn from playing the game with direct-response play (when $\delta = 0.75$ as in our data).[47]

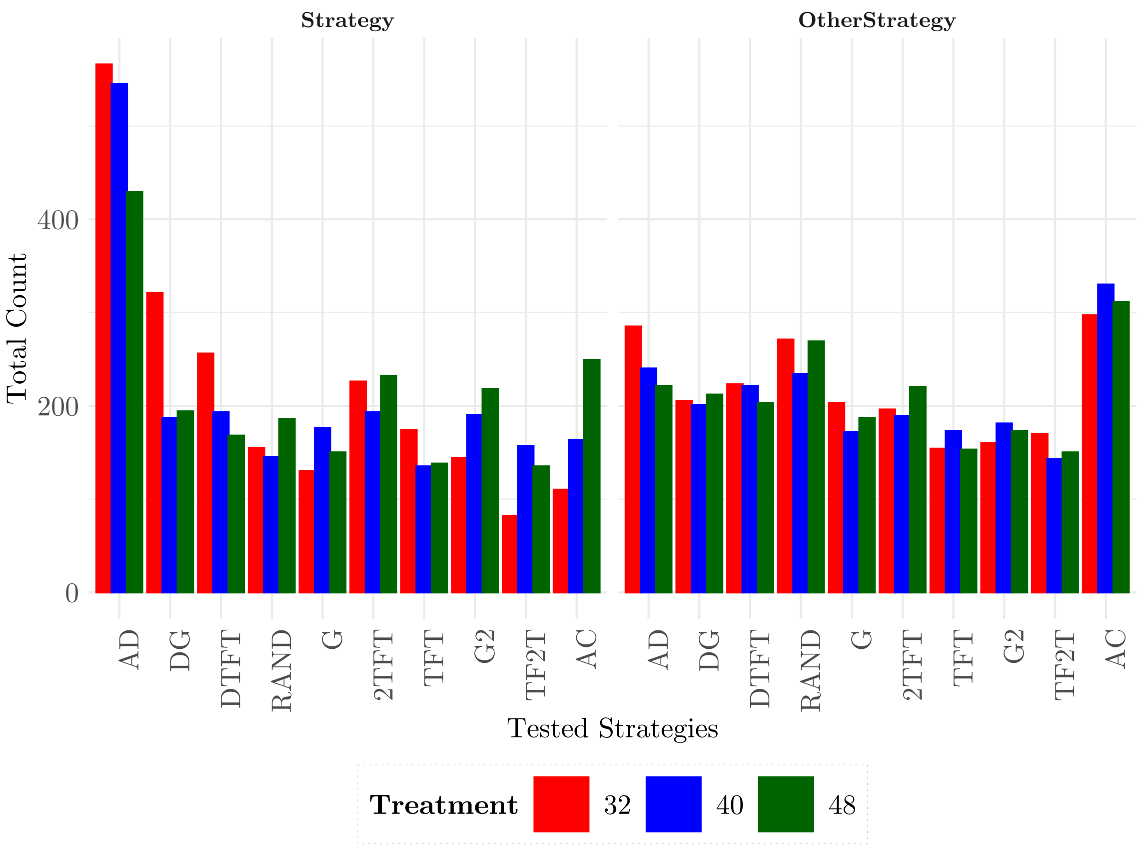

Figure A.24: Number of times each strategy was tested during the training phase

Notes: As described in Section 2.5, in each 'test match' the subject chose one of the ten 'plans' for herself (labeled 'Strategy' here) and one of the ten for the 'other' (labeled 'OtherStrategy' here); see Figure A.6 in Web Appendix I for a screenshot.

[47] Table 10 of Dal Bó and Fréchette (2018) lists the five key strategies and reports the results from Dal Bó and Fréchette (2011, 2019), while Figures 3(a) and 3(c) show our results. In all three cases, the frequency of AC, G and TFT goes up, while the frequency of AD and DTFT goes down (with the exception that the frequency of DTFT does not vary in the data from Dal Bó and Fréchette, 2011).

We elicited supergame strategies, and so subjects were unable to experiment within supergame. Romero and Rosokha (2019a)'s results suggest that this feature of strategy elicitation can further help to explain why, in our data, cooperation is broadly stable over supergames when the return to joint cooperation is high. In particular, Romero and Rosokha (2019a) elicit supergame strategies in the indefinitely repeated prisoner's dilemma, but allow subjects to make strategy adjustments during the course of a supergame. When strategy adjustments are costless, Romero and Rosokha (2019a) find that cooperation increases over supergames, but when strategy adjustments are costly (so that subjects rarely adjust their strategy within supergame), Romero and Rosokha (2019a) find that cooperation tends to be flat over supergames.

Although cooperation remains stable at the aggregate level, we find substantial individual-level learning (see Section 4.2). One aspect pertains to beliefs, whereby subjects whose opponents in Supergames 1 to 24 choose more cooperative strategies end up with more optimistic beliefs (see Panel II of Table 3). At the same time, subjects whose opponents in Supergames 1 to 24 choose less cooperative strategies end up with less optimistic beliefs. The heterogeneous experience and response are consistent with the stable aggregate level of cooperation: even as beliefs become more accurate with experience, considerable heterogeneity in beliefs and cooperation remains in the final supergame (see Web Appendix V.2 and Web Appendix V.6).

Web Appendix VIII Full set of estimates

In order to improve precision, our design included a short matrix reasoning test and four demographic variables that we use as controls in our regressions (Web Appendix III.10 describes the matrix test of cognitive ability; Web Appendix III.11 describes the four demographic variables). Since these variables were not the main focus of the paper, we do not include or discuss the corresponding estimates in the main text. For completeness, Table A.9 reports the full set of estimates for our regressions of cooperation and optimism presented in the main text (we also include Supergame 1 regressions from Table A.6 in Web Appendix VI.1):

- Column 1 of Table A.9 corresponds to Column 1 of Table A.6.
- Column 2 of Table A.9 corresponds to Column 3 of Table A.6.
- Column 3 of Table A.9 corresponds to Column 5 of Table 2.
- Column 4 of Table A.9 corresponds to Column 5 of Panel I of Table 3 and to Column 1 of Table 4.
- Column 5 of Table A.9 corresponds to Column 2 of Table 4.
- Column 6 of Table A.9 corresponds to Column 5 of Panel II of Table 3 and to Column 3 of Table 4.

Table A.9 shows that our matrix reasoning test does not predict cooperation or optimism. Table A.10 further shows that the matrix reasoning test predicts accuracy of beliefs in Supergame 25 but not in Supergame 1, while the matrix test does not predict good responding to beliefs.[48] When interpreting these results, recall that we included only a short eleven-question matrix reasoning test. By contrast, Gill and Prowse (2016)'s study of cognitive ability in the repeated beauty contest game used the much longer Raven test of matrix reasoning that includes 60 questions and took 30 minutes to complete (Fe et al., forthcoming, also use the 60-question Raven test to study how cognitive ability affects strategic behavior in children). Proto et al. (2019, forthcoming) find that cognitive ability predicts cooperation: because we elicit strategies, our design turns off the main mechanism that explains how cognitive ability affects cooperation in Proto et al. (2019, forthcoming), namely that more intelligent subjects make fewer errors in implementing their strategy.

Turning to the demographic variables, the results in Table A.9 provide some evidence that STEM majors, older subjects and males cooperate more; we interpret these effects with caution

[48] To help interpret the effect size, note that our measure of accuracy is defined from -1 to 0 (see Web Appendix IV.3). Since subjects learn within their session, in Supergame 25 we define accuracy at the session level instead of the treatment level: if we define accuracy of beliefs at the session level also in Supergame 1, the effect of the matrix test on accuracy in Supergame 1 remains statistically insignificant.

since the effects on cooperation are not always statistically significant at the 5% level and the effects on beliefs are never statistically significant at the 5% level. The effects of having attended an American high school are never statistically significant at the 5% level.

Dependent Variable:	(1) Strategy cooperation	(2) Optimism	(3) Round 1 cooperation	(4) Strategy cooperation	(5) Strategy cooperation	(6) Optimism
Data from Supergames:	**1**	**1**	**2-25**	**2-25**	**21-25**	**25**
R40	0.105**	0.098***	0.040	0.039	0.067**	-0.016
	(0.044)	(0.028)	(0.043)	(0.029)	(0.031)	(0.024)
R48	0.280***	0.146***	0.113***	0.068**	0.084**	0.005
	(0.041)	(0.026)	(0.038)	(0.026)	(0.034)	(0.035)
Length of Supergame $t-1$			0.007***	0.004***	0.008***	
			(0.001)	(0.001)	(0.003)	
Length of Supergames 1 to 24						0.025*
						(0.013)
Other's Round 1 coop in Supergame $t-1$			0.080***			
			(0.016)			
Other's strategy coop in Supergame $t-1$				0.066***	0.076***	
				(0.014)	(0.018)	
Others' strategy coop in Supergames 1 to 24						0.760***
						(0.097)
Own Round 1 coop in Supergame 1			0.299***			
			(0.035)			
Own strategy coop in Supergame 1				0.334***	0.298***	0.069**
				(0.036)	(0.046)	(0.033)
Own optimism in Supergame 1			0.345***	0.232***	0.196***	0.411***
			(0.061)	(0.039)	(0.055)	(0.040)
Supergame number			-0.000	-0.000	-0.009***	
			(0.001)	(0.001)	(0.003)	
Anxiety	0.009	-0.008	0.002	-0.004	-0.022	0.003
	(0.019)	(0.011)	(0.013)	(0.009)	(0.014)	(0.010)
Cautiousness	-0.002	-0.002	-0.008	-0.009	-0.010	-0.008
	(0.017)	(0.012)	(0.012)	(0.008)	(0.012)	(0.012)
Kindness	-0.022	-0.017	0.003	0.008	0.001	0.017*
	(0.017)	(0.011)	(0.013)	(0.011)	(0.015)	(0.008)
Manipulativeness	-0.027	-0.020*	-0.001	0.004	0.006	-0.003
	(0.019)	(0.012)	(0.014)	(0.010)	(0.013)	(0.011)
Trust	-0.005	0.005	0.024**	0.021**	0.022**	0.017**
	(0.018)	(0.012)	(0.010)	(0.009)	(0.009)	(0.007)
Matrix Test (standardized)	0.014	-0.011	0.009	0.003	-0.001	-0.011
	(0.019)	(0.012)	(0.012)	(0.010)	(0.012)	(0.010)
STEM Major	0.048	0.045*	0.029	0.035**	0.040*	0.014
	(0.038)	(0.024)	(0.024)	(0.016)	(0.023)	(0.016)
Age Under 20	-0.066*	-0.005	-0.072***	-0.066***	-0.070**	-0.029
	(0.037)	(0.022)	(0.024)	(0.016)	(0.029)	(0.022)
High School in USA	-0.056	-0.028	-0.046*	-0.015	-0.008	0.005
	(0.039)	(0.027)	(0.023)	(0.017)	(0.022)	(0.019)
Male	0.027	-0.017	0.083***	0.038*	0.007	-0.019
	(0.040)	(0.024)	(0.027)	(0.021)	(0.026)	(0.019)
Intercept	0.352***	0.448***	0.001	0.094***	0.292***	-0.267***
	(0.045)	(0.028)	(0.037)	(0.029)	(0.072)	(0.063)
Mean of dependent variable	0.478	0.536	0.464	0.467	0.461	0.430
N	390	390	9360	9360	1950	390

Table A.9: Cooperation and optimism regressions: full set of estimates

Notes: The bullet points at the end of the first paragraph of Web Appendix VIII match each column to the corresponding regression reported in the paper: the notes to those regressions describe the regression and the variables. Section 2.8.1 describes the five personality factors. Web Appendix III.10 describes the matrix test of cognitive ability. Web Appendix III.11 describes the four demographic variables.

	(1)	(2)	(3)	(4)
Dependent Variable:	Belief Accuracy (vs. Treatment)	Belief Accuracy (vs. Session)	Good Responder To Beliefs	Good Responder To Beliefs
Data from Supergames:	**1**	**25**	**1**	**25**
R40	0.022	0.009	0.066	-0.272***
	(0.019)	(0.018)	(0.062)	(0.070)
R48	0.102***	0.012	-0.111*	-0.387***
	(0.018)	(0.014)	(0.062)	(0.069)
Anxiety	-0.006	-0.002	0.015	0.002
	(0.008)	(0.009)	(0.027)	(0.029)
Cautiousness	0.014*	0.004	0.028	0.001
	(0.008)	(0.009)	(0.025)	(0.020)
Kindness	0.003	-0.007	-0.032	-0.042
	(0.007)	(0.006)	(0.026)	(0.025)
Manipulativeness	-0.005	-0.002	0.027	-0.019
	(0.008)	(0.005)	(0.026)	(0.025)
Trust	0.002	0.007	-0.043*	-0.029
	(0.008)	(0.008)	(0.025)	(0.021)
Matrix Test (standardized)	-0.011	0.021***	0.020	0.000
	(0.008)	(0.007)	(0.026)	(0.029)
STEM Major	-0.001	-0.010	-0.085	-0.044
	(0.016)	(0.012)	(0.054)	(0.047)
Age Under 20	0.026	0.023	-0.068	0.013
	(0.016)	(0.014)	(0.052)	(0.051)
High School in USA	-0.026	0.000	0.052	0.030
	(0.018)	(0.011)	(0.056)	(0.059)
Male	0.006	-0.020	0.069	0.065
	(0.016)	(0.020)	(0.054)	(0.063)
Intercept	-0.274***	-0.201***	0.541***	0.691***
	(0.019)	(0.016)	(0.066)	(0.063)
Mean of dependent variable	-0.228	-0.202	0.500	0.497
N	390	390	390	390

Table A.10: Accuracy and good responding: full set of estimates

Notes: Each column reports a linear OLS regression, with $R = 32$ as the omitted category. Accuracy of beliefs at the treatment level is defined in Web Appendix IV.3; we use this measure in Column 1, while in Column 2 we use the same measure but calculated at the session level (because subjects learn within their session). Good responding to beliefs is a binary variable defined in the second paragraph of Section 3.3. Section 2.8.1 describes the five personality factors. Web Appendix III.10 describes the matrix test of cognitive ability. Web Appendix III.11 describes the four demographic variables. $N = 390$ because four subjects did not complete the demographic questionnaire. Heteroskedasticity-robust standard errors are shown in parentheses (with clustering at the session level for the second and fourth columns). ***, ** and * denote significance at the 1 percent, 5 percent and 10 percent levels (two-sided tests).

Web Appendix IX Additional figures

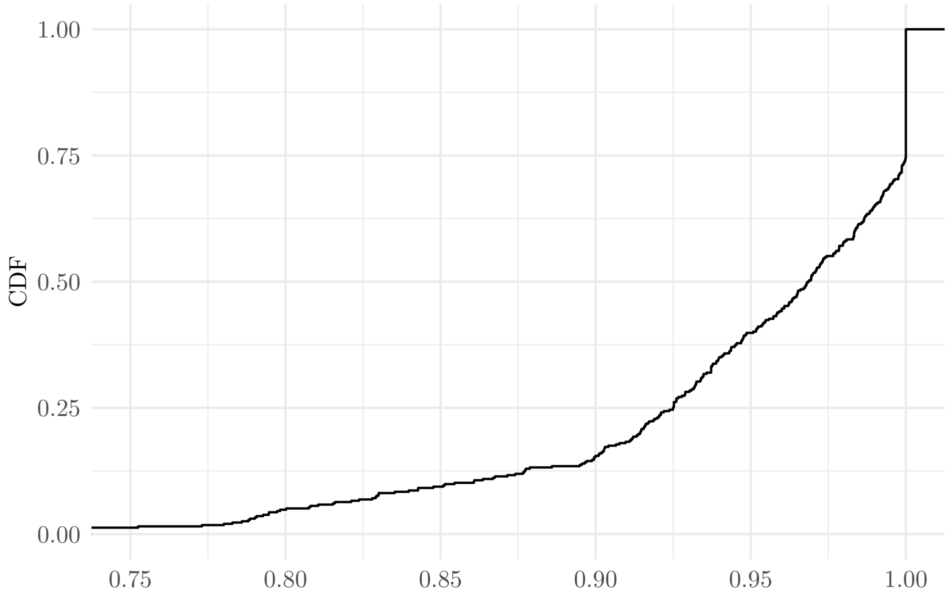

Figure A.25: CDF of payoffs relative to best responding to beliefs in Supergame 1

Notes: The figure shows the cumulative distribution function across subjects of the expected payoff of the subject's chosen strategy given her beliefs as a proportion of the expected payoff of the best response to her beliefs.

Web Appendix X Additional tables

	AD	*DG*	*DTFT*	*RAND*	*G*	*2TFT*	*TFT*	*G2*	*TF2T*	*AC*
AD	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
DG	0.0000	0.0000	0.0000	0.1500	0.1875	0.1875	0.1875	0.7500	0.7500	0.7500
DTFT	0.0000	0.0000	0.0000	0.3750	0.1875	0.1875	0.4286	0.7500	0.7500	0.7500
RAND	0.5000	0.5000	0.5000	0.5000	0.5000	0.5000	0.5000	0.5000	0.5000	0.5000
G	0.2500	0.2500	0.2500	0.4000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
2TFT	0.2500	0.2500	0.2500	0.4844	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
TFT	0.2500	0.3906	0.5714	0.6250	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
G2	0.4375	1.0000	1.0000	0.7097	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
TF2T	0.4375	1.0000	1.0000	0.8594	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
AC	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000

Table A.11: Cooperation rates

Notes: Each cell reports the cooperation rate of the strategy on the vertical axis when playing against the strategy on the horizontal axis. Cooperation rates measure the expected number of rounds of cooperation divided by the expected number of rounds, and are based on analytical calculations.

R	*AD*	*DG*	*DTFT*	*RAND*	*G*	*2TFT*	*TFT*	*G2*	*TF2T*	*AC*
32	0.91	0.02	0.05	0.01	0.01	0.01	0.02	0.02	0.02	0.02
40	0.23	0.08	0.04	0.00	0.05	0.01	0.01	0.67	0.15	0.05
48	0.04	0.03	0.01	0.00	0.10	0.04	0.05	0.85	0.18	0.04

Table A.12: Frequency each strategy is a best response in Supergame 1

Notes: This table replicates Table A.1, but using best responding instead of good responding.

	R	*AD*	*DG*	*DTFT*	*RAND*	*G*	*2TFT*	*TFT*	*G2*	*TF2T*	*AC*
I: All	32	0.52	0.82	0.78	0.50	0.78	0.76	0.88	0.71	0.75	0.67
	40	0.76	1.00	1.00	1.00	0.71	0.72	0.70	0.55	0.55	0.61
	48	0.56	0.42	0.42	0.00	0.67	0.68	0.71	0.41	0.41	0.51
II: Pessimists	32	0.80	0.88	0.86	–	0.62	0.62	0.86	0.50	0.50	0.33
	40	0.84	1.00	1.00	–	0.62	0.64	0.57	0.34	0.34	0.12
	48	1.00	0.71	0.71	–	0.41	0.43	0.47	0.22	0.22	0.16
III: Optimists	32	0.26	0.71	0.67	0.50	0.90	0.89	0.89	0.88	0.88	1.00
	40	0.50	1.00	1.00	1.00	0.77	0.76	0.76	0.68	0.68	0.69
	48	0.00	0.00	0.00	0.00	0.91	0.91	0.91	0.60	0.60	0.61

Table A.13: Probability of good responding in Supergame 1
(conditional on each strategy being a good response)

Notes: For each strategy, the table shows the proportion of subjects who good respond to their beliefs, conditional on that strategy being in the set of good responses to the subject's beliefs. Good responding is defined in the second paragraph of Section 3.3. Panels II and III split subjects according to whether they are above or below the median level of optimism within-treatment, where optimism is defined in the notes to Figure 4 (subjects at the median are allocated to the above-median category).

	AD	*DG*	*DTFT*	*RAND*	*G*	*2TFT*	*TFT*	*G2*	*TF2T*	*AC*
AD	100.00	100.00	100.00	150.00	125.00	125.00	125.00	143.75	143.75	200.00
DG	100.00	100.00	100.00	140.70	115.25	115.25	129.31	146.00	146.00	146.00
DTFT	100.00	100.00	100.00	126.75	115.25	115.25	134.86	146.00	146.00	146.00
RAND	74.00	87.50	107.75	119.00	110.00	117.59	130.25	137.87	151.34	164.00
G	87.00	105.75	105.75	125.20	128.00	128.00	128.00	128.00	128.00	128.00
2TFT	87.00	105.75	105.75	119.97	128.00	128.00	128.00	128.00	128.00	128.00
TFT	87.00	98.44	113.14	111.25	128.00	128.00	128.00	128.00	128.00	128.00
G2	77.25	108.00	108.00	106.00	128.00	128.00	128.00	128.00	128.00	128.00
TF2T	77.25	108.00	108.00	96.72	128.00	128.00	128.00	128.00	128.00	128.00
AC	48.00	108.00	108.00	88.00	128.00	128.00	128.00	128.00	128.00	128.00

Table A.14: Payoffs when $R = 32$

Notes: Each cell reports the expected payoff in points of the strategy on the vertical axis when playing against the strategy on the horizontal axis. Expected payoffs are based on analytical calculations.

	AD	*DG*	*DTFT*	*RAND*	*G*	*2TFT*	*TFT*	*G2*	*TF2T*	*AC*
AD	100.00	100.00	100.00	150.00	125.00	125.00	125.00	143.75	143.75	200.00
DG	100.00	100.00	100.00	143.10	115.25	115.25	129.31	170.00	170.00	170.00
DTFT	100.00	100.00	100.00	132.75	115.25	115.25	134.86	170.00	170.00	170.00
RAND	74.00	89.90	113.75	127.00	116.40	125.34	140.25	149.23	165.09	180.00
G	87.00	105.75	105.75	131.60	160.00	160.00	160.00	160.00	160.00	160.00
2TFT	87.00	105.75	105.75	127.72	160.00	160.00	160.00	160.00	160.00	160.00
TFT	87.00	98.44	113.14	121.25	160.00	160.00	160.00	160.00	160.00	160.00
G2	77.25	132.00	132.00	117.35	160.00	160.00	160.00	160.00	160.00	160.00
TF2T	77.25	132.00	132.00	110.47	160.00	160.00	160.00	160.00	160.00	160.00
AC	48.00	132.00	132.00	104.00	160.00	160.00	160.00	160.00	160.00	160.00

Table A.15: Payoffs when $R = 40$

Notes: Each cell reports the expected payoff in points of the strategy on the vertical axis when playing against the strategy on the horizontal axis. Expected payoffs are based on analytical calculations.

	AD	*DG*	*DTFT*	*RAND*	*G*	*2TFT*	*TFT*	*G2*	*TF2T*	*AC*
AD	100.00	100.00	100.00	150.00	125.00	125.00	125.00	143.75	143.75	200.00
DG	100.00	100.00	100.00	145.50	115.25	115.25	129.31	194.00	194.00	194.00
DTFT	100.00	100.00	100.00	138.75	115.25	115.25	134.86	194.00	194.00	194.00
RAND	74.00	92.30	119.75	135.00	122.80	133.09	150.25	160.58	178.84	196.00
G	87.00	105.75	105.75	138.00	192.00	192.00	192.00	192.00	192.00	192.00
2TFT	87.00	105.75	105.75	135.47	192.00	192.00	192.00	192.00	192.00	192.00
TFT	87.00	98.44	113.14	131.25	192.00	192.00	192.00	192.00	192.00	192.00
G2	77.25	156.00	156.00	128.71	192.00	192.00	192.00	192.00	192.00	192.00
TF2T	77.25	156.00	156.00	124.22	192.00	192.00	192.00	192.00	192.00	192.00
AC	48.00	156.00	156.00	120.00	192.00	192.00	192.00	192.00	192.00	192.00

Table A.16: Payoffs when $R = 48$

Notes: Each cell reports the expected payoff in points of the strategy on the vertical axis when playing against the strategy on the horizontal axis. Expected payoffs are based on analytical calculations.

	(1)	(2)
R40	12.79***	12.58***
	(0.49)	(0.52)
R48	48.05***	46.99***
	(0.84)	(0.94)
Accuracy of beliefs		10.01***
		(2.71)
Proximity to best response	88.49***	87.14***
	(7.20)	(7.43)
Mean of dependent variable	128.15	128.15
N	390	390

Table A.17: Expected earnings in Supergame 1 (robustness)

Notes: The regressions reported here are the same as those reported in Columns 2 and 3 of Table A.2 except that we replace 'Good responder to beliefs' with 'Proximity to best response', which measures the expected payoff of a subject's chosen strategy given her beliefs as a proportion of the expected payoff of the best response to her beliefs. To help interpret the effect size, note that our measure of proximity to the best response is defined from 0 to 1. Figure A.25 in Web Appendix IX shows the cumulative distribution function of proximity.

	(1)	(2)	(3)	(4)	(5)
R40	0.148***	0.135***	0.064	0.077	0.035
	(0.056)	(0.051)	(0.046)	(0.053)	(0.042)
R48	0.300***	0.275***	0.154***	0.199***	0.111***
	(0.042)	(0.038)	(0.043)	(0.045)	(0.038)
Length of Supergame $t-1$	0.007***				0.006***
	(0.001)				(0.001)
Other's Round 1 coop	0.087***				0.079***
in Supergame $t-1$		(0.019)			(0.016)
Own Round 1 coop			0.385***		0.286***
in Supergame 1			(0.036)		(0.034)
Own optimism				0.691***	0.353***
in Supergame 1				(0.073)	(0.064)
Mean of dependent variable	0.464	0.464	0.464	0.464	0.464
N	9360	9360	9360	9360	9360

Table A.18: Round 1 cooperation in Supergame t (robustness)

Notes: This table replicates Table 2, except that we use Probit regressions instead of linear OLS regressions. The table reports average marginal effects and heteroskedasticity-robust standard errors with clustering at the session level.

I: Strategy cooperation in Supergame t

	(1)	(2)	(3)	(4)	(5)
R40	0.103**	0.094**	0.059*	0.049	0.037
	(0.041)	(0.038)	(0.032)	(0.037)	(0.029)
R48	0.209***	0.191***	0.090***	0.128***	0.065**
	(0.033)	(0.030)	(0.029)	(0.029)	(0.025)
Length of Supergame $t-1$	0.004***				0.004***
	(0.001)				(0.001)
Other's inferred coop in Supergame $t-1$		0.095***			0.088***
		(0.021)			(0.017)
Own strategy coop in Supergame 1			0.422***		0.333***
			(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**	0.011	0.086*	0.062	-0.028
	(0.045)	(0.023)	(0.042)	(0.043)	(0.023)
R48	0.244***	0.049	0.178***	0.171***	-0.014
	(0.043)	(0.033)	(0.045)	(0.041)	(0.033)
Length of Supergames 1 to 24	0.040				0.024**
	(0.024)				(0.012)
Others' inferred coop in Supergames 1 to 24		1.027***			0.947***
		(0.122)			(0.110)
Own strategy coop in Supergame 1			0.232***		0.061*
			(0.036)		(0.033)
Own optimism in Supergame 1				0.493***	0.419***
				(0.042)	(0.042)
Mean of dependent variable	0.430	0.430	0.430	0.430	0.430
N	390	390	390	390	390

Table A.19: Effect of experience on behavior and beliefs (robustness)

Notes: The regressions reported in Panel I are the same as those reported in Panel I of Table 3, except that we replace 'Other's strategy coop in Supergame $t-1$' with 'Other's inferred coop in Supergame $t-1$'. The regressions reported in Panel II are the same as those reported in Panel II of Table 3, except that we replace 'Others' strategy coop in Supergames 1 to 24' with 'Others' inferred coop in Supergames 1 to 24'. 'Other's inferred coop in Supergame t' measures the weighted average of the strategy cooperation of each of the ten strategies, where the weights come from the posterior distribution over the opponent's strategies after Bayesian updating from the realized sequence of play in Supergame t; the Bayesian update for Supergame t uses the uniform distribution as the prior, and so the Bayesian update for Supergame t does not use information from sequences of play in prior supergames. See the notes to Table 3 for the definition of strategy cooperation. 'Others' inferred coop in Supergames 1 to 24' is the mean of 'Other's inferred coop in Supergame t' over the first 24 supergames.

I: Strategy cooperation in Supergame t

	(1)	(2)	(3)	(4)	(5)
R40	0.110**	0.102**	0.064*	0.057	0.044
	(0.045)	(0.042)	(0.034)	(0.042)	(0.033)
R48	0.210***	0.193***	0.086***	0.130***	0.064**
	(0.036)	(0.033)	(0.030)	(0.032)	(0.027)
Length of Supergame $t-1$	0.003***				0.004***
	(0.001)				(0.001)
Other's strategy coop in Supergame $t-1$		0.079***			0.071***
		(0.019)			(0.015)
Own strategy coop in Supergame 1			0.434***		0.348***
			(0.035)		(0.034)
Own optimism in Supergame 1				0.564***	0.225***
				(0.053)	(0.041)
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**	0.023	0.085*	0.064	-0.016
	(0.047)	(0.023)	(0.043)	(0.045)	(0.023)
R48	0.246***	0.068**	0.178***	0.174***	0.006
	(0.044)	(0.031)	(0.046)	(0.041)	(0.033)
Length of Supergames 1 to 24	0.037				0.021
	(0.024)				(0.014)
Others' strategy coop in Supergames 1 to 24		0.844***			0.771***
		(0.109)			(0.097)
Own strategy coop in Supergame 1			0.234***		0.066*
			(0.037)		(0.033)
Own optimism in Supergame 1				0.497***	0.417***
				(0.042)	(0.036)
Mean of dependent variable	0.430	0.430	0.430	0.430	0.430
N	390	390	390	390	390

Table A.20: Effect of experience on behavior and beliefs (robustness)

Notes: The regressions reported here are the same as those reported in Table 3, except that the regressions here do not include controls for the personality factors, demographic characteristics, or cognitive ability.

Unfriendly, provocable and lenient strategies

- Unfriendly: defects in the very first round.
- Provocable: starts by cooperating but defects immediately in response to the opponent's first defection.
- Lenient: starts by cooperating and does not defect immediately in response to the opponent's first defection.
- See Web Appendix III.3 for further details.

Optimism, Cooperativeness and OptimismRelTruth

- Optimism: how often a subject expects others to cooperate when they play against each other (measured by the expected cooperation rate of a subject's belief distribution playing against itself).
- Cooperativeness: how often a subject expects her chosen strategy to cooperate (measured by the expected cooperation rate of the chosen strategy playing against the subject's belief distribution).
- OptimismRelTruth: optimism relative to how often others cooperate (measured by the difference between optimism and the expected cooperation rate of the treatment-level strategy distribution (excluding the subject's own choice) playing against itself).
- See the notes to Figure 4 for further details.

Good responding

- 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).
- See the second paragraph of Section 3.3 for further details.

Strategy cooperation

- How much a strategy cooperates on average against a uniform distribution over the available strategies (measured by the expected cooperation rate of the strategy playing against the uniform distribution over the ten available strategies).
- See the notes to Table 3 for further details.

Cooperative evidence and uncooperative evidence

- A subject receives 'cooperative evidence' ('uncooperative evidence') when her opponent's strategy cooperation in the previous supergame was above (below) the treatment-level median.
- See the notes to Figure 6 for further details.

Table A.21: Summary of main definitions