

It's the thought that counts:
The role of intentions in noisy repeated games*

David G. Rand, Drew Fudenberg and Anna Dreber

First version: April 13 2013
This version: January 7 2014

Abstract:

We examine cooperation in repeated interactions where intended actions are implemented with noise but intentions are perfectly observable. Observable intentions lead to more cooperation compared to control games where intentions are unobserved, allowing subjects to reach similar cooperation levels as in games without noise. Most subjects condition exclusively on intentions, and use simpler, lower-memory strategies compared to games where intentions are unobservable. When the returns to cooperation are high, some subjects are tolerant, using good outcomes to forgive attempted defections; when the returns to cooperation are low, some subjects are punitive, using bad outcomes to punish accidental defections.

* Rand: Department of Psychology, Department of Economics, Cognitive Science Program, School of Management, Yale University, Box 208205, New Haven CT 06520-8205 (e-mail: david.rand@yale.edu); Fudenberg: Department of Economics, Harvard University, 1805 Cambridge Street, Cambridge, MA 02138 (e-mail: dfudenberg@harvard.edu); Dreber: Department of Economics, Stockholm School of Economics, Box 6501, 113 83 Stockholm, Sweden (e-mail: anna.dreber@hhs.se). Rand and Fudenberg are co-lead authors. National Science Foundation Grant SES-0954162, the John Templeton Foundation and the Jan Wallander and Tom Hedelius Foundation provided financial support. We are grateful for comments from Yoella Bereby-Meyer, Fiery Cushman, Guillaume Frechette, Moshe Hoffman, Alexander Peysakhovich, Bjørn-Atle Reme, Madison Storm, and Sevgi Yuksel, and we thank Yoella Bereby-Meyer and Al Roth for sharing their data with us.

1. Introduction

This paper studies cooperation in infinitely repeated games where the intended actions are implemented with error, so that the actions played are only a noisy or implicit signal of what was intended. The possibility of error is pervasive in social interactions, and many if not most of these interactions do not have a fixed and known termination date. The resulting imperfect public monitoring has received a large amount of attention in the theoretical literature on infinitely repeated games (e.g., Green and Porter 1984, Radner et al. 1986, Abreu et al. 1990, Fudenberg et al. 1994), but only a handful of experimental studies have explored infinitely repeated games with errors (e.g. Aoyagi and Frechette 2009, Bigoni et al. 2012, Fudenberg et al. 2012, Aoyagi et al. 2013).

Our setup differs from that of these past studies in that we consider the effect of players directly observing the intended actions of their opponents, in addition to the realized ones. This sort of information is available in some real-world settings, for example compensation for hedge fund managers where both the positions taken and the actual outcomes are observable and thus explicit, or in a homicide when it is clear that the accused shot the victim but extenuating circumstances may exist – here the legal system pays attention to both intentions and outcomes, differentiating between manslaughter and various levels of murder.

From a theoretical standpoint, the impact of explicitly observing intentions is clear: the highest equilibrium payoff can be obtained with strategies that completely ignore the realized outcomes and condition only on intended play, and moreover this best equilibrium is the same as when actions are implemented without error. Note that this is very different from the situation in one-shot games, where maximizing monetary payoffs would lead subjects to ignore intentions entirely. Even in those games, a substantial proportion of subjects do condition on intentions in addition to outcomes when both pieces of information are available.¹ One possible explanation for this apparent “preference for reciprocity” is that it reflects a heuristic that fosters cooperation in repeated interactions. If so, we might expect to see even more reliance on intentions in settings where conditioning on intentions leads to a cooperative equilibrium even in the absence of a preference for reciprocity. At an empirical level, the question of how extensively people

¹ Past work on intentions in one-shot games is discussed in Section 2, as well as the Bereby-Meyer and Roth (2006) and Kunreuther et al. (2009) studies of intentions in the finitely-repeated prisoner’s dilemma.

condition their play on intentions in infinitely repeated games remains open, as does the extent to which they also condition on outcomes, and the effect of all this on the level of cooperation.

To begin to understand these issues, we study the experimental play of the repeated prisoner's dilemma when intended actions are implemented with error. Our main goals are to understand when and in what ways subjects use data on intentions and outcomes, and how cooperation when intentions are revealed compares to either a setting with error when intentions are not observed, or one in which error is not present (so the actions themselves reveal the intentions). We present evidence from a set of infinitely repeated prisoner's dilemma games with a continuation probability of $7/8$ and an error rate of $1/8$. In our main treatments, intentions are explicit; as controls, we also consider the same games but where only actions are observable (thus leaving intentions implicit), as well as the same games without exogenously imposed error (where the observed action corresponds to the intended one). We explore two different payoff specifications for the stage game actions "Cooperate" ("C") and "Defect" ("D") (neutral language was used in the experiment itself). In the "high benefit" treatment, the benefit that playing "C" gives to the other player is high enough that there is a cooperative equilibrium in the game with errors whether or not intentions are observed. In the "low benefit" treatment, the benefit that C gives is low enough that the only equilibrium with errors and unobserved intentions is for both players to always defect, although cooperation remains an equilibrium outcome when intentions are observed.

Summary of results

We use two different methods to analyze the data: a structural estimation of the distribution of strategies using the "structural frequency estimation method" (SFEM) of Dal Bó and Frechette (2011), and a descriptive analysis that relates play in a given period of a supergame to the opponent's intention and action in the period before (which implicitly assumes subjects use strategies that mostly depend on that information). Both methods show that most subjects condition almost exclusively on intentions and thus play consistently with predictions based on maximizing money payoffs: In our descriptive analysis, the effect of opponent's intention is dramatically larger than that of the actual outcome. Similarly, in the strategy estimation, more than two thirds of subjects use strategies that do not condition on outcomes.

To the extent that subjects do condition on outcomes, interestingly, they do so in different ways depending on the payoff specification. In the treatment where there is less of an incentive

to cooperate, some people (about 15%) are punitive in treating both accidental cooperation (partner meant to play D but played C) and accidental defection (partner meant to play C but played D) as defection; only when the partner both intended to play C and actually did so was this treated as cooperation. This behavior is not observed in the treatment with high returns to cooperation, where instead some people (about 19%) are tolerant in that they only retaliate against intentional defections – these subjects forgive both accidental defection (partner meant to play C but played D) as well as accidental cooperation (partner meant to play D but played C). Thus the “punitive” subjects in the low-benefit treatment use realized outcomes to punish cooperators that defect by accident, while in the high-benefit treatment “tolerant” subjects use the realized outcomes to forgive defectors that accidentally cooperated.

By conditioning largely on intentions, subjects are able to achieve high levels of cooperation in both treatments. Compared with the controls in which intentions are implicit, explicitly revealing intentions lead to significantly more cooperation. Interestingly, this increase in cooperation is not associated with more leniency (where the subject overlooks the partner’s first defection) but instead with an increase in simple strategies that conditioned on at most the previous period. This suggests that many of the longer memory strategies seen in Fudenberg et al. (2012) were the result of subjects trying to infer the intentions of their opponent, either because doing so leads to higher monetary payoffs or because preferences depend on the intentions of others.

In principle, games with errors but explicit intentions are distinct from games with no errors, so people might use different strategies in each. To evaluate this possibility, we compare play when intentions are explicit to play in games where there are no errors.² We find that not only does revealing intentions increase cooperation compared to games where intentions are implicit, but it successfully moves cooperation levels all the way back up to the level seen in the absence of errors. Revealing intentions also leads subjects to use similar strategies to those observed in the absence of errors; in particular in both cases most subjects condition only on play in previous period.

² Previous work on infinitely repeated games without errors has shown that subjects learn to cooperate, as long as the returns on cooperation are large enough relative to the continuation probability (Dal Bó 2005, Dreber et al. 2008, Dal Bó & Frechette 2011, Fudenberg et al. 2012, Rand & Nowak 2013). Furthermore, cooperation is significantly higher without errors compared to the case with errors where intentions are implicit (Fudenberg et al. 2012).

2. Experimental Design and Empirical Questions

In our experiments, the infinitely repeated prisoner’s dilemma was induced by having a known constant probability of $\delta = 7/8$ that a supergame would continue between two players following each period; with probability $1-\delta$, the supergame ended and subjects were informed that they have been re-matched with a new partner. In the main treatments, there was also a known constant error probability of $E=1/8$ that an intended move is changed to the opposite move. In this “explicit intentions” treatment, subjects were informed of the intended action of the other player, the other player’s realized action, and whether their own move had been changed (i.e. when they make an error). We also have an “implicit intentions” control, where subjects were told their own realized action and the realized action of the other player but not the other player’s intended action. Finally, we have a set of control conditions without errors.

Realized payoffs	Expected payoffs																		
$b/c = 1.5$ <table style="margin-left: auto; margin-right: auto; border-collapse: collapse;"> <thead> <tr> <th style="padding: 5px;"></th> <th style="padding: 5px; text-align: center;">C</th> <th style="padding: 5px; text-align: center;">D</th> </tr> </thead> <tbody> <tr> <th style="padding: 5px; text-align: left;">C</th> <td style="padding: 5px; text-align: center;">1,1</td> <td style="padding: 5px; text-align: center;">-2,3</td> </tr> <tr> <th style="padding: 5px; text-align: left;">D</th> <td style="padding: 5px; text-align: center;">3,-2</td> <td style="padding: 5px; text-align: center;">0,0</td> </tr> </tbody> </table>		C	D	C	1,1	-2,3	D	3,-2	0,0	$b/c = 1.5, E = 1/8$ <table style="margin-left: auto; margin-right: auto; border-collapse: collapse;"> <thead> <tr> <th style="padding: 5px;"></th> <th style="padding: 5px; text-align: center;">C</th> <th style="padding: 5px; text-align: center;">D</th> </tr> </thead> <tbody> <tr> <th style="padding: 5px; text-align: left;">C</th> <td style="padding: 5px; text-align: center;">0.875, 0.875</td> <td style="padding: 5px; text-align: center;">-1.375, 2.375</td> </tr> <tr> <th style="padding: 5px; text-align: left;">D</th> <td style="padding: 5px; text-align: center;">2.375, -1.375</td> <td style="padding: 5px; text-align: center;">0.125, 0.125</td> </tr> </tbody> </table>		C	D	C	0.875, 0.875	-1.375, 2.375	D	2.375, -1.375	0.125, 0.125
	C	D																	
C	1,1	-2,3																	
D	3,-2	0,0																	
	C	D																	
C	0.875, 0.875	-1.375, 2.375																	
D	2.375, -1.375	0.125, 0.125																	
$b/c = 4$ <table style="margin-left: auto; margin-right: auto; border-collapse: collapse;"> <thead> <tr> <th style="padding: 5px;"></th> <th style="padding: 5px; text-align: center;">C</th> <th style="padding: 5px; text-align: center;">D</th> </tr> </thead> <tbody> <tr> <th style="padding: 5px; text-align: left;">C</th> <td style="padding: 5px; text-align: center;">6,6</td> <td style="padding: 5px; text-align: center;">-2,8</td> </tr> <tr> <th style="padding: 5px; text-align: left;">D</th> <td style="padding: 5px; text-align: center;">8,-2</td> <td style="padding: 5px; text-align: center;">0,0</td> </tr> </tbody> </table>		C	D	C	6,6	-2,8	D	8,-2	0,0	$b/c = 4, E = 1/8$ <table style="margin-left: auto; margin-right: auto; border-collapse: collapse;"> <thead> <tr> <th style="padding: 5px;"></th> <th style="padding: 5px; text-align: center;">C</th> <th style="padding: 5px; text-align: center;">D</th> </tr> </thead> <tbody> <tr> <th style="padding: 5px; text-align: left;">C</th> <td style="padding: 5px; text-align: center;">5.25, 5.25</td> <td style="padding: 5px; text-align: center;">-0.75, 6.75</td> </tr> <tr> <th style="padding: 5px; text-align: left;">D</th> <td style="padding: 5px; text-align: center;">6.75, -0.75</td> <td style="padding: 5px; text-align: center;">0.75, 0.75</td> </tr> </tbody> </table>		C	D	C	5.25, 5.25	-0.75, 6.75	D	6.75, -0.75	0.75, 0.75
	C	D																	
C	6,6	-2,8																	
D	8,-2	0,0																	
	C	D																	
C	5.25, 5.25	-0.75, 6.75																	
D	6.75, -0.75	0.75, 0.75																	

Figure 1. Payoff matrices for each condition. Payoffs are denoted in points.

Subjects were informed of the specifics of their treatment (but not the existence of other treatments) in the experimental instructions, which are included in the Online Appendix.

The stage game is the prisoner’s dilemma in Figure 1 where the payoffs are denoted in points. Cooperation and defection take the “benefit/cost” (b/c) form, where cooperation means paying a cost c to give a benefit b to the other player, while defection gives 0 to each party; b/c

took the values 1.5 and 4.³ The expected payoff tables in Figure 1 incorporate the noise probability $E=1/8$. Subjects were presented with both the b/c representation and the resulting pre-error payoff matrix, in neutral language (the choices were labeled A and B as opposed to the “C” and “D” that is standard in the prisoner’s dilemma).

As noted earlier, in the explicit-intentions treatment, the highest equilibrium payoffs can be supported with strategies that condition only on intentions and ignore outcomes; moreover, the set of such equilibria is the same as in a game with the same expected payoff matrix and explicit actions. Under both of the payoff specifications we used, the explicit-intentions game has subgame-perfect equilibria in which both players cooperate each period, including for example the strategy profile where both players use the “Grim” strategy, which is “Play C iff either player has never played D”. However, Dal Bó (2005) shows that the existence of a cooperative equilibrium in a repeated game without noise is not sufficient for there to be much cooperation, and subsequent work by Blonski et al. (2011), Dal Bó and Frechette (2011), (2013) and Rand & Nowak (2013) suggests that a key determinant is whether Grim risk-dominates the strategy “Always Defect” (ALLD) in a 2x2 game. We might suspect that a similar pattern would apply to games with noise and observed intentions, so we note that “Grim-I” (the grim strategy that conditions only on intentions and ignores outcomes) risk-dominates ALLD even in the low-benefit treatment.⁴

Turning to the game with errors and implicit intentions, we note that in the low-benefit treatment $b/c=1.5$ “both play Grim” is not a Nash equilibrium but this is an equilibrium when $b/c=4$.⁵ Consistent with this observation, the experiments reported in Fudenberg et al. (2012) found substantially more cooperation at $b/c=4$ than when $b/c=1.5$.

A total of 338 subjects participated at the Harvard Decision Science Laboratory in Cambridge, MA. In each session, 12-32 subjects interacted anonymously via computer using the

³ Each session used a single payoff specification. Note that the benefit/cost specification implies that the short-run gain to playing D instead of C is independent of the other player’s action. The prisoner’s dilemma is more general than this; its defining characteristics are that D is a dominant strategy and that both playing C yields the highest payoff - in particular both playing C should be more efficient than alternating between (C,D) and (D,C).

⁴ In the noisy repeated game Grim-I earns a discounted average payoff of $7/8$ when facing itself. Grim-I vs ALLD yields $-11/8$ the first period, then $1/8$ afterwards for discounted average of $-4/64$; ALLD vs Grim-I earns $26/64$; and ALLD vs ALLD earns $1/8$. Thus facing a 50-50 mixture between the two strategies, Grim-I earns $26/64$, while ALLD gets $17/64$.

⁵ See the online appendix to Fudenberg et al. (2012) for equilibrium calculations of the implicit-intentions game. Note that there are many other cooperative equilibria in this game when $b/c=4$, including “perfect Tit for tat”, which says to play C if yesterday’s outcome was (C,C) or (D,D) and otherwise play D.

software z-Tree (Fischbacher 2007) in a sequence of infinitely repeated prisoner’s dilemmas (see Table 1 for summary statistics on the different conditions). We conducted a total of 16 sessions between September 2009 and December 2012.⁶ We only implemented one condition during a given session, so each subject participated in only one condition. We used the exchange rate of 30 units = \$1. Subjects were given a show-up fee of \$10 plus their winnings from the repeated prisoner’s dilemma.⁷ To allow for negative stage-game payoffs, subjects began the session with an “endowment” of 50 units (in addition to the show-up fee).⁸ On average subjects made \$18 per session, with a range from \$11 to \$32. Sessions lasted approximately 60 minutes.⁹

Table 1. Summary statistics per condition and b/c.

	b/c=1.5			b/c=4		
	Explicit intentions	No error	Implicit intentions	Explicit intentions	No error	Implicit intentions
Sessions per condition	2	2	3	2	3	4
Subjects per condition	44	44	72	40	48	90
Average number of supergames	9.5	9	11.25	9.8	8.25	11.9
Average number of periods per supergame	8.2	8.2	8.4	8.1	8.2	8.1

To implement random game lengths, we followed the procedure of Dreber et al. (2008) and Fudenberg et al. (2012): In each session every first supergame lasted t_1 periods, every second supergame lasted t_2 etc. For comparability between the implicit and explicit intentions data, we used the sequence of game lengths generated in Fudenberg et al. (2012), completing as many games as possible within the allotted session time.¹⁰

⁶ All sessions were conducted during the academic year, and all subjects were recruited through the CLER lab at Harvard Business School using the same recruitment procedure. Some of the data in our control treatments were originally reported in Fudenberg et al. (2012); these earlier controls used the more demanding “turnpike protocol” as a way to rule out contagion effects. However as the turnpike protocol restricts the number of supergames to one-half of the subjects in the room, and Dal Bó and Frechette (2011) argue that it doesn’t matter, our subsequent work has replaced it with the more common random-matching protocol.

⁷ Subjects also received earnings from a post-PD allocation decision that they were unaware of when playing the PD.

⁸ No subject finished with an endowment of fewer than 19 units, and only 2 out of 338 subjects had fewer than 50 units.

⁹ Subjects were given at most 30 seconds to make their decision, and informed that after 30 seconds a random choice would be made. The frequency of random decisions was very low, 0.006.

¹⁰ Note that from the viewpoint of the subjects, it was irrelevant when the game lengths were determined.

Past experiments on revealed intentions in games with errors have only studied one-shot or finitely repeated games. Bereby-Meyer and Roth (2006) explore cooperation in the one-shot and finitely repeated prisoner's dilemma where actions are implemented without noise and payoffs are either a deterministic or stochastic function of the actions played; since the end period is common knowledge, there is a substantial last-period effect in each supergame (as in Kunreuther et al. 2009 who also explore random payoffs). They find that the outcome due to the random shock in the previous period matters for the decision to cooperate this period, but less so than whether the other player cooperated or defected in the previous period.¹¹ Charness and Levine (2007), Cushman et al. (2009) and Schächtele et al. (2011) study one-shot games where intentions and outcomes can be each be either good or bad, can be in conflict due to a random device, and can either be rewarded or punished.¹² While a significant fraction of subjects in these games at least partially condition on intentions, there is also a tendency for them to condition on outcomes. Because these settings do not have a cooperative equilibrium, this work offers little guidance as to the fraction of players that will follow the theoretically optimal policy of conditioning only on intentions in the explicit-intentions treatment, on how the players who do respond to outcomes will do so, or on how play in the explicit-intentions and no-error infinitely repeated games will compare. Note that although conditioning only on intentions yields the highest equilibrium payoffs, a subject who believes that other subjects will condition on outcomes as well as intentions will find it optimal to do so as well, as may a subject who is uncertain whether others respond to outcomes as well as intentions.

Inspired both by past experimental findings and theoretical concerns, we organize our analysis around the following questions:

QUESTION 1: Does observing intentions allow more cooperative play compared to no intentions?

QUESTION 2: How similar are cooperation rates in the explicit-intentions treatments compared to the no-error treatments?

¹¹ Aoyagi and Frechette (2009) and Ambrus and Greiner (2012) study finitely repeated games with imperfect public monitoring; but subjects do not observe their opponents' intended actions...

¹² Additional work exploring the role of intentions in one-shot games is described in papers such as Blount (1995), Brandts and Solà (2001), Andreoni et al. (2002), Falk et al. (2003), McCabe et al. (2003) and Falk et al. (2008), which either compare how subjects respond to offers made by humans versus randomly generated offers (thus removing intentionality) or vary the strategy space of one player (thus changing the intentionality associated with a given outcome). In these papers intentions and outcomes are not in conflict.

QUESTION 3: How close do subjects come to basing their play solely on intentions?

QUESTION 4: To the extent that subjects condition on realized outcomes as well as intentions in the explicit-intentions treatment, how do they do this?

QUESTION 5: How do the strategies used in the explicit-intentions games compare to those used in games with implicit intentions or no errors?

3. Results




Figure 2. First period cooperation across supergames for each condition and b/c .

Before addressing our main questions of interest, we investigate the extent to which behavior changed over the course of a session as the result of learning. Figure 2 displays measures of aggregate behavior in each supergame, and suggests that some learning occurred, especially in the no-error and explicit-intentions conditions. Examining each condition separately, we find a significant increase in first period cooperation over supergame number in the explicit-intentions treatments ($b/c=1.5$: $p=0.010$; $b/c=4$, $p<0.001$) and the no-error controls ($b/c=1.5$: $p=0.003$; $b/c=4$, $p=0.034$), but not in the implicit-intentions controls ($b/c=1.5$: $p=0.788$; $b/c=4$, $p=0.166$).¹³ Consistent with this, a regression of all data together¹⁴ shows that learning (as measured by the

¹³ Logistic regression with robust standard errors clustered on subjects and group (i.e. subject pairing). Two-level clustering in all regressions follows the procedure described in Thompson (2011). See Appendix A Table A1 for full regression details.

¹⁴ Logistic regression with robust standard errors clustered on subjects and group, including dummies for control type (no errors, implicit-intentions) and b/c ratio, and interacting control type with supergame number. P-values are those associated with the condition dummy X supergame number coefficients. We note that this regression finds significantly more cooperation at $b/c=4$ than $b/c=1.5$ ($p<0.001$), consistent with previous work (Dal Bó 2005, Dreber et al. 2008).

effect of supergame number on first-period cooperation) is significantly slower in the implicit-intentions control compared to the no-error control ($p < 0.001$), but that learning is equally fast ($p = 0.615$) in the no-error control and the explicit-intentions treatment.¹⁵ To balance the need for data with the evidence of learning, we focus our analysis on the last four supergames of each session, as in Fudenberg et al. (2012).

To address Questions 1 and 2, we now examine the aggregate level of cooperation over the last four supergames, as well as the fraction of the time that subjects cooperated in the first period of a new supergame.




Figure 3. First period cooperation and overall cooperation for each condition and b/c .

QUESTION 1 Does observing intentions allow more cooperative play compared to no intentions?

QUESTION 2 How similar are cooperation rates in the explicit-intentions treatments compared to the no-error treatments?

Figure 3 suggests that there is roughly equal cooperation in the explicit-intentions treatments and the no-error controls, and less cooperation in the implicit-intentions controls (Q1-Q2). Statistical tests confirm that there is indeed more cooperation when intentions are explicit

¹⁵ These results differ from those of Bereby-Meyer and Roth (2006), who find that learning is slower in a finitely repeated game with probabilistic payoffs where intentions are explicit than in one with deterministic payoffs. Their analysis compares first period cooperation in the first supergame with that of the last one (rather than regressing across all supergames). Analyzing our data in that way does not change our results qualitatively: we still find that learning is significantly faster in explicit-intentions than implicit-intentions ($p = 0.018$) but that there is no significant difference between explicit-intentions and no error ($p = 0.808$). We return to this difference between our results and those of Bereby-Meyer and Roth in the Discussion and Appendix C.

than when they are implicit (significant difference at $p=0.02$ or less for all comparisons, except for first period cooperation in the cooperative $b/c=4$ treatment, where the differences with implicit-intentions are not significant, although of the same sign (vs. explicit-intentions, $p=0.14$; versus no error, $p=0.294$)).¹⁶

To answer Questions 3 through 5, we turn from aggregate behavior to considering the particular strategies used by subjects in our experiments. As one way to do this, we use the SFEM of Dal Bó and Frechette (2011) to assign probability weights to a predefined set of strategies. We complement this method with descriptive statistics that do not require the specification of a particular strategy set, but instead make assumptions about the general form of strategies employed.

Before addressing the remaining experimental questions in turn, we briefly describe the SFEM of strategy estimation and present its results. We then draw from these results to answer our questions. We will only summarize this method here (see Dal Bó and Frechette 2011 and Fudenberg et al. 2012 for more information). The idea is to restrict attention to a relatively large but finite set S of strategies, and suppose that each subject chooses a fixed element of S in the last 4 supergames, and moreover that regardless of whether there are exogenous errors, subjects make mistakes or “mental errors” when choosing their intended action. These mistakes let us assign a positive likelihood to any history for player and any strategy, and we can then assign an aggregate likelihood to any probability distribution p on S . We estimate p by MLE, and compute the standard errors by bootstrap; Appendix C presents the likelihood function we use.

A key aspect of this approach is choosing the set of strategies S to include in the estimation. Given the available data it is not possible to distinguish all possible strategies, as some histories arise only rarely and infinitely many can never occur at all in any finite sample. Guided by theoretical considerations and past empirical work we begin with a set of 38 of strategies, and then discard those that did not seem to be present in at least one payoff specification (including the no-error and implicit-intention controls). Roughly speaking, we start from the strategies that Dal Bó and Frechette (2011) and Fudenberg et al. (2012) found had a non-negligible share in at least one treatment, and add similar ones that condition on intentions

¹⁶ Unless otherwise noted, all subsequent p-values are generated using logistic regressions taking cooperation choice (0=D, 1=C) as the dependent variable and a treatment dummy as the independent variable, with robust standard errors clustered on subjects and group. To generate the p-values reported in this paragraph, we performed pairwise comparisons including the data from each relevant pair of bars in Figure 3. See Appendix A Table A2 for regression details.

alone or both intentions and outcomes. Appendix E lists all of the 38 original strategies and the procedure for discarding strategies.

Our final strategy set includes 17 strategies, which are described in Table 2. In addition to describing each strategy, Table 2 also indicates which strategies are lenient, in that they wait for multiple defections to punish, and which are forgiving, in that they are willing to return to cooperating following a breakdown in cooperation. We are particularly interested in these lenient strategies given their prevalence in the implicit-intentions controls, as reported in Fudenberg et al. (2012).

Table 2. Descriptions of the 17 strategies included in the main SFEM analysis.

Strategy	Abbreviation	Description
<i>Unconditional strategies</i>		
Always Cooperate	ALLC	Always play C (Lenient & forgiving)
Always Defect	ALLD	Always play D
<i>Strategies that condition on outcomes</i>		
Tit-for-Tat	TFT	Play C unless partner's action was D last period (Forgiving)
Tit-for-2-Tats	TF2T	Play C unless partner's action was D in both of the last 2 periods (Lenient & forgiving)
Tit-for-3-Tats	TF3T	Play C unless partner's action was D in all of the last 3 periods (Lenient & forgiving)
2-Tits-for-1-Tat	2TFT	Play C unless partner's action was D in either of the last 2 periods (2 periods of punishment if partner plays D) (Forgiving)
2-Tits-for-2-Tats	2TF2T	Play C unless there were 2 consecutive periods out of the last 3 periods in which either the partner's action was D (2 periods of punishment if partner plays D twice in a row) (Lenient & forgiving)
Grim	Grim	Play C until either player's action is D, then play D forever
Grim 2	Grim2	Play C until 2 consecutive periods occur in which either player's action was D, then play D forever (Lenient)
Grim 3	Grim3	Play C until 3 consecutive periods occur in which either player's action was D, then play D forever (Lenient)
Exploitive Tit-for-Tat	D-TFT	Play D in the first period, then play TFT
<i>Strategies that condition on intentions</i>		
Intention based Tit-for-Tat	TFT-I	Play C unless partner's intention was D last period (Forgiving)
Intention based Tit-for-3-Tats	TF3T-I	Play C unless partner's intention was D in all of the last 3 periods (Lenient & forgiving)
Intention based Grim	Grim-I	Play C until either player's intention is D, then play D forever
Intention based Grim 2	Grim2-I	Play C until 2 consecutive periods occur in which either player's intention was D, then play D forever (Lenient)
<i>Strategies that condition on both intentions and outcomes</i>		
Tolerant Tit-for-Tat	TFT-T	Play C unless partner's intention and action were both D last period (Forgiving)
Punitive 2-Tits-for-2-Tats	2TF2T-P	Play C unless there were 2 consecutive periods out of the last 3 periods in which either the partner's intention or action was D (2 periods of punishment if partner intends or actually plays D twice in a row) (Lenient & forgiving)

The probability assigned to these 17 strategies by the SFEM procedure in each of our 6 conditions is shown in Table 3.¹⁷ Only strategies that condition on outcomes are included in the SFEM for the no-error control (because intention and outcome are the same) and in the implicit-intention control (because intention information was unavailable to the subjects).

Table 3. SFEM results by condition. Bootstrapped standard errors shown in parentheses. Difference from 0 indicated by ** $p < 0.01$, * $p < 0.05$, † $p < 0.1$

	b/c=1.5			b/c=4		
	No Error	Explicit-Intentions	Implicit-Intentions	No Error	Explicit-Intentions	Implicit-Intentions
ALLC	0 (0)	0.02 (0.03)	0 (0)	0.24* (0.1)	0.07 (0.06)	0.06† (0.03)
TFT	0.27** (0.09)	0 (0)	0.19** (0.05)	0.14† (0.07)	0 (0)	0.07* (0.03)
TF2T	0 (0.02)	0.04 (0.04)	0.05 (0.03)	0 (0.04)	0.02 (0.02)	0.2** (0.07)
TF3T	0 (0)	0 (0.02)	0.01 (0.01)	0 (0.04)	0 (0.02)	0.09† (0.05)
2TFT	0 (0)	0 (0)	0.06 (0.04)	0.15* (0.07)	0 (0)	0.03 (0.02)
2TF2T	0.06 (0.04)	0 (0)	0 (0.02)	0 (0)	0 (0.01)	0.12* (0.06)
Grim	0.43** (0.08)	0 (0)	0.14** (0.04)	0.15† (0.08)	0 (0)	0.04 (0.02)
Grim2	0.01 (0.02)	0 (0)	0.06† (0.03)	0.16† (0.08)	0.04 (0.04)	0.05† (0.03)
Grim3	0 (0)	0.01 (0.02)	0.06† (0.03)	0 (0.05)	0.07 (0.06)	0.11** (0.04)
ALLD	0.18** (0.06)	0.18** (0.06)	0.29** (0.07)	0.07† (0.04)	0.10* (0.05)	0.23** (0.05)
D-TFT	0.05 (0.04)	0 (0)	0.14** (0.05)	0.09† (0.05)	0 (0)	0 (0)
TFT-I		0.20* (0.09)			0.04 (0.09)	
TF3T-I		0.07 (0.05)			0.16† (0.08)	
Grim-I		0.26** (0.08)			0.14† (0.08)	
Grim2-I		0.02 (0.04)			0.18† (0.09)	
TFT-T		0.04 (0.04)			0.19† (0.10)	
2TF2T-P		0.15* (0.06)			0 (0.02)	
<i>Gamma</i>	0.36** (0.02)	0.31** (0.02)	0.46** (0.02)	0.35** (0.03)	0.39** (0.04)	0.43** (0.02)

¹⁷ We note that similar strategies were estimated to be present in (i) our no-error controls and Dal Bó & Frechette's (2012) no-error games (they use a modified strategy method), and in (ii) our results and the games with no error and with public errors of Aoyagi et al. 2013. This provides evidence of the validity of the SFEM procedure.

We now turn to our remaining experimental questions.

QUESTION 3: How close do subjects come to basing their play solely on intentions?

Examining Table 3, we see that a majority of subjects in the explicit-intentions treatments disregard outcomes (i.e. play unconditional strategies or strategies that condition exclusively on intentions): 77% of probability weight at $b/c=1.5$, 69% of probability weight at $b/c=4$. Strategies conditioning exclusively on outcomes account for only 4% at $b/c=1.5$ and 13% at $b/c=4$, and strategies that condition on both intentions and outcomes account for 19% in each payoff specification.

As an additional way to examine this question, we consider all histories in which the opponent's intent and actual move in the last period differed. In 82% of such cases ($b/c=1.5$: 84%, $b/c=4$: 80%), the subject's decision matched the opponent's intent rather than the opponent's actual move. The results are qualitatively unchanged if we exclude subjects who cooperated in fewer than 25% of all decisions: the subject's decision then matched the opponent's intent in the previous period in 85% of cases ($b/c=1.5$: 85%, $b/c=4$: 85%).

Thus both the SFEM and the descriptive statistics suggest that a large majority of subjects base their play solely on intentions. This contrasts with the findings in some experiments on one-shot games, where maximizing money payoffs requires ignoring intentions entirely. Even though many subjects do condition at least partially on intentions in these games, there is much more of a tendency for them to condition on outcome as well compared to our results.¹⁸ The fact that intentions play a larger role in repeated games is consistent with the view that reciprocity in one-shot games stems from the application of a heuristic that was developed for repeated interactions.

We now focus our attention on the subset of players who *do* actually condition on outcomes to answer Question 4.

QUESTION 4: To the extent that subjects condition on realized outcomes as well as intentions in the explicit-intentions treatment, how do they do this?

¹⁸ For example, pooling data from the two treatments of Charness and Levine (2007), 59% of subjects (23 out of 39) rewarded when both intent and outcome were good, while 28.3% (15 out of 53) rewarded when intent was good but outcome was bad. Thus if all the subjects who rewarded after (good intent, bad outcome) would have done so after (good intent, good outcome), 28.3% of subjects used a purely intention-based strategy while 30.7% also conditioned on outcomes. The results of Cushman et al. (2009) are even more extreme, with no subjects conditioning purely on intentions, 36.7% conditioning purely on outcomes, and 46.7% conditioning on both intentions and outcomes.

Considering the results in Table 3, we see that subjects in the two payoff specifications condition on outcomes differently. At $b/c=1.5$, a non-negligible probability weight (15%) is assigned to the strategy 2TF2T-P. In general, 2TF2T waits for two periods of defection by the partner and then punishes for two periods. The 2TF2T-P variant of 2TF2T is “punitive,” in that it uses outcomes to punish an accidental D, but does not use a realization of C to forgive an intended D. That is, these players condition on outcome when the intention was C, but not when the intention was D.¹⁹ At $b/c=4$, conversely, a non-negligible probability weight (19%) is assigned to the strategy TFT-T (4% of subjects at $b/c=1.5$ also play TFT-T). This strategy is a variant of TFT that is “tolerant,” in that it uses outcomes to forgive would-be defectors who cooperated by accident, but not to punish unintended defections. That is, in contrast to the punitive version of 2TF2T, this tolerant strategy conditions on outcomes when the intention was D, but not when the intention was C.

To complement the SFEM analysis, we use simple descriptive measures that implicitly suppose that subjects ignore observations from two or more periods ago and only condition on observations from the previous period (Figure 4). Consistent with the MLE results, Figure 4 suggests that at $b/c=1.5$ subjects are punitive and condition on outcome when the opponent’s intention was C, but not when the opponent’s intention was D. This visual impression is confirmed by a positive relationship between a player’s probability of cooperating and the opponent’s actual move last period when the opponent intended to play C at $b/c=1.5$ ($p<0.001$).²⁰ Again consistent with the SFEM, we see a different pattern at $b/c=4$; here, subjects are tolerant and condition on outcome when the opponent’s intention was D but not C. The relationship between cooperation and opponent’s actual move last period when the opponent intended to play D at $b/c=4$ is not statistically significant ($p=0.236$), although the magnitude of the difference is not insubstantial (36.4%C when outcome is C, 27.7%C when outcome is D). As discussed below, we provide further evidence for tolerant outcome-based behavior by examining play in a repeated Ultimatum Game with errors and explicit intentions (see Discussion and Appendix B).

¹⁹ Charness and Levine (2007) find something similar in their one-shot experiments where some subjects only reward if both the intended and actual outcome are good.

²⁰ See Appendix A Table A3 for regression results.




Figure 4. Probability of cooperating in the last 4 supergames of the explicit-intentions conditions, as a function of opponent’s intention and actual move in the previous period.²¹

A natural question that arises from these results is whether subjects learn to ignore outcomes over time. We do not find evidence of such learning: Interacting opponent’s actual move last period with supergame number finds non-significant coefficients trending in the *positive* direction, both for $b/c=1.5$ when opponent intended C ($p=0.094$; controlling for period, player’s intended action last period, player’s overall frequency of cooperation, and player’s frequency of cooperation across first periods, $p=0.193$) and for $b/c=4$ when opponent intended D ($p=0.161$; with controls, $p=0.218$).²² The conclusion that subjects do not learn to ignore outcome is reinforced by comparing an SFEM analysis on the first four supergames to that on the last four: Fewer subjects condition on both intentions and outcomes in early supergames ($b/c=1.5$: 9%, $b/c=4$: 4%) than in later supergames ($b/c=1.5$: 19%; $b/c=4$: 19%). Thus there is no evidence that reliance on outcomes decreases with experience; if anything, it seems that reliance on outcomes may increase over time. See Appendix F for full regression tables and first four supergame SFEM results.

QUESTION 5: How do the strategies used in the explicit intentions games compare to those used in games with implicit intentions or no errors?

²¹ Error bars indicate standard errors of the mean clustered on subject and pairing. Average behavior over all subjects is shown.

²² We find qualitatively equivalent results when, instead of assuming that learning is linear in supergame number, we compare play in the first four supergames with play in the last four. (Interacting a “last 4 supergames” dummy with opponent’s actual move finds a non-significant positive coefficient in all cases; $b/c=1/5$, opponent intended C: $p=0.13$ without controls, $p=0.33$ with controls; $b/c=4$, opponent intended D: $p=0.174$ without controls, $p=0.160$ with controls).

In addition to asking whether and how subjects used realized outcomes in addition to intentions to guide their play, we are interested in how various strategic features of play with explicit intentions compare to those in the no-error and implicit-intentions controls. Specifically, we compare the memory length of the strategies used in the three conditions, and the extent to which play is “lenient” in the sense of not punishing the first deviation by an opponent. Intuitively, the high share of lenient strategies observed in the implicit-intentions controls may correspond to subjects giving their partner the benefit of the doubt that a defection could have occurred by accident, and to combine this sort of leniency with punishment for persistent defections requires strategies to look back more than one period.²³ Thus we might expect less leniency and more simple strategies in the explicit-intentions conditions compared to the implicit-intentions conditions.

Table 4 shows the relevant aggregated SFEM frequencies, as well as various descriptive statistics. First we consider the SFEM aggregations. In terms of strategy complexity, we see that the explicit-intentions treatments look very similar to the no-error controls: a large majority of cooperative strategies are simple in that they are either unconditional or condition/trigger based on the previous period only ($b/c=1.5$: 74% explicit-intentions, 75% no error; $b/c=4$, 81% explicit-intentions, 77% no error). This stands in stark contrast to the implicit-intentions controls, where the frequency of simple strategies is cut nearly in half (43% at both $b/c=1.5$ and $b/c=4$).

Next we consider leniency. We know from Fudenberg et al. (2012) that when intentions are implicit, leniency is common (and very successful) at b/c ratios where cooperative equilibria exist (such as $b/c=4$), but relatively rare at the low b/c ratio of 1.5 where there are no cooperative equilibria. If leniency reflects an attempt to infer the intentions of one’s partner, we would expect less leniency at $b/c=4$ when intentions are explicit than when they are implicit. Consistent with that expectation, at $b/c=4$, the fraction of cooperative strategies that are lenient in the explicit-intentions condition is much more similar to the no-error control than the implicit-intentions control. When $b/c=1.5$, it is not clear what to expect, because revealing intentions creates cooperative equilibria where none existed with implicit intentions: while there is less need to infer intentions when intentions are explicit, it is also much less costly to be lenient (since most

²³ Fudenberg et al. (2012) also considers the strategic element of ‘forgiveness,’ or willingness to return to cooperate after punishing a defection. Unlike leniency, there is not a clear a priori prediction about the effect of observing intentions on forgiveness. Thus we do not analyze forgiveness here, but include it in the Appendix G for completeness, where we show that there is not a clear relationship between it and whether intentions are observable.

others are cooperative). We find that when $b/c=1.5$, leniency in the explicit-intentions treatment is similar to the implicit-intentions control (and actually slightly higher), and both are higher than the no-error control.²⁴

We now complement these results with descriptive statistics. For maximum comparability, these measures use intentions in the explicit-intentions treatments, and outcomes in the implicit-intentions and no-error controls. In each case, we measure leniency by examining all histories in which C (either intentional or realized, depending on the measure) occurred in all but the previous period, while in the previous period one subject played D.²⁵ We then ask how frequently the subject who had hitherto cooperated showed leniency by continuing to cooperate. The results are similar to the SFEM. At $b/c=4$, leniency in the explicit-intentions condition is lower than the implicit-intentions control, whereas at $b/c=1.5$, the amount of leniency is similar in explicit-intentions and implicit-intentions.

In sum, we find evidence that in the presence of errors, making intentions explicit increases the frequency of cooperative strategies, reduces the complexity of those cooperative strategies, and also reduces the extent of leniency (at least at $b/c=4$, the specification in which leniency is common when intentions are implicit).

²⁴ As the fraction of cooperative strategies varies across condition (in particular, at $b/c=1.5$ the implicit-intention condition is much lower than the other conditions), we report the fraction of cooperative strategies that are lenient, rather than the fraction of all strategies that are lenient. For completeness, we report the un-normalized values here: $b/c=1.5$: no-error=7%, explicit-intentions=31%, implicit-intentions=17%; $b/c=4$: no-error=40%, explicit-intentions=53%, implicit-intentions=63%.

²⁵ We also include second round decisions in which the first round's outcome was CD.

Table 4. Aggregated SFEM frequencies and descriptive results by condition.

	b/c=1.5			b/c=4		
	No error	Explicit intentions	Implicit intentions	No error	Explicit intentions	Implicit intentions
Cooperative strategies in SFEM	77%	82%	56%	84%	90%	77%
<i>% of Cooperative strategies in SFEM that are:</i>						
Memory at most 1	75%	74%	43%	77%	81%	43%
Lenient	9%	38%	31%	47%	59%	81%
<i>Descriptive statistics</i>						
%C first period	73%	75%	54%	83%	86%	76%
%C all periods	49%	60%	32%	78%	75%	59%
Leniency	15%	28%	29%	42%	55%	66%

Discussion

We begin by asking how well subjects did in terms of maximizing their payoffs, both overall and by type of strategy used. This provides some insight into which sorts of strategies were (ex post) mistakes, and gives us a rough sense of how close the distribution of play is to an equilibrium - e.g. what percentage of players are receiving close to the best possible payoff given the distribution of play.

Table 5 shows the expected payoff of each strategy given the distribution estimated by the SFEM. In the explicit-intentions treatment at $b/c=1.5$, the two most prevalent strategies are TFT-I and Grim2-I; these purely intention-based strategies also yield the highest payoff of 5.5. The payoff of the punitive 2TF2T-P, which was also somewhat common, was slightly lower, but this difference is not statistically significant.

Table 5. SFEM frequency and expected payoffs for each strategy in each condition. Highest payoff strategies, and strategies with payoffs that are not statistically different from the highest payoff (based on bootstrapped standard errors using a significance level of $p < 0.1$), are highlighted in gray.

	b/c=1.5						b/c=4					
	No Error		Explicit Intentions		Implicit intentions		No error		Explicit Intentions		Implicit intentions	
	MLE	Payoff	MLE	Payoff	MLE	Payoff	MLE	Payoff	MLE	Payoff	MLE	Payoff
ALLC		3.5	0.02	3.6		-1.3	0.24	43.0	0.07	36.5	0.06	28.1
TFT	0.27	5.9		4.6	0.19	2.4	0.14	42.3		33.6	0.07	29.0
TF2T		5.7	0.04	4.8	0.05	1.5		43.9	0.02	36.7	0.20	29.6
TF3T		5.5		4.6	0.01	0.9		43.8		37.0	0.09	29.5
2TFT		5.8		4.7	0.06	2.9	0.15	40.8		30.3	0.03	27.0
2TF2T	0.06	5.7		4.9		1.9		43.9		36.3	0.12	29.6
Grim	0.43	5.8		4.4	0.14	3.0	0.15	40.8		25.9	0.04	24.0
Grim2	0.01	5.7		4.8	0.06	2.4	0.16	43.9	0.04	33.5	0.05	27.9
Grim-3		5.5	0.01	4.9	0.06	1.8		43.8	0.07	36.4	0.11	29.2
ALLD	0.18	2.5	0.18	4.1	0.29	3.7	0.07	21.1	0.10	19.1	0.23	21.0
D-TFT	0.05	2.7		4.1	0.15	2.9	0.09	25.4		29.8		28.7
TFT-I			0.20	5.5					0.04	37.7		
TF3T-I			0.07	5.1					0.16	37.4		
Grim-I			0.26	5.5					0.14	37.6		
Grim2-I			0.02	5.3					0.18	37.5		
TFT-T			0.04	5.3					0.19	37.5		
2TF2T-P			0.15	5.2						36.5		

The most common poorly performing strategy here is ALLD; this is a common feature of repeated games experiments. The payoff loss to ALLD is even higher in the no-error control, due to the smaller share of lenient strategies and the increased share of the unforgiving strategy set Grim and its variants. Conversely, ALLD yields the highest payoff in the implicit-intentions control, where it is also the most commonly used strategy. Mistaken extrapolation from that case could help explain the play of ALLD in the explicit-intentions and no-error treatments.

At $b/c=4$, in the explicit-intentions treatment most subjects play some sort of conditional cooperation strategy based on intentions only or intentions and outcomes; all of these strategies do fairly well, earning payoffs not statistically different from the highest performing strategy. In particular the expected payoff of the tolerant strategy TFT-T is statistically indistinguishable from that of TFT-I. Once again, the most common “mistake” is to play ALLD, which yields a

payoff of about 19 versus the high-30's payoffs obtained with conditional cooperation. Thus in both payoff treatments, a high fraction of the subjects do quite well, and subjects who condition on outcomes as well as intentions do so at essentially no cost.

As shown above in Figure 2, learning is significantly slower in the implicit-intentions control compared to the no-error control, but that learning is equally fast in the no-error control and the explicit-intentions treatment. This latter fact contrasts with the finding of Bereby-Meyer and Roth (2006), who found that adding a stochastic shock to payoffs (as opposed to actions) resulted in slower learning. One possible explanation is that our procedure speeds learning because it focuses attention on the opponent's intentions, which are what subjects need to be learning about in order to reach the best equilibrium.²⁶ This focus on intentions might be due either to the fact that implementation errors alter both players' payoffs in our case while the lotteries in Bereby-Meyer and Roth were independent, or because the framing of our game suggests to subjects that intentions are what matters.

To directly test this latter possibility, we ran a follow-up experiment on Amazon Mechanical Turk (Horton et al. 2011), recruiting 96 subjects and randomizing them into one of two ways to explain the structure of the random errors, the "Error" and "Lottery" conditions. In both conditions, subjects were told the payoff structure of our $b/c=4$ explicit-intentions PD. Then they were told to imagine the other player had chosen C and the low-probability D outcome had occurred, and asked to what extent they thought the other person intended to pay 0 cents and give 0 cents (i.e. had intended the D outcome) using a 7-point Likert scale (1="Do not intend it at all" to 7="Completely intended it").

In the Error condition, the probabilistic mechanism was explained with the same language as in our explicit-intentions condition. "There is a $7/8$ probability that the move you choose actually occurs. But with probability $1/8$, your move is changed to the opposite of what you picked." In the Lottery condition, the probabilistic mechanism was instead explained by saying that there were two options: A = $7/8$ chance of [-2 for you, +8 for other] and $1/8$ chance of

²⁶ Another possibility is that the difference in expected payoffs between cooperative and non-cooperative strategies may have been smaller in Bereby-Meyer and Roth's experiments than in our explicit-intentions conditions, thus providing a weaker signal for learning. To evaluate this possibility, one could use SFEM to estimate the distribution of strategies in Bereby-Meyer and Roth's data and then calculate the expected payoffs of each strategy. However, given that they used fixed length games, it is not clear to us which strategies should be included in the SFEM (e.g. strategies which open with cooperation but then switch to defection after some number of periods) so we do not explore this possibility here.

[0 for you, 0 for other] or B = 7/8 chance of [0 for you, 0 for other] and 1/8 chance of [-2 for you, +8 for other].²⁷

As predicted, we found that subjects in the Error condition thought that the [0,0] outcome was significantly less intentional than subjects in the Lottery condition did (mean intentionality ratings: Error: 2.27, Lottery: 3.33; Rank-sum, $p=0.036$; Tobit regression with robust standard errors: $p=0.020$; including controls for age, gender and education: $p=0.009$). This result supports our hypothesis that framing noise as execution errors emphasizes the ‘accidental’ nature of bad outcomes relative to framing noise as a lottery, and so increases the subjects’ attention to the intentions of their partner. Put differently, the execution-error framing decreases subjects’ sense of ‘causal control’ (Cushman et al. 2009) relative to the lottery framing, in that the error frame makes it seem as though some other agent (the computer) is causing the bad outcome, rather than the player.²⁸ We conjecture that this increased the subjects’ ability to focus on the intended good outcome and that this is why learning proceeded more quickly.

We also see that subjects use simpler, lower-memory strategies with explicit intentions than when intentions are implicit. This suggests that the more complex strategies found in the implicit-intentions condition use longer memory in part as a way to attempt to learn and track the intentions of other subjects. This is particularly true for $b/c=4$, where there is a high return to cooperation, and long memory lenient cooperative strategies were most prevalent with implicit intentions.

To further investigate strategic complexity, we examine how response times vary with play and condition. As in Rand et al (2012), we see that faster decisions are more cooperative ($p<0.001$).²⁹ Considering variation by condition, we would predict based on the complexity of the decision setting that decision times should be fastest in the no-error control, slowest in the implicit-intentions control, and intermediate with explicit intentions. When we examine the data, we see that the decision times conform to this prediction when $b/c=4$.³⁰ However, when $b/c=1.5$,

²⁷ See the Online Appendix for the full instructions.

²⁸ See also Bolton et al. (2005) who explore procedural fairness versus outcome fairness.

²⁹ Logistic regression with robust standard errors clustered on subject and pairing, considering the last four supergames. Log-10 transformed response time is taken as the independent variable, and controls for condition (dummies for explicit-intention, no-error or implicit-intention), b/c , supergame number and period number are included. Decision times are log-10 transformed as in Rand et al 2012 to account for the heavily skewed nature of the response time distributions. Equivalent results are found when using untransformed response times, when including all supergames, or both. See Online Appendix A Table OA1.

³⁰ Log-10 transformed response times: no-error, 0.144; explicit-intentions, 0.185; implicit-intentions, 0.263. An equivalent ordering is obtained when controlling for supergame, period, and whether the decision was C or D.

decision time is longest in the explicit-intentions treatment, with the no-error treatment coming second.³¹

We also see interesting differences across conditions in learning: In the no-error and explicit-intentions conditions, decision times decrease with experience as measured by supergame number on ($p < 0.001$ for both), as we would expect; but with implicit intentions, decision times actually *increase* ($p < 0.001$).³² Further investigation of this increasing response time in the implicit-intentions condition finds a significant negative interaction between supergame number and subjects' overall frequency of cooperation when predicting response time ($p < 0.001$; controlling for the period 1 intention of the partner in the previous supergame and current partner's action in the previous period, $p = 0.001$)³³: the less often a subject cooperated overall, the more her response time tended to increase with experience (controlling for whether the current decision was C or D). As a result, regressing decision time against supergame number finds no significant relationship when only examining subjects who were largely cooperative (i.e. cooperated in more than 2/3 of all decisions, $p = 0.890$).³⁴ This increase in reaction times by non-cooperative subjects may reflect the effects of learning: these subjects are generally inclined to defect, but over time gather evidence that cooperation might be in their interest. This makes cooperation somewhat more attractive, moving the expected utility of cooperation and defection closer together, and as a result faces these subjects with a more difficult and time consuming choice (in contrast to the initially cooperative subjects, whose inclination is reinforced by experience).

³¹ Log-10 transformed response times: no-error, 0.294; explicit-intentions, 0.224; implicit-intentions, 0.235. An equivalent ordering is obtained when controlling for supergame period, and whether the decision was C or D, except that the very similar response times of explicit-intentions and implicit-intentions flip.

³² Linear regressions with robust standard errors clustered on subject and pairing, taking $\log_{10}(\text{response time})$ as the DV, supergame as the IV, and including controls for b/c, period and whether the decision was C or D. See Online Appendix A Table OA2..

³³ Linear regression with robust standard errors clustered on subject and pairing, considering the implicit-intentions condition, taking $\log_{10}(\text{response time})$ as the DV, supergame number and frequency of cooperation over all decisions as the IVs, and including an interaction between these two terms as well as controls for b/c, period and whether the decision was C or D.

³⁴ The cutoff of 2/3 was determined by testing at which frequency of overall cooperation the net coefficient on supergame number became non-significant. P-value reflects linear regression with robust standard errors clustered on subject and pairing, considering the implicit-intentions condition, taking $\log_{10}(\text{response time})$ as the DV, supergame number and frequency of cooperation over all decisions as the IVs, and controls for b/c, period and whether the decision was C or D. The finding of no effect among these subjects persists when also controlling for the period 1 intention of the partner in the previous supergame and current partner's action in the previous period ($p = 0.192$).

Although we do not have a theoretical explanation for all of the response time data, we feel that the connection between response time and choice of strategy, and how this varies with the strategic environment, is an interesting topic that merits future study

Our data shows that while subjects focus primarily on intentions, outcomes can matter too. To further understand when and how subjects rely on intentions versus outcomes, we also ran a series of repeated ultimatum games, where the proposer had only two possible offers (an equal split and a selfish one), the chosen offer was proposed with probability $7/8$, and using the strategy method player 2 specified responses for each of the four possible pairs of intended and actual offer. (The details of this experiment are in Appendix B.) This is a repeated game and so has equilibria in which players condition their actions on observations from past periods; in an informal sense the game resembles the $b/c=4$ treatment of the RPD, because the gains from cooperative behavior (e.g. fair offers) are high relative to other factors.³⁵ There is indeed some evidence of conditioning on past period: When faced with either an intentional or accidental unfair offer, P2s were more likely to accept if P1 intended A in the previous period. Regardless of P1's previous play, however, the use of outcomes looked qualitatively similar to that in the $b/c=4$ "cooperative" treatment of the RPD (among those who did condition on outcomes): play was largely tolerant, in that only intentional low offers were consistently rejected.³⁶ Thus our repeated ultimatum game experiment provides further evidence of the use of tolerant strategies that condition on both intentions and outcomes, as was suggested in the $b/c=4$ explicit-intentions treatment of our main experiment.

Conclusion

We conclude that making intentions explicit allows subjects to achieve the same level of cooperation under errors as if errors were not present, because subjects largely ignore outcomes and condition only on intentions. This finding is consistent with the predictions of theory in the sense that the highest payoff equilibrium ignores outcomes. Moreover, intention-based strategies

³⁵ For example there is an equilibrium in which Player 2s condition only on intentions and use the following strategy: "Accept when Player 1 intended the equal split; reject when he intended the selfish one, as long as I have always done so in the past, and otherwise always accept."

³⁶ Both intentional and accidental fair offers were accepted 100% of the time. Accidental low offers were accepted 88% of the time if P1 intended to make a fair offer in the previous period and 67% of the time if not. Intentional low offers were accepted 17% of the time if P1 intended to make a fair offer in the previous period and 7% of the time if not.

are both common and earn high payoffs given the observed distribution of play. Thus institutions that increase the observability of intentions may help to mitigate the negative consequences of errors: when your aim is true, accidents will by and large be forgiven and forgotten.

References

- Abreu, Dilip, David Pearce, and Ennio Stacchetti. 1990. "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring." *Econometrica* 58 (5): 1041–63.
- Ambrus, Attila, and Ben Greiner. 2012. "Imperfect Public Monitoring with Costly Punishment - An Experimental Study." *American Economic Review* 102 (7): 3317-32
- Andreoni, James, Paul M. Brown, and Lise Vesterlund. 2002. "What Makes an Allocation Fair? Some Experimental Evidence." *Games and Economic Behavior* 40, 1–24.
- Aoyagi, Masaki, V. Bhaskar, and Guillaume Frechette (2013), "The Impact of Monitoring in Infinitely Repeated Games: Perfect, Public, and Private." Manuscript in preparation.
- Aoyagi, Masaki, and Guillaume Frechette. 2009. "Collusion as Public Monitoring Becomes Noisy: Experimental Evidence." *Journal of Economic Theory* 144 (3): 1135–65.
- Bereby-Meyer, Yoella, and Alvin E. Roth. 2006. "The Speed of Learning in Noisy Games: Partial Reinforcement and the Sustainability of Cooperation." *American Economic Review* 96(4) 1029-1042.
- Bigoni, Maria, Jan Potters, and Giancarlo Spagnolo. 2012. "Flexibility and Collusion with Imperfect Monitoring." Working Paper.
- Blonski, Matthias, Peter Ockenfels, and Giancarlo Spagnolo. 2011. "Equilibrium Selection in the Repeated Prisoner's Dilemma: Axiomatic Approach and Experimental Evidence." *American Economic Journal: Microeconomics* 3 (3): 164–92.
- Blount, Sally. 1995. "When Social Outcomes aren't Fair: The Effect of Causal Attributions on Preferences." *Organizational Behavior and Human Decision Processes* 63: 131–144.
- Bolton, Gary E., Jordi Brandts, and Axel Ockenfels. 2005. "Fair Procedures: Evidence from Games Involving Lotteries." *Economic Journal* 115: 1054-1076.
- Brandts, Jordi, and Carles Solà. 2001. "Reference Points and Negative Reciprocity in Simple Sequential Games." *Games and Economic Behavior* 36: 138–157.
- Charness, Gary, and David I. Levine. 2007. "Intention and Stochastic Outcomes: An Experimental Study." *Economic Journal* 117: 1051-1072.
- Cushman, Fiery, Anna Dreber, Ying Wang, and Jay Costa. 2009. "Accidental Outcomes Guide Punishment in a 'Trembling Hand' Game." *PLoS ONE* 8, 1–7.
- Dal Bó, Pedro. 2005. "Cooperation under the Shadow of the Future: Experimental Evidence from Infinitely Repeated Games." *American Economic Review* 95 (5): 1591–1604.
- Dal Bó, Pedro, and Guillaume R. Frechette. 2011. "The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence." *American Economic Review* 101 (1): 411–29.

- Dal Bó, Pedro, and Guillaume R. Frechette. 2012. "Strategy Choice In The Infinitely Repeated Prisoners Dilemma." Working paper.
- Dreber, Anna, David G. Rand, Drew Fudenberg, and Martin A. Nowak. 2008. "Winners Don't Punish." *Nature* 452: 348–351.
- Falk, Armin, Ernst Fehr, and Urs Fischbacher. 2003. "On the Nature of Fair Behavior." *Economic Inquiry* 41, 20–26.
- Falk, Armin, Ernst Fehr, and Urs Fischbacher. 2008. "Testing Theories of Fairness — Intentions Matter." *Games and Economic Behavior* 62: 287–303.
- Fischbacher, Urs. 2007. "Z-Tree: Zurich Toolbox for Ready-Made Economic Experiments." *Experimental Economics* 10 (2): 171–78.
- Fudenberg, Drew, David G. Rand, and Anna Dreber 2012. "Slow to Anger and Fast to Forgive: Cooperation in an Uncertain World." *American Economic Review* 102 (2), 720-749.
- Fudenberg, Drew, David K. Levine, and Eric Maskin. 1994. "The Folk Theorem with Imperfect Public Information." *Econometrica* 62 (5): 997–1039.
- Green, Edward J., and Robert H. Porter. 1984. "Noncooperative Collusion under Imperfect Price Information." *Econometrica* 52 (1): 87–100.
- Kunreuther, Howard, Gabriel Silvasi, Eric Bradlow, and Dylan Small. 2009. "Bayesian Analysis of Deterministic and Stochastic Prisoner's Dilemma Games." *Judgment and Decision Making* 4 (5): 363–384.
- Horton John J., David G. Rand, and Richard J Zeckhauser. 2011. "The Online Laboratory: Conducting Experiments in a Real Labor Market." *Experimental Economics* 14 (3): 399–425.
- McCabe, Kevin A., Mary L. Rigdon, and Vernon L. Smith. 2003. "Positive Reciprocity and Intentions in Trust Games." *Journal of Economic Behavior and Organization* 52: 267–275.
- Radner, Roy, Roger Myerson, and Eric Maskin. 1986. "An Example of a Repeated Partnership Game with Discounting and with Uniformly Inefficient Equilibria." *Review of Economic Studies* 53: 59-69.
- Rand, David G., Joshua D. Greene, and Martin A. Nowak. 2012. "Spontaneous Giving and Calculated Greed." *Nature* 489: 427–430.
- Rand, David G., and Martin A. Nowak. 2013. "Human cooperation." *Trends in Cognitive Sciences* 17: 413-425.
- Schächtele, Simeon, Tobias Gerstenberg, and David Lagnado. 2011. "Beyond Outcomes: The Influence of Intentions and Deception." Working Paper, UCL.
- Thompson, Samuel B. 2011. "Simple formulas for standard errors that cluster by both firm and time." *Journal of Financial Economics* 99: 1-10.

Appendix A:

Table A1. First period cooperation by supergame. Logistic regression with robust standard errors clustered on subject and pairing.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Explicit-intentions		No-error		Implicit-intentions		All data
	b/c=1.5	b/c=4	b/c=1.5	b/c=4	b/c=1.5	b/c=4	
Supergame	0.118*** (0.0456)	0.189*** (0.0262)	0.118*** (0.0396)	0.129** (0.0607)	0.00583 (0.0217)	0.0336 (0.0243)	0.123*** (0.0332)
Explicit							0.0307 (0.305)
Implicit							0.0260 (0.263)
Explicit X Supergame							0.0230 (0.0458)
Implicit X Supergame							-0.101*** (0.0369)
b/c							0.318*** (0.0834)
Constant	0.214 (0.269)	0.658* (0.344)	0.0841 (0.298)	0.780*** (0.300)	0.152 (0.218)	0.793*** (0.226)	-0.439 (0.310)
Observations	416	392	396	396	810	1,072	3,482
Subject-clusters	44	40	44	48	72	90	338
Pairing-clusters	208	196	198	198	405	536	1,741

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table A2. Mean comparisons, first period and overall C. Logistic regression with robust standard errors clustered on subject and pairing.

	(1)	(2)	(3)	(4)
	First period C		Overall C	
	b/c=1.5	b/c=4	b/c=1.5	b/c=4
Explicit	0.946** (0.392)	0.708 (0.480)	1.167*** (0.288)	0.735** (0.317)
Constant	0.153 (0.208)	1.128*** (0.232)	-0.752*** (0.147)	0.348** (0.162)
Observations	464	520	4,032	4,444
Subject-clusters	116	130	116	130
Pairing-clusters	232	260	232	260

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table A3. Effect of intentions vs outcomes on overall cooperation in explicit intentions condition, by b/c. Logistic regression with robust standard errors clustered on subject and pairing.

	(1)	(2)	(3)	(4)	(5)	(6)
	b/c=1.5, Other Intended C			b/c=4, Other Intended D		
Other's Previous Outcome (0=D, 1=C)	1.104*** (0.289)	-0.0156 (0.474)	0.0786 (1.186)	0.399 (0.337)	-0.493 (0.520)	-0.456 (0.492)
Supergame		0.0667 (0.0823)	-0.0942 (0.142)		0.00376 (0.0406)	0.0484 (0.0312)
Outcome X Supergame		0.148* (0.0883)	0.227 (0.174)		0.132 (0.0942)	0.118 (0.0958)
Period			-0.131** (0.0591)			-0.0999* (0.0515)
Your Previous Intention (0=D, 1=C)			4.003*** (0.467)			1.253*** (0.269)
Your C Freq			8.702*** (1.147)			3.648*** (1.135)
Your Period 1 C Freq			-2.579*** (0.837)			-1.750** (0.806)
Constant	1.283*** (0.328)	0.957* (0.516)	-3.475*** (0.968)	-0.959*** (0.267)	-0.956*** (0.310)	-2.175*** (0.596)
Observations	795	1,708	1,708	293	761	761
Subject-clusters	44	44	44	31	38	38
Pairing-clusters	84	197	197	39	105	105

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Appendix B: Repeated Ultimatum Game

In our repeated Ultimatum Game experiments, 68 subjects were randomly assigned to be either Player 1 (P1) or Player 2 (P2) in a series of stochastically repeated Ultimatum Games (UG). Players had the same role in all supergames. Each supergame involved a random number of periods played with the same opponent. In each period of the UG, P1 chose one of two offers, A or B. Option A gave both players 50 MUs each.³⁷ Option B gave P1 80 MUs and P12 20 MUs. There was a 1/8 chance that the actual offer would be opposite of the offer chosen by P1. Using the strategy method, P2 indicated whether she would accept or reject each of the 4 possible [intention]x[outcome] pairs. Then both players were informed of P1's chosen offer, the actual offer, and whether P2 accepted or rejected. With 7/8 probability, the supergame would continue and the two players would play another period. With 1/8 probability, the supergame would end and players would be randomly rematched.³⁸




Figure B1. Mean P1 offer and P2 acceptance for each of the 4 possible [intention-outcome] pairs. Panel A shows behavior in the first period of each supergame. Panel B shows behavior over last 4 supergames by period.

We find some evidence of learning: the probability of P1 choosing the fair offer A increases over supergame (period 1: $p=0.011$; all periods, $p<0.001$), as does the probability of P2 rejecting accidental unfair offers (“Accept A-B” in Fig B1; period 1: $p=0.008$; all periods, $p=0.002$) (all other accept probabilities $p>0.20$). Thus we follow the main text PD analysis and focus on the last four supergames in our subsequent analyses.

We now consider the choices of the 34 P2s individually. We find that

³⁷ We used an exchange rate of 30 units = \$1.

³⁸ See the Online Appendix for the experimental instructions.

- 70.6% of subjects (N=24) condition on both intentions and outcomes and are tolerant (i.e. reject in more than 95% of cases where both P1's intention and the outcome were B, accept in more than 95% of cases otherwise).
- 14.7% of subjects (N=5) accept in 95% of cases or more regardless of P1's intention and outcome)
- 8.8% of subjects (N=3) condition on outcomes only, (i.e. reject in >95% of cases where the outcome was B regardless of P1's intention, and accept in >95% of the cases where the outcome was A regardless of P1's intention)
- 5.9% of subjects (N=2) were not easily classified, having acceptance probabilities >5% but less than 95% in one or more cases.

We also examine whether P2s condition their play on P1's intention in the previous period. In the last four supergames, no P2 ever rejected a fair offer, either intentional (A-A) or accidental (B-A). Considering unfair offers, however, we do find a significant effect of P1's intended offer in the previous period³⁹: Both for accidental low offers (A-B, $p=0.009$) and purposeful low offers (B-B, $p=0.004$), P2s were more likely to accept if P1 intended A in the previous period, as shown in Figure B2. Regardless of P1's intention in the previous period, however, we see a large amount of tolerance (i.e. good outcomes are used to forgive bad intentions much more than bad outcomes are used to punish good intentions).




Figure B2. Acceptance probabilities as a function of P1's intended offer in the previous period.

³⁹ P-values from logistic regression with robust standard errors clustered on subject and pairing, including controls for supergame number and period number.

In sum, our repeated UG experiment provides further evidence of the use of tolerant strategies that condition on both intentions and outcomes, as was suggested in the $b/c=4$ explicit-intentions treatment of our main experiment.

Appendix C: Re-analysis of learning in Bereby-Meyer and Roth (2006)

Bereby-Meyer and Roth (2006) (BMR) compare learning rates in finitely repeated games where payoffs are either deterministic or probabilistic.⁴⁰ They have also include a ‘sun-spots’ control in which payoffs are deterministic, but players are also presented with the outcome of two random lotteries (that do not effect payoffs) each turn. They conclude that learning to cooperate in period 1 occurs more slowly in the probabilistic condition than in either the deterministic or sun-spot conditions. As the main text explains, they used a different analysis strategy. Instead of regressing first period cooperation against supergame number, BMR compared just the first and the last supergame. This difference in methods did not change the analysis of our data.

Now, we re-analyze their data using our learning metric: comparing the coefficient for supergame number when predicting first period cooperation across conditions.⁴¹ Doing so, we also find that there is significantly slower learning when payoffs are probabilistic compared to deterministic in a repeated game (condition X supergame: coeff = 0.120, $p < 0.001$), but we do not find a significant difference in learning speed between their probabilistic condition and their sun-spot control (coeff = 0.032, $p = 0.224$). As can be seen in Figure C1, cooperation in the sun-spot control climbs rapidly in supergames 2 and 3, but then stabilizes.




Figure C1. Cooperation in period 1 of the finitely repeated games from Bereby-Meyer and Roth (2006).

⁴⁰ They also study one-shot games.

⁴¹ As in the main text, we use logistic regression with standard errors clustered on subject and group.

If we restrict our re-analysis to supergames 1 and 20 (as in their original analysis), we replicate their result: significantly more change from supergame 1 to 20 in both the deterministic ($p=0.006$) and the sun-spot ($p=0.005$) conditions compared to the probabilistic condition. These two methods of analysis give different results in the sun-spots condition because most of the learning there occurs in the first three supergames; regressing over all supergames shows relatively little change in first period cooperation per supergame, despite that fact that first period cooperation increases substantially from the first supergame to the last.

Overall, it seems that the noise in BMR had a substantially different effect on learning than the shocks to actions in our experiments. One possible explanation is that the noise in their probabilistic condition psychologically feels different from how we introduce noise.⁴² Although these two processes are mathematically equivalent, our procedure places more emphasis on the role of intentions: it makes it feel like the other person didn't *mean* to choose the outcome in cases where errors occur. In BMR's setup, although the players do not have direct control over the outcomes, it may not feel to the opponent that the randomness of lottery is changing the actor's intent. As discussed in the main text, we explore this possibility with an additional experiment.

⁴² Recall that in their probabilistic condition, players select between two lotteries, and then an outcome is drawn from the chosen lottery. In our explicit-intention conditions, players choose a fixed outcome, and the computer then switches their choice to the opposite with some probability.

Appendix D: Details of the Structural Frequency Estimation Method (SFEM)

We use the SFEM introduced by Dal Bo & Frechette (2011) and used in Fudenberg et al. (2012). We suppose that if subject i uses strategy s , her chosen action in round r of supergame k is C if $s_{ikr}(s) + \gamma \varepsilon_{ikr} \geq 0$, where $s_{ikr}(s) = 1$ if strategy s says to play C in round r of supergame k given the history to that point, and $s_{ikr}(s) = -1$ if s says to play D. Here ε_{ikr} is an error term that is independent across subjects, rounds, supergames, and histories, γ parameterizes the probability of mistakes, and the density of the error term is such that the overall likelihood that subject i uses strategy s is

$$(1) \quad p_i(s) = \prod_k \prod_r \left(\frac{1}{1 + \exp(-s_{ikr}(s) / \gamma)} \right)^{y_{ikr}} \left(\frac{1}{1 + \exp(s_{ikr}(s) / \gamma)} \right)^{1 - y_{ikr}},$$

where y_{ikr} is 1 if the subject chose C and 0 if the subject chose D.⁴³

For any given set of strategies S and proportions p , we then derive the likelihood for the entire sample as a mixture model, namely $\sum_I \ln \left(\sum_{s \in S} p(s) p_i(s) \right)$. Note that the specification assumes that all subjects are ex-ante identical with the same probability distribution over strategies and the same distribution over errors; one could relax this at the cost of adding more parameters. Because p describes a distribution over strategies, this likelihood function implies that in a very large sample we expect fraction $p(s)$ of subjects to use strategy s , though for finite samples there will be a non-zero variance in the population shares. We use maximum likelihood estimation (MLE) to estimate the prevalence of the various strategies, and bootstrapping to associate standard errors with each of our frequency estimates. We construct 100 bootstrap samples for each treatment by randomly sampling the appropriate number of subjects with replacement. We then determine the standard deviation of the MLE estimates for each strategy frequency across the 100 bootstrap samples. The validity of this procedure was demonstrated using simulated data in Fudenberg et al 2012.

⁴³ Thus the probability of an error in implementing one's strategy is $1/(1+\exp(1/\gamma))$. Note that this represents error in intention, rather than the experimentally imposed error in execution. This formulation assumes that all strategies have an equal rate of implementation error. In Fudenberg et al. (2012) we show that the MLE estimates of strategy shares are robust to allowing each strategy have a different value of γ .

**Appendix E: SFEM including all intent-outcome hybrids
(only the explicit-intentions conditions)**

	b/c=1.5	b/c=4
ALLC	0.02 (0.03)	0.07 (0.06)
TFT	0 (0)	0 (0)
TF2T	0 (0.01)	0 (0)
TF3T	0 (0)	0 (0)
2TFT	0 (0)	0 (0)
2TF2T	0 (0)	0 (0)
G	0 (0)	0 (0)
G2	0 (0)	0.04 (0.04)
G3	0 (0)	0.07 (0.06)
ALLD	0.18** (0.06)	0.1* (0.05)
DTFT	0 (0)	0 (0)
TFTI	0.19* (0.09)	0.03 (0.09)
TF2TI	0.02 (0.04)	0 (0.02)
TF3TI	0 (0.03)	0.16† (0.09)
2TFTI	0.09 (0.07)	0 (0.01)
2TF2TI	0 (0.04)	0 (0)
GI	0.24** (0.08)	0.15† (0.09)
G2I	0 (0.01)	0.17 (0.11)
G3I	0 (0.01)	0 (0.01)
D-TFTI	0.02 (0.03)	0 (0)

	b/c=1.5	b/c=4
TFTI CC	0.01 (0.02)	0 (0)
TF2TI CC	0 (0.03)	0.03 (0.04)
TF3TI CC	0.08 (0.06)	0 (0.02)
2TFTI CC	0 (0)	0 (0)
2TF2TI CC	0.08 (0.06)	0 (0.01)
GI CC	0 (0)	0 (0)
G2I CC	0 (0)	0 (0)
G3I CC	0 (0)	0 (0.04)
D-TFTI CC	0 (0)	0 (0)
TFTI DD	0 (0.04)	0.18† (0.1)
TF2TI DD	0.06 (0.05)	0 (0.01)
TF3TI DD	0 (0)	0 (0.05)
2TFTI DD	0 (0)	0 (0)
2TF2TI DD	0 (0.01)	0 (0.02)
GI DD	0 (0)	0 (0)
G2I DD	0 (0.04)	0.01 (0.09)
G3I DD	0 (0)	0 (0.01)
D-TFTI DD	0 (0)	0 (0)
Gamma	0.3** (0.02)	0.39** (0.04)

This table shows the results of a first MLE with all 38 possible strategies. Next we performed a second estimation including only strategies that had weight great than 0.05 in at least one condition. For our final MLE shown in the main text, we then included only the strategies that were present at $p < 0.10$ in at least one condition in the second MLE.

Appendix F – Change in conditioning on outcomes over period

Table F1. Cooperation as a function of opponent’s actual move in the previous period.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	b/c=1.5, Opponent intended C last period				b/c=4, Opponent intended D last period			
	Last 4 supergames		All supergames		Last 4 supergames		All supergames	
Opponent's Outcome	1.104*** (0.289)	1.666*** (0.556)	-0.0156 (0.474)	0.0786 (1.186)	0.399 (0.337)	0.366 (0.433)	-0.493 (0.520)	-0.456 (0.492)
Supergame Number		-0.0882 (0.165)	0.0667 (0.0823)	-0.0942 (0.142)		0.0584 (0.0942)	0.00376 (0.0406)	0.0484 (0.0312)
Period		-0.124 (0.100)		-0.131** (0.0591)		-0.211*** (0.0788)		-0.0999* (0.0515)
Player's intended C last period		3.915*** (0.466)		4.003*** (0.467)		1.237*** (0.464)		1.253*** (0.269)
Your overall C		6.013*** (1.010)		8.702*** (1.147)		3.326*** (0.781)		3.648*** (1.135)
Your first period C		-1.209 (0.837)		-2.579*** (0.837)		-1.229*** (0.428)		-1.750** (0.806)
Opponent's Outcome X Supergame Number			0.148* (0.0883)	0.227 (0.174)			0.132 (0.0942)	0.118 (0.0958)
Constant	1.283*** (0.328)	-3.136* (1.672)	0.957* (0.516)	-3.475*** (0.968)	-0.959*** (0.267)	-1.606* (0.905)	-0.956*** (0.310)	-2.175*** (0.596)
Observations	795	795	1,708	1,708	293	293	761	761

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table F2. SFEM analyzing first four supergames.

	b/c=1.5			b/c=4		
	No Error	Explicit Intentions	Implicit Intentions	No Error	Explicit Intentions	Implicit Intentions
ALLC	0 (0)	0 (0)	0.01 (0.01)	0.11† (0.06)	0.06 (0.05)	0.06† (0.03)
TFT	0.19* (0.09)	0.03 (0.03)	0.12* (0.05)	0.38** (0.09)	0 (0)	0.1** (0.04)
TF2T	0.08† (0.05)	0 (0.01)	0.05 (0.04)	0.13 (0.09)	0 (0)	0.16** (0.06)
TF3T	0 (0)	0.03 (0.03)	0.01 (0.01)	0 (0)	0 (0.02)	0.05 (0.04)
2TFT	0.04 (0.05)	0 (0)	0.11* (0.05)	0 (0)	0 (0)	0.03 (0.03)
2TF2T	0 (0.02)	0 (0)	0.11* (0.05)	0 (0.03)	0 (0.01)	0.15** (0.06)
Grim	0.22** (0.08)	0 (0)	0.07* (0.03)	0.14* (0.07)	0 (0)	0 (0.01)
Grim2	0.06 (0.04)	0 (0)	0.02 (0.02)	0.06 (0.06)	0 (0.02)	0.11** (0.04)
Grim3	0.02 (0.03)	0 (0)	0.04 (0.03)	0 (0.04)	0.11 (0.07)	0.07† (0.04)
ALLD	0.20** (0.06)	0.25** (0.07)	0.40** (0.06)	0.09* (0.04)	0.13* (0.05)	0.23** (0.05)
D-TFT	0.18** (0.06)	0.06 (0.04)	0.05 (0.03)	0.09* (0.05)	0.02 (0.03)	0.03 (0.02)
TFT-I		0.29** (0.1)			0.34** (0.11)	
TF3T-I		0.06 (0.04)			0 (0)	
Grim-I		0.18** (0.07)			0.19† (0.1)	
Grim2-I		0.03 (0.03)			0.11 (0.08)	
TFT-T		0.05 (0.04)			0.04 (0.05)	
2TF2T-P		0.04 (0.05)			0 (0.04)	
<i>Gamma</i>	0.43** (0.03)	0.43** (0.02)	0.57** (0.03)	0.37** (0.03)	0.51** (0.06)	0.51** (0.03)

Appendix G: Forgiveness across conditions

	b/c=1.5			b/c=4		
	No error	Explicit intentions	Implicit intentions	No error	Explicit intentions	Implicit intentions
% Cooperative strategies in SFEM that are forgiving	43%	65%	55%	63%	52%	73%
Descriptive statistics	5%	6%	13%	19%	17%	30%

Considering the SFEM aggregations and forgiveness, it is unclear what the effect of making intentions explicit is: at $b/c=1.5$, forgiveness is most frequent when intentions are explicit; at $b/c=4$, forgiveness is least frequent when intentions are explicit.

The results for descriptive statistics are somewhat more consistent. To measure forgiveness using descriptive statistics, we first identify all histories in which (i) at least one subject chose C in the first period, (ii) in at least one previous period, the initially cooperative subject chose C while the other subject chose D and (iii) in the immediately previous period the formerly cooperative subject played D. We then ask how frequently this formerly cooperative subject showed forgiveness by returning to C. We see that in both payoff specifications, forgiveness is similar in the explicit-intentions and the no-error conditions, and lower in the implicit-intentions control.

For Online Publication

Online Appendix A: Decision time analysis regression tables

Table OA1. Predicting cooperation based on decision time. Logistic regression with robust standard errors clustered on subject and pairing.

	(1)	(2)	(3)	(4)
	Last 4 Supergames		All Supergames	
Decision time (log10(sec))	-0.937*** (0.214)		-0.977*** (0.168)	
Decision time (sec)		-0.0911*** (0.0275)		-0.0991*** (0.0223)
b/c	0.418*** (0.0673)	0.444*** (0.0677)	0.371*** (0.0532)	0.388*** (0.0536)
Supergame	0.0420 (0.0371)	0.0383 (0.0376)	0.0395*** (0.0118)	0.0419*** (0.0119)
Period	-0.148*** (0.0143)	-0.147*** (0.0142)	-0.150*** (0.0104)	-0.147*** (0.0103)
Explicit-Intentions	0.141 (0.257)	0.154 (0.258)	0.104 (0.195)	0.109 (0.195)
Implicit-Intentions	-0.879*** (0.220)	-0.814*** (0.221)	-0.763*** (0.161)	-0.698*** (0.160)
Constant	0.116 (0.357)	-0.0663 (0.358)	0.227 (0.206)	0.0120 (0.202)
Observations	11,537	11,604	28,501	28,664
Subject-clusters	338	338	338	338
Pairing-clusters	676	676	1741	1741

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table OA2. Predicting decision time by supergame. Linear regression with robust standard errors clustered on subject and pairing.

	(1) Explicit- Intentions	(2) No-Error	(3) Implicit- Intentions	(4) Implicit- Intentions	(5) Implicit- Intentions	(6) Implicit- Intentions	(7) Implicit- Intentions
	All subjects	All subjects	All subjects	All subjects	All subjects	>67%C Subjects	>67%C Subjects
Supergame	- 0.0126*** (0.00140)	-0.00998*** (0.00219)	0.00672*** (0.00169)	0.0163*** (0.00360)	0.0191*** (0.00396)	-0.000265 (0.00192)	0.00261 (0.00200)
b/c	-0.00536 (0.00782)	-0.0576*** (0.00673)	0.0240** (0.0113)	0.0414*** (0.0132)	0.0368*** (0.0137)	0.0294 (0.0197)	0.0201 (0.0195)
Period	- 0.0093*** (0.00119)	-0.00421*** (0.00111)	-0.0114*** (0.00139)	-0.00858*** (0.00104)	-0.000738 (0.00103)	-0.0102*** (0.00168)	-0.00274 (0.00169)
Your Intended Decision	-0.0208 (0.0144)	-0.0411*** (0.0122)	-0.0923*** (0.0206)	-0.0119 (0.00765)	-0.00924 (0.00892)	0.00537 (0.0140)	0.00211 (0.0124)
Your C Frequency				-0.138*** (0.0535)	-0.129** (0.0564)	-0.588*** (0.221)	-0.496** (0.209)
Supergame X Your C Freq				-0.0199*** (0.00549)	-0.0198*** (0.00612)		
Prev Partner's First Move					-0.0239* (0.0124)		-0.0167 (0.0125)
Partner's Intention Last Period					-0.0159* (0.00823)		0.00647 (0.00962)
Constant	0.374*** (0.0285)	0.493*** (0.0267)	0.199*** (0.0297)	0.160*** (0.0341)	0.107*** (0.0392)	0.548*** (0.142)	0.439*** (0.136)
Observations	6,552	6,466	15,483	15,483	12,545	4,866	3,949
Subject-clusters	84	92	162	162	162	50	50
Pairing-clusters	404	396	941	941	849	452	410
R-squared	0.061	0.155	0.034	0.072	0.074	0.042	0.026

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Online Appendix B: Experimental Instructions

Main treatments (Explicit intentions)

Instructions:

Thank you for participating in this experiment.

Please read the following instructions carefully. If you have any questions, do not hesitate to ask us. Aside from this, no communication is allowed during the experiment.

This experiment is about decision making. You will be randomly matched with other people in the room. None of you will ever know the identity of the others. Everyone will receive a fixed show-up amount of \$10 for participating in the experiment. In addition, you will be able to earn more money based on the decisions you and others make in the experiment. Everything will be paid to you in cash immediately after the experiment.

You will interact numerous times with different people. Based on the choices made by you and the other participants over the course of these interactions, you will receive between \$0 and \$30, in addition to the \$10 show-up amount.

You begin the session with 50 units in your account. Units are then added and/or subtracted to that amount over the course of the session as described below. At the end of the session, the total number of units in your account will be converted into cash at an exchange rate of 30 units = \$1.

The Session:

The session is divided into a series of interactions between you and other participants in the room.

In each interaction, you play a random number of rounds with another person. In each round you and the person you are interacting with can choose one of two options. Once the interaction ends, you get randomly re-matched with another person in the room to play another interaction.

The setup will now be explained in more detail.

The round

In each round of the experiment, the same two possible options are available to both you and the other person you interact with: A or B.

The payoffs of the options (in units)

Option	You will get	The other person will get
--------	-----------------	------------------------------

A: -2 +8

B: 0 0

If your move is A then you will get -2 units, and the other person will get +8 units.

If your move is B then you will get 0 units, and the other person will get 0 units.

Calculation of your income in each round:

Your income in each round is the sum of two components:

- the number of units you get from the move you played
- the number of units you get from the move played by the other person.

Your round-total income for each possible action by you and the other player is thus

		Other person	
		A	B
You	A	+6	-2
	B	+8	0

For example:

If you play A and the other person plays A, you would both get +6 units.

If you play A and the other person plays B, you would get -2 units, and they would get +8 units.

If you play B and the other person plays A, you would get +8 units, and they would get -2 units.

If you play B and the other person plays B, you would both get 0 units.

Your income for each round will be calculated and presented to you on your computer screen.

The total number of units you have at the end of the session will determine how much money you earn, at an exchange rate of 30 units = \$1.

Each round you must enter your choice within 30 seconds, or a random choice will be made.

A chance that the your choice is changed

There is a $7/8$ probability that the move you choose actually occurs. But with probability $1/8$, your move is changed to the opposite of what you picked. That is:

When you choose A, there is a $7/8$ chance that you will actually play A, and $1/8$ chance that instead you play B. The same is true for the other player.

When you choose B, there is a $7/8$ chance that you will actually play B, and $1/8$ chance that instead you play A. The same is true for the other player.

Both players are informed of the moves which actually occur, as well as the moves chosen by each player. Thus with $1/8$ probability, an error in execution occurs, and you know whether the other person's action was what they chose, or an error.

For example, if you choose A and the other player chooses B then:

- With probability $(7/8)*(7/8)=0.766$, no changes occur. You will both be told that your move is A and the other person's move is B, and that you chose A and the other person chose B. You will get -2 units, and the other player will get +8 units.
- With probability $(7/8)*(1/8)=0.109$, the other person's move is changed. You will both be told that your move is A and the other person's move is A, and that you chose A and the other person chose B. You both will get +6 units.
- With probability $(1/8)*(7/8)=0.109$, your move is changed. You will both be told that your move is B and the other person's move is B, and that you chose A and the other person chose B. You will both get +0 units.
- With probability $(1/8)*(1/8)=0.016$, both your move and the other person's moves are changed. You will both be told that your move is B and the other person's move is A, and that you chose A and the other person chose B. You will get +8 units and the other person will get -2 units.

Random number of rounds in each interaction

After each round, there is a $7/8$ probability of another round, and $1/8$ probability that the interaction will end. Successive rounds will occur with probability $7/8$ each time, until the interaction ends (with probability $1/8$ after each round). Once the interaction ends, you will be randomly re-matched with a different person in the room for another interaction. Each interaction has the same setup. You will play a number of such interactions with different people.

Summary

To summarize, every interaction you have with another person in the experiment includes a random number of rounds. After every round, there is a $7/8$ probability of another round. There will be a number of such interactions, and your behavior has no effect on the number of rounds or the number of interactions.

There is a $1/8$ probability that the option you choose will not happen and the opposite option occurs instead, and the same is true for the person you interact with. You will be told which moves actually occur, and you will know what move the other person actually chose.

At the beginning of the session, you have 50 units in your account. At the end of the session, you will receive \$1 for every 30 units in your account.

You will now take a very short quiz to make sure you understand the setup.

The session will then begin with one practice round. This round will not count towards your final payoff.

Screenshot of the information screen:

Explicit intentions

Remaining Time [sec]: 3

ROUND SUMMARY

Your desired move: A
Your actual move: A

Other's desired move: A
Other's actual move: A

Your income this round: 6

OK

This screenshot shows a 'ROUND SUMMARY' window with a light gray background. At the top right, it says 'Remaining Time [sec]: 3'. The main content is centered and lists: 'Your desired move: A', 'Your actual move: A', 'Other's desired move: A', 'Other's actual move: A', and 'Your income this round: 6'. An 'OK' button is at the bottom right.

Implicit intentions:

Remaining Time [sec]: 4

ROUND SUMMARY

Your desired move: B
Your actual move: A

Other's desired move: ?
Other's actual move: A

Your income this round: 6

OK

This screenshot shows a 'ROUND SUMMARY' window with a light beige background. At the top right, it says 'Remaining Time [sec]: 4'. The main content is centered and lists: 'Your desired move: B', 'Your actual move: A', 'Other's desired move: ?', 'Other's actual move: A', and 'Your income this round: 6'. An 'OK' button is at the bottom right.

No error:

Remaining Time [sec]: 11

ROUND SUMMARY

Your move: A
Other's move: A

Your income this round: 6

OK

This screenshot shows a 'ROUND SUMMARY' window with a light gray background. At the top right, it says 'Remaining Time [sec]: 11'. The main content is centered and lists: 'Your move: A', 'Other's move: A', and 'Your income this round: 6'. An 'OK' button is at the bottom right.

Amazon Mechanical Turk test of the framing of the error term

Error condition:

Imagine you are playing a game with another worker on mTurk. In each round of the experiment, the same two possible options are available to both you and the other person you interact with: A or B.

The payoffs of the options (in units)

Option	You	The other person
A	-2	+8
B	0	0

If your move is A then you will get -2 units, and the other person will get +8 units.

If you move is B then you will get 0 units, and the other person will get 0 units.

Calculation of your income in the game:

Your income in the game is the sum of two components:

- the number of units you get from the move you played
- the number of units you get from the move played by the other person.

A chance that the your choice is changed

There is a 7/8 probability that the move you choose actually occurs. But with probability 1/8, your move is changed to the opposite of what you picked. That is:

When you choose A, there is a 7/8 chance that you will actually play A, and 1/8 chance that instead you play B. The same is true for the other player.

When you choose B, there is a 7/8 chance that you will actually play B, and 1/8 chance that instead you play A. The same is true for the other player.

Both players are informed of the moves which actually occur, as well as the moves chosen by each player. Thus with 1/8 probability, an error in execution occurs, and you know whether the other person's action was what they chose, or an error.

[New page]

Imagine the other player chooses A, and the 1/8 probability switch takes effect, so the other player's choice is changed to B. Therefore, you get 0 cents and the other player gets 0 cents from their action.

To what extent do you think the other person intended to pay 0 cents and give you 0 cents?

1 - Did not intend it at all 2 3 4 5 6 7 - Completely intended it

Lottery condition:

Imagine you are playing a game with another worker on mTurk. In each round of the experiment, the same two possible options are available to both you and the other person you interact with: A or B.

The payoffs of the options (in units)

Option	7/8 chance	1/8 chance
A	-2 you, +8 other	0 you, 0 other
B	0 you, 0 other	-2 you, +8 other

When you choose A, there is a 7/8 chance that you will get -2 units, and the other person will get +8 units, and 1/8 chance that you will get 0 units, and the other person will get 0 units. The same is true for the other player.

When you choose B, there is a 7/8 chance that you will get 0 units, and the other person will get 0 units, and 1/8 chance that you will get -2 units, and the other person will get +8 units. The same is true for the other player.

Both players are informed of the moves which actually occur, as well as the moves chosen by each player.

Calculation of your income in the game:

Your income in each round is the sum of two components:

- the number of units you get from the move you played
- the number of units you get from the move played by the other person.

[New page]

Imagine the other player chooses A, and the 1/8 probability takes effect. Therefore, you get 0 cents and the other player gets 0 cents from their action.

To what extent do you think the other person intended to pay 0 cents and give you 0 cents?

1 - Did not intend it at all 2 3 4 5 6 7 - Completely intended it

Ultimatum Game experiment Player 1

Instructions:

Thank you for participating in this experiment.

Please read the following instructions carefully. If you have any questions, do not hesitate to ask us. Aside from this, no communication is allowed during the experiment.

This experiment is about decision making. You will be randomly matched with other people in the room. None of you will ever know the identity of the others. Everyone will receive a fixed show-up amount of \$10 for participating in the experiment. In addition, you will be able to earn more money based on the decisions you and others make in the experiment. Everything will be paid to you in cash immediately after the experiment.

You will interact numerous times with different people. Based on the choices made by you and the other participants over the course of these interactions, you will receive between \$0 and \$30, in addition to the \$10 show-up amount.

You begin the session with 0 units in your account. Units are then added and/or subtracted to that amount over the course of the session as described below. At the end of the session, the total number of units in your account will be converted into cash at an exchange rate of 50 units = \$1.

The Session:

The session is divided into a series of interactions between you and other participants in the room.

In each interaction, you play a random number of rounds with another person. **You will always be in the role of Player 1.** The other people you interact with will be in the role of Player 2. In each round Player 1 will make a decision, and Player 2 will respond. Once the interaction ends, you get randomly re-matched with another person in the room to play another interaction.

The setup will now be explained in more detail.

The round

In each round of the experiment, Player 1 makes a proposal of how to split 10 units between him/herself and Player 2. Player 1 has two options to choose between: A or B.

Option	Player 1 gets	Player 2 gets
A:	5	5
B:	8	2

Once Player 1 chooses a proposal, A or B, Player 2 gets to respond.

- Player 2 can accept the proposal, in which case both players each earn the number of units specified by Player 1's offer.
- Or Player 2 can reject the proposal, in which case both players earn 0 units.

The outcome of each round will be presented to you on your computer screen.

The total number of units you have at the end of the session will determine how much money you earn, at an exchange rate of 50 units = \$1.

A chance that Player 1's choice is changed

There is an 80% chance that the offer Player 1 chooses is the one that is actually made. But 20% of the time, Player 1's offer is changed, and the opposite offer is made to Player 2.

When Player 1 chooses A, there is an 80% chance that Player 1 will actually make offer A, and a 20% chance that instead the actual offer is B.

When Player 1 chooses B, there is an 80% chance that Player 1 will actually make offer B, and a 20% chance that instead the actual offer is A.

If Player 2 accepts, then both players receive the number of units specified by the actual offer. If Player 2 rejects, then both players get 0 units.

Although only the actual decision effects the potential payoffs, Player 2 can base his/her decision on both Player 1's intended offer and the actual offer. Specifically, in every round, Player 2 indicates whether he/she would accept or reject in each of the 4 possible cases:

1. Player 1 intends to offer A, and the actual offer is A
2. Player 1 intends to offer A, but the actual offer is B
3. Player 1 intends to offer B, and the actual offer is B
4. Player 1 intends to offer B, but the actual offer is A

After Player 1 has chosen an offer, and Player 2 has indicated his/her choice in each of these 4 cases, both players are informed of Player 1's intended and actual offers, Player 2's resulting response, and both players' payoffs for the round.

Random number of rounds in each interaction

After each round, there is a $7/8$ probability of another round, and $1/8$ probability that the interaction will end. Successive rounds will occur with probability $7/8$ each time, until the interaction ends (with probability $1/8$ after each round). Once the interaction ends, you will be randomly re-matched with a different person in the room for another interaction. Each interaction has the same setup. You will play a number of such interactions with different people. You will always be in the same role in each round of every interaction.

Summary

To summarize, every interaction you have with another person in the experiment includes a random number of rounds. After every round, there is a $7/8$ probability of another round. In every round, Player 1 makes a proposal to Player 2, which Player 2 can either accept or reject. There will be a number of such interactions, and your behavior has no effect on the number of rounds or the number of interactions.

There is a 20% chance that the offer Player 1 chooses will not be made, and that instead the opposite offer will actually be made. Player 2 can base his/her decision to accept or reject on both the offer Player 1 intended and the offer that is actually made.

At the beginning of the session, you have 0 units in your account. At the end of the session, you will receive \$1 for every 50 units in your account.

You will always be in the role of Player 1.

You will now take a very short quiz to make sure you understand the setup.

The session will then begin with one practice interaction. This interaction will not count towards your final payoff.

Ultimatum Game experiment Player 2

Instructions:

Thank you for participating in this experiment.

Please read the following instructions carefully. If you have any questions, do not hesitate to ask us. Aside from this, no communication is allowed during the experiment.

This experiment is about decision making. You will be randomly matched with other people in the room. None of you will ever know the identity of the others. Everyone will receive a fixed show-up amount of \$10 for participating in the experiment. In addition, you will be able to earn more money based on the decisions you and others make in the experiment. Everything will be paid to you in cash immediately after the experiment.

You will interact numerous times with different people. Based on the choices made by you and the other participants over the course of these interactions, you will receive between \$0 and \$30, in addition to the \$10 show-up amount.

You begin the session with 0 units in your account. Units are then added and/or subtracted to that amount over the course of the session as described below. At the end of the session, the total number of units in your account will be converted into cash at an exchange rate of 50 units = \$1.

The Session:

The session is divided into a series of interactions between you and other participants in the room.

In each interaction, you play a random number of rounds with another person. **You will always be in the role of Player 2.** The other people you interact with will be in the role of Player 1. In each round Player 1 will make a decision, and Player 2 will respond. Once the interaction ends, you get randomly re-matched with another person in the room to play another interaction.

The setup will now be explained in more detail.

The round

In each round of the experiment, Player 1 makes a proposal of how to split 10 units between him/herself and Player 2. Player 1 has two options to choose between: A or B.

Option	Player 1 gets	Player 2 gets
A:	5	5
B:	8	2

Once Player 1 chooses a proposal, A or B, Player 2 gets to respond.

- Player 2 can accept the proposal, in which case both players each earn the number of units specified by Player 1's offer.
- Or Player 2 can reject the proposal, in which case both players earn 0 units.

The outcome of each round will be presented to you on your computer screen.

The total number of units you have at the end of the session will determine how much money you earn, at an exchange rate of 50 units = \$1.

A chance that Player 1's choice is changed

There is an 80% chance that the offer Player 1 chooses is the one that is actually made. But 20% of the time, Player 1's offer is changed, and the opposite offer is made to Player 2.

When Player 1 chooses A, there is an 80% chance that Player 1 will actually make offer A, and a 20% chance that instead the actual offer is B.

When Player 1 chooses B, there is an 80% chance that Player 1 will actually make offer B, and a 20% chance that instead the actual offer is A.

If Player 2 accepts, then both players receive the number of units specified by the actual offer. If Player 2 rejects, then both players get 0 units.

Although only the actual decision effects the potential payoffs, Player 2 can base his/her decision on both Player 1's intended offer and the actual offer. Specifically, in every round, Player 2 indicates whether he/she would accept or reject in each of the 4 possible cases:

5. Player 1 intends to offer A, and the actual offer is A
6. Player 1 intends to offer A, but the actual offer is B
7. Player 1 intends to offer B, and the actual offer is B
8. Player 1 intends to offer B, but the actual offer is A

After Player 1 has chosen an offer, and Player 2 has indicated his/her choice in each of these 4 cases, both players are informed of Player 1's intended and actual offers, Player 2's resulting response, and both players' payoffs for the round.

Random number of rounds in each interaction

After each round, there is a $7/8$ probability of another round, and $1/8$ probability that the interaction will end. Successive rounds will occur with probability $7/8$ each time, until the interaction ends (with probability $1/8$ after each round). Once the interaction ends, you will be randomly re-matched with a different person in the room for another interaction. Each interaction has the same setup. You will play a number of such interactions with different people. You will always be in the same role in each round of every interaction.

Summary

To summarize, every interaction you have with another person in the experiment includes a random number of rounds. After every round, there is a $7/8$ probability of another round. In every round, Player 1 makes a proposal to Player 2, which Player 2 can either accept or reject. There will be a number of such interactions, and your behavior has no effect on the number of rounds or the number of interactions.

There is a 20% chance that the offer Player 1 chooses will not be made, and that instead the opposite offer will actually be made. Player 2 can base his/her decision to accept or reject on both the offer Player 1 intended and the offer that is actually made.

At the beginning of the session, you have 0 units in your account. At the end of the session, you will receive \$1 for every 50 units in your account.

You will always be in the role of Player 2.

You will now take a very short quiz to make sure you understand the setup.

The session will then begin with one practice interaction. This interaction will not count towards your final payoff.

Slow to Anger and Fast to Forgive: Cooperation in an Uncertain World[†]

By DREW FUDENBERG, DAVID G. RAND, AND ANNA DREBER*

We study the experimental play of the repeated prisoner's dilemma when intended actions are implemented with noise. In treatments where cooperation is an equilibrium, subjects cooperate substantially more than in treatments without cooperative equilibria. In all settings there was considerable strategic diversity, indicating that subjects had not fully learned the distribution of play. Furthermore, cooperative strategies yielded higher payoffs than uncooperative strategies in the treatments with cooperative equilibria. In these treatments successful strategies were "lenient" in not retaliating for the first defection, and many were "forgiving" in trying to return to cooperation after inflicting a punishment. (JEL C72, C73, D81)

Repeated games with observed actions have a great many equilibrium outcomes when players are patient, as shown by various folk theorems.¹ These theorems show that cooperative play is *possible* when players are concerned about future rewards and punishments, but since repeated play of a static equilibrium is always an equilibrium of the repeated game, the folk theorems do not predict that cooperation will in fact occur. Intuition and evidence (e.g., Axelrod 1984; Dal Bó 2005; Dreber et al. 2008; Dal Bó and Frechette 2011; Duffy and Ochs 2009) suggest that in repeated games with observed actions players do indeed tend to cooperate when there is a cooperative equilibrium, at least if the gains to cooperation are sufficiently large.²

Outside of the laboratory, actions are often observed with noise: someone who claims they worked hard, or that they were too busy or sick to help, may or may not be telling the truth, and an awkward or inconvenient action may have been well-intentioned; similarly, a self-interested action may wind up accidentally benefiting

*Fudenberg: Department of Economics, Harvard University, 1805 Cambridge Street, Cambridge, MA 02138 (e-mail: dfudenberg@harvard.edu); Rand: Program for Evolutionary Dynamics, Harvard University, 1 Brattle Square Suite 6, Cambridge, MA 02138 (e-mail: drand@fas.harvard.edu); Dreber: Department of Economics, Stockholm School of Economics, Box 6501, 113 83 Stockholm, Sweden (e-mail: anna.dreber@hhs.se). Fudenberg and Rand are co-lead authors. We are grateful to Pedro Dal Bó and Guillaume Frechette for sharing their data and code with us, and for taking the time to reply to our numerous queries. We thank Rob Boyd, Armin Falk, Simon Gächter, Edward Glaeser, Stephanie Hurder, Guido Imbens, Magnus Johannesson, Martin Nowak, Parag Pathak, Thomas Pfeiffer, and John Scholz for helpful conversations and comments, and three referees for useful reports. Research support was provided by National Science Foundation grants SES-064816 and SES-0951462, the Dean of the Faculty of Arts and Sciences, and David Rand is supported by a grant from the John Templeton Foundation.

[†] To view additional materials, visit the article page at <http://dx.doi.org/10.1257/aer.102.2.720>.

¹ Friedman (1971); Aumann and Shapley (1994); and Fudenberg and Maskin (1986).

² Dal Bó (2005) and Dal Bó and Frechette (2011) find that there need not be cooperation when the gain from cooperation is small. Earlier papers had found less cooperative play than these studies, but as Dal Bó (2005) discusses, these papers had various methodological flaws, such as subjects playing versus an experimenter instead of each other, or extremely low payments.

another. In this setting, too, there can be equilibria in which players do better than in the one-shot play of the game, as seen, for example, in the trigger-strategy equilibria constructed by Green and Porter (1984), and indeed a version of the folk theorem applies to repeated games with imperfect public monitoring provided, as shown by Fudenberg, Levine, and Maskin (1994).³ Moreover, while there are evolutionary arguments for cooperation in repeated games with perfectly observed actions, the evolutionary arguments for cooperative equilibria are even stronger with imperfect observations, as the possibility that punishment may be triggered by “mistake” decreases the viability of unrelenting or grim strategies that respond to a single bad observation by never cooperating again.⁴

This paper studies experimental play of the repeated prisoner’s dilemma when intended actions are implemented with noise. Our main goals are to understand whether and when subjects play cooperatively, and also to get a sense of the sorts of strategies that they use. We present evidence from four different payoff specifications for a repeated prisoner’s dilemma, with stage-game actions “Cooperate” (“C”) and “Defect” (“D”) (neutral language was used in the experiment itself); the difference in specifications was the benefit that playing “C” gives to the other player. We consider these four payoff specifications with a continuation probability of $7/8$ and an error rate of $1/8$. As controls we also consider error rates of $1/16$ and 0 (no exogenously imposed error at all) under the “most cooperative” payoff specification, i.e., that with highest benefit to cost ratio. We find that there is much more cooperation in specifications where there is a “cooperative equilibrium” (an equilibrium in which players initially cooperate, and continue to do so at least until a D is observed): players cooperate 49–61 percent of the time in treatments with cooperative equilibria, compared to 32 percent in the specification without cooperative equilibria. In these specifications, we also find that cooperative strategies yielded higher payoffs than uncooperative ones. Conversely, in the one treatment where “Always Defect” is the only equilibrium, this strategy was the most prevalent and had the highest payoff.

As compared to experiments on the prisoner’s dilemma without noise, which we review in Section II, subjects were markedly more lenient (slower to resort to punishment). For example, in the payoff treatments that support cooperation, we find that subjects played C in 63–77 percent of the rounds immediately following their partner’s first defection, compared to only 13–42 percent in the cooperative treatments of Dal Bó and Frechette (2011), Dreber et al. (2008), and our no-error control, which are the no-error games we use for comparison throughout.⁵ Subjects

³Public monitoring means that all players observe a public signal whose distribution depends on the actions played, as in the work of Abreu, Pearce, and Stachetti (1990). The noisy games we study have public monitoring, and satisfy the “identifiability” conditions that are needed for the folk theorem.

⁴Axelrod and Hamilton (1981) applied the ESS solution concept to a small subset of the repeated game strategies. Fudenberg and Maskin (1990, 1994) and Binmore and Samuelson (1992) show that ESS predicts cooperation in the prisoner’s dilemma with noisy observations even when all finite-complexity strategies are allowed. Boyd (1989) notes that noise permits cooperative strategies to be evolutionarily stable, but in his framework noncooperative strategies are stable as well. A parallel literature considers explicit evolutionary dynamics on a restricted strategy space; this includes Feldman and Thomas (1987); Nowak and Sigmund (1990); and Nowak et al. (2004) for the case of observed actions, and Nowak and Sigmund (1993) and Imhof, Fudenberg, and Nowak (2007) for repeated games with noise.

⁵These statistics describe cooperation after one’s partner defects for the first time, rather than overall cooperation. We find similar levels of overall cooperation without noise (78 percent) and with low noise of $1/16$ (82 percent), but substantially less overall cooperation with high noise of $1/8$ (59 percent).

also showed a considerable level of forgiveness (willingness to give cooperation a second chance): in the noisy specifications with cooperative equilibria, our subjects returned to playing C in 18–33 percent of rounds following the breakdown of cooperation, as compared to 6–28 percent in the games without error.⁶ Consistent with these findings, subjects are more likely to condition on play more than 1 round ago in the noisy treatments than in the no-error games; 65 percent use strategies with longer memories in presence of error, compared to 28 percent in the games without error.

In addition to such descriptive statistics, we more explicitly explore which strategies our subjects used, using the techniques of Dal Bó and Frechette (2011).⁷ Relatively few subjects used the strategy “Tit-for-Tat” (TFT), which was prevalent in the Dal Bó and Frechette (2011) experiment, or the strategy “Perfect Tit-for-Tat” (PTFT, also called “Win-Stay, Lose-Shift”), which is favored by evolutionary game theory in treatments where it is an equilibrium *if* players are restricted to strategies that depend only on the previous round’s outcome, as is commonly assumed in that literature.⁸ Instead, the most prevalent cooperative strategies were “Tit-for-2-Tats” (TF2T: punish once after two defections) and “2-Tits-for-2-Tats” (2TF2T: punish two times after two defections), both of which are forgiving, and modified, lenient versions of the Grim strategy, which wait for two or three defections before abandoning cooperation. These results, and other descriptive measures of subjects’ play, show that subjects can and do use strategies that look back more than one round, at least in games with noise.⁹

We find considerable strategic diversity in all settings: no strategy has probability of greater than 30 percent in any treatment, and typically 3 or 4 strategies seem prevalent. Moreover, in every treatment a substantial number of players seem to use the “Always Defect” (ALLD) strategy. That strategy does quite poorly in treatments where most of the subjects are playing strategies that are conditionally cooperative, but it is a best response to the belief that most other subjects play ALLD. Similarly, a substantial fraction of agents cooperate in the treatment where ALLD earns the highest payoff. We take this as evidence that learning was incomplete, and that it is difficult to learn the optimal response to the prevailing distribution of play.

An alternative explanation for the diversity of play is that it reflects a distribution of social preferences, with some subjects preferring to cooperate even if it does not maximize their own payoffs, and others playing to maximize the difference between their partner’s payoff and their own. To test this alternative hypothesis, we had subjects play a dictator game at the end of the session, with payoffs going to

⁶This describes histories in which (i) at least one subject chose C in the first round; (ii) in at least one previous round, the initially cooperative subject chose C while the other subject chose D; and (iii) in the immediately previous round the formerly cooperative subject played D.

⁷In Section III we summarize the findings of Wedekind and Milinski (1996) and Aoyagi and Frechette (2009), who also analyzed the strategies used in a repeated prisoner’s dilemma, and of Engle-Warnick and Slonim (2006), who examined the strategies used in a trust game. With the exception of Aoyagi and Frechette, these experiments considered environments without noise; the introduction of noise leads more information sets to be reached and so makes it easier to distinguish between strategies.

⁸The explicit analysis of evolutionary dynamics becomes quite difficult when longer memories are possible.

⁹We used a continuation probability of $\frac{7}{8}$, instead of the $\frac{1}{2}$ and $\frac{3}{4}$ in Dal Bó and Frechette (2011), to investigate the extent to which players condition on observations before the previous round. When the continuation probability is $\frac{1}{2}$, many interactions will last three or fewer rounds, which makes it hard to study how far back players look in choosing their actions.

recipients recruited at a later experimental session, and we also asked subjects to fill out a post-experimental survey on attitudes and motivations. In another paper (Dreber, Fudenberg, and Rand 2010) we explore this possibility; our main conclusion is that in the treatments with cooperative equilibria, social preferences do not seem to be a key factor in explaining who cooperates and what strategies they use. Leniency and forgiveness seem to be motivated by strategic concerns rather than social preferences.

I. Experimental Design

The purpose of the experimental design is to test what happens when subjects play an infinitely repeated prisoner's dilemma with error. The infinitely repeated game is induced by having a known constant probability that the interaction will continue between two players following each round. We let the continuation probability be $\delta = 7/8$. With probability $1 - \delta$, the interaction ends and subjects are informed that they have been rematched with a new partner. There is also a known constant error probability that an intended move is changed to the opposite move. Our main conditions use $E = 1/8$; we also ran control conditions with $E = 1/16$ and $E = 0$. Subjects were informed when their own move was changed (i.e., when they made an error), but not when the other player's move was changed; they were only notified of the other player's actual move, not the other's intended move. Subjects were informed of all of the above in the experimental instructions, which are included in the online Appendix.

The stage game is the prisoner's dilemma in Figure 1 where the payoffs are denoted in points. Cooperation and defection take the "benefit/cost" (b/c) form, where cooperation means paying a cost c to give a benefit b to the other player, while defection gives 0 to each party; c was fixed at 2, and b/c took the values 1.5, 2, 2.5, and 4 in our four different treatments.¹⁰ Subjects were presented with both the b/c representation and the resulting pre-error payoff matrix, in neutral language. (The choices were labeled A and B as opposed to the "C versus D" choice that is standard in the prisoner's dilemma.) We used the exchange rate of 30 units = \$1. Subjects were given a show-up fee of \$10 plus their winnings from the repeated prisoner's dilemma and an end-of-session dictator game. To allow for negative stage-game payoffs, subjects began the session with an "endowment" of 50 units (in addition to the show-up fee).¹¹ On average subjects made \$22 per session, with a range from \$14 to \$36. Sessions lasted approximately 90 minutes.¹²

A total of 384 subjects participated voluntarily at the Harvard Decision Science Laboratory in Cambridge, Massachusetts. In each session, 12–32 subjects interacted anonymously via computer using the software Z-Tree (Fischbacher 2007) in

¹⁰This payoff specification gives us a simple one-parameter ordering of the treatments; we do not think it is essential for our results. Note that the specification implies that the short-run gain to playing D instead of C is independent of the other player's action. The prisoner's dilemma is more general than this; its defining characteristics are that D is a dominant strategy and that both playing C yields the highest payoff—in particular both playing C should be more efficient than alternating between (C, D) and (D, C).

¹¹No subject ever had fewer than 19 units, and only 4 out of 384 subjects ever dropped below 40 units.

¹²Subjects were given at most 30 seconds to make their decision, and informed that after 30 seconds a random choice would be made. The average decision time was 1.3 seconds, much less than the 30 second limit, and the frequency of random decisions was very low, 0.0055.

Realized payoffs			Expected payoffs		
$b/c = 1.5$	C	D	$b/c = 1.5, E = 1/8$	C	D
	C	D		C	D
	D			D	
$b/c = 2$	C	D	$b/c = 2, E = 1/8$	C	D
	C	D		C	D
	D			D	
$b/c = 2.5$	C	D	$b/c = 2.5, E = 1/8$	C	D
	C	D		C	D
	D			D	
$b/c = 4$	C	D	$b/c = 4, E = 1/8$	C	D
	C	D		C	D
	D			D	
			$b/c = 4, E = 1/16$	C	D
				C	D
				D	

FIGURE 1. PAYOFF MATRICES FOR EACH SPECIFICATION

Note: Payoffs are denoted in points.

a sequence of infinitely repeated prisoner's dilemmas. (See Table 1 for summary statistics on the different treatments.) We conducted a total of 18 sessions between September 2009 and October 2010.¹³ We only implemented one treatment during a given session, so each subject participated in only one treatment. To rematch subjects after the end of each repeated game, we used the turnpike protocol as in Dal Bó (2005). Subjects were divided into two equal-sized groups, A and B. A-subjects only interacted with B-subjects and vice versa, so that no subject ever played twice with another subject, or with a subject who played with a subject they had played with, so that subjects could not influence the play of subjects they interacted with in the future.¹⁴ Subjects were informed about this setup. To implement random game lengths, we pregenerated a sequence of integers t_1, t_2, \dots , according to the specified geometric distribution to use in all sessions, such that in each session every first interaction lasted t_1 rounds, every second interaction lasted t_2 , etc.¹⁵

¹³ All sessions were conducted during the academic year, and all subjects were recruited through the CLER lab at Harvard Business School using the same recruitment procedure. Subject demographics varied relatively little across sessions and treatments. See the online Appendix for demographic summary statistics by session.

¹⁴ Thus, the maximum number of interactions in a session with N subjects was $N/2$, and in each session we ran the maximum number of interactions. This explains why the average number of interactions differs between the different treatments. The average numbers of interactions per subject in Table 1 are not integers because there were multiple sessions per treatment.

¹⁵ The starting place in the sequence of random game lengths that was used in the experiment was picked by the programmer, and the sequence following the chosen starting place had an unusually low number of short games. As a result, the overall distribution of game lengths differed from what would be expected from a geometric distribution, which raises the concern that the subjects could have noticed this and adjusted their play. Analysis of the data

TABLE 1—SUMMARY STATISTICS FOR EACH TREATMENT

	b/c = 1.5 E = 1/8	b/c = 2 E = 1/8	b/c = 2.5 E = 1/8	b/c = 4 E = 1/8	b/c = 4 E = 1/16	b/c = 4 E = 0
Sessions per treatment	3	2	3	4	3	3
Subjects per treatment	72	52	64	90	58	48
Average number of interactions	11	11.5	10.7	11.3	9.9	7.8
Average number of rounds per interaction	8.4	8.3	8.3	8.1	8.0	8.2

Following the end of the series of repeated prisoner's dilemmas, subjects played a dictator game and answered survey questions related to prosocial behavior, motivation, strategies, and demographics. (See Dreber, Fudenberg, and Rand 2010 for more information.)

II. Theoretical and Experimental Background

We begin by analyzing the set of equilibria of the various specifications. In all of the treatments, the only static equilibrium is to defect. In the treatment with $b/c = 1.5$, the only Nash equilibrium is ALLD, while the other treatments all allow cooperative equilibria.¹⁶ As there are no explicit characterization theorems for the entire set of equilibrium outcomes for noisy repeated games with fixed discount factors, our initial analysis focused on a few repeated game strategies that have previously received attention.

In particular, we chose the payoffs so that when $b/c = 4$, the memory-1 strategy PTFT—"Play C if yesterday's outcome was (C,C) or (D,D) and otherwise play D"—is an equilibrium. This strategy has received a great deal of attention in the literature on evolutionary game theory, where it is also called "Win-Stay, Lose-Shift" (Nowak and Sigmund 1993; Wedekind and Milinski 1996; Posch 1999; Imhof, Fudenberg, and Nowak 2007). When both players use PTFT, one play of D (either intentionally or by mistake) leads to one round of (D,D) followed by a return to the equilibrium path; the strategy is called "perfect" because this error-correcting property allows it to be subgame-perfect, in contrast to TFT, which typically is not.¹⁷ PTFT has theoretical appeal because it is error-correcting and has only memory-1, but we conjecture that most subjects will view it as counterintuitive to cooperate after mutual defection, which raises the question of how widely the strategy is actually used.

Standard equilibrium analysis predicts no cooperation when $b/c = 1.5$, but offers little guidance when b/c is large enough that there are cooperative equilibria.

shows, however, that subjects did not become less likely to cooperate in later rounds over the course of the session; see the online Appendix for details.

¹⁶ Because the error term is strictly positive regardless of the actions played, every information set is reached with positive probability, and Nash equilibrium implies sequential rationality. Thus, in the games with errors every Nash equilibrium is a sequential equilibrium, and every pure-strategy Nash equilibrium is equivalent to a perfect public equilibrium.

¹⁷ In the game without errors, PTFT is a subgame-perfect equilibrium if $c < \delta(b - c)$ or $\delta > 1/(b/c - 1)$ which is the case when $b/c = 2.5$ or 4. Analysis of the game with errors shows that PTFT is not an equilibrium when $b/c = 2$ or 2.5, essentially because the errors lower the expected value of cooperation, but PTFT is an equilibrium of the game with errors when $b/c = 4$. See the online Appendix for details.

The evolutionary game theory models of Nowak and Sigmund (1993) and Imhof, Fudenberg, and Nowak (2007) restrict their attention to memory-1 strategies and predict that PTFT will be the most common cooperative strategy when cooperation succeeds. This suggests that we would observe cooperation (and PTFT in particular) when PTFT is an equilibrium at $b/c = 4$, but not at $b/c = 2$ or 2.5 . The evolutionary analysis of Fudenberg and Maskin (1990, 1994) predicts cooperation in all three treatments with cooperative equilibria, but does not provide a precise prediction of what strategies will be played.

Experimental work on repeated games without errors also suggests that cooperation is more likely when it is more beneficial, and in particular that the existence of a cooperative equilibrium is necessary but not sufficient for there to be a substantial amount of cooperation (e.g., Roth and Murnighan 1978; Murnighan and Roth 1983; Feinberg and Husted 1993; Dal Bó 2005; Dreber et al. 2008; Duffy and Ochs 2009; Dal Bó and Frechette 2011; Blonski, Ockenfels, and Spagnolo 2011). Blonski and Spagnolo (2004) proposed that the key to whether cooperation occurs is whether TFT is a best response to a $1/2$ - $1/2$ probability distribution over TFT and ALLD; i.e., whether TFT risk-dominates ALLD in the game with only those two strategies.¹⁸ Dal Bó and Frechette (2011) and Blonski, Ockenfels, and Spagnolo (2011) find empirical support for this risk-dominance criterion in games without noise; and we find the same in a reanalysis of the prisoner's dilemma experiments of Dreber et al. (2008).

The success of the "risk dominance by TFT" criterion in experiments without noise raises the question of whether a similar criterion will explain cooperation in games with noise. One complication is that in our experiment, noise lowers the payoff of TFT against itself sufficiently that TFT is not an equilibrium of the overall game.¹⁹ It is, however, an equilibrium of the 2×2 game where players are restricted to play either TFT or ALLD; in this 2×2 game TFT is only risk dominant if $b/c = 2.5$ or 4 . Thus, to the extent that the risk-dominance criterion extends to games with noise, it predicts substantially more cooperation when $b/c = 2.5$ than when $b/c = 2$.

This combination of observations about equilibria of the game and insights from past experiments leads to our first set of experimental questions:

QUESTION 1: Is cooperation more frequent in treatments where there are cooperative equilibria? Is risk dominance of TFT over ALLD a good predictor of cooperation?

QUESTION 2: Is there substantially more cooperation when cooperation is the outcome of an equilibrium in strategies that base their play on outcomes in the previous round?

The answers to questions 1 and 2 provide some indirect evidence on the strategies that subjects use. To get a more precise understanding of the particular strategies

¹⁸This is the case when the present value of cooperation is sufficiently high compared to the loss caused by one period of (C,D). Note that Blonski, Ockenfels, and Spagnolo (2011) offer an alternative theoretical justification for this equilibrium selection criterion.

¹⁹See the online Appendix for TFT equilibrium calculations.

being employed by subjects, we examine more direct evidence on their play. This leads to our four additional experimental questions:

QUESTION 3: *What strategies do subjects use in the noisy prisoner's dilemma? Do subjects use PTFT when it is an equilibrium strategy?*

QUESTION 4: *How do the strategies used vary with the gains to cooperation?*

QUESTION 5: *How do the strategies used vary with the level of noise?*

III. Methodology

The general theory of repeated games, like that of extensive form games, views strategies as complete contingent plans, which specify how the player will act in every possible information state. In practice, cognitive constraints may lead subjects to use relatively simple strategies, corresponding to automata with a small number of internal states. It is unclear, however, what a priori restrictions one should impose on subjects' strategies, and one of the goals of our experiment is to let the data reveal what sorts of strategies are actually used. For this reason we did not want to use the "strategy method," where subjects are asked to pick a strategy that is implemented for them: the full set of strategies is infinite, so forcing subjects to choose a strategy that depends only on the previous round's outcome (memory-1) is much too restrictive while allowing for all strategies that depend on the last two periods (memory-2) gives too large a strategy set to present explicitly. In addition, we would like to consider some simple strategies such as Grim, which can be viewed as a two-state automata but has arbitrarily long memory, as its current play depends on whether "D" was played in any previous round. As the data cannot discriminate between all possible repeated game strategies, we used a combination of prior intuition, survey responses, and data analysis to identify a small set of strategies that seem to best describe actual play.

There has been comparatively little past work on identifying the strategies subjects use in repeated game experiments. In the repeated prisoner's dilemma, Wedekind and Milinski (1996) note that subjects rarely play C in the round after they played D and the opponent played C, and take this as evidence of PTFT, but since subjects rarely played C following (D,D), their data seems more consistent with some sort of grim strategy. Dal Bó and Frechette (2011) used maximum likelihood to estimate the proportions of subjects using one of six ex ante relevant strategies. They find that ALLD and TFT account for the majority of their data. Aoyagi and Frechette (2009) study experimental play of a prisoner's dilemma where subjects do not observe their partner's actions but instead observe a noisy symmetric signal of it. The signal is a real number, and has the same distribution under (C,D) and (D,C), so that commonly discussed prisoner's dilemma strategies such as TFT are not implementable. They find that subjects play "trigger strategies" of memory-1, except in the limit no-noise case where signals from two rounds ago also have an impact. Engle-Warnick and Slonim (2006) study the strategies used in a repeated sequential-move trust game by counting how many observations of a subject's play in a given interaction is described exactly by a given strategy. They find that in most of the interactions, the investors choose

actions that are consistent with a grim trigger strategy, while the play of the trustees is more diverse. Camera, Casari, and Bigoni (2010) on the other hand find little evidence of grim trigger strategies on the subject level when subjects are put randomly in groups of four to play an indefinitely repeated PD where in each round they are matched randomly with one subject in the group. Even though behavior at the aggregate looks like grim trigger, individual behavior is far more heterogeneous.

An advantage of studying repeated games with errors is that we can more easily identify different strategies: in the absence of errors a number of repeated game strategies are observationally equivalent; for example, if a pair of subjects cooperates with each other in every round, we see no data on how they would have responded to defections. Thus, the introduction of errors has a methodological advantage as well as a substantive one, as the errors will lead more histories to occur and thus make it easier to distinguish between histories.

To have any hope of inferring the subjects' strategies from their play, we must focus our attention on a subset of the infinitely many repeated game strategies. We begin with strategies that have received particular attention in the theoretical literature: ALLD, ALLC, Grim,²⁰ TFT, and PTFT. Because one round of punishment is only enough to sustain cooperation in one of our four treatments (when $b/c = 4$) we also include modifications of TFT and PTFT that react to D with two rounds of defection; we call these 2TFT and 2PTFT. We also include the strategy T2 used by Dal Bó and Frechette (2011).²¹

To inform our extension of this strategy set, we asked subjects to describe their strategies in a post-experimental survey. Several regularities emerged from these descriptions. Many subjects reported "giving the benefit of the doubt" to an opponent on the first defection, assuming that it was a result of noise rather than purposeful malfeasance; only after two or three defections by their partner would they switch to defection themselves.²² We refer to this slow-to-anger behavior as "leniency." None of the strategies mentioned above is lenient; note that leniency requires looking further into the past than permitted by memory-1 strategies. Subjects also varied in the extent to which they reported being willing to return to cooperation following a partner's defection. We refer to this strategic feature as "forgiveness," which is an often-discussed aspect of TFT (as opposed to Grim, for example); 2TFT also shows forgiveness, as do PTFT and 2PTFT, although only following mutual defection.

In response to the subjects' strategy descriptions, we added several lenient strategies to our analysis. Because our games were on average only eight rounds in length, we have limited power to explore intermediate levels of forgiveness between TFT and Grim, so we restrict strategies to either forgive after one to three rounds or to never forgive (as with Grim and its lenient variants).²³

²⁰ As in Dal Bó and Frechette (2011), our specification of Grim begins by playing C and then switches permanently to D as soon as either player defects.

²¹ 2TFT initially plays C, then afterwards plays C if opponent has never played D or if the opponent played C in both of the previous two rounds. A defection by the partner triggers two rounds of punishment defection in both 2TFT and T2. T2, however, automatically returns to C following the two Ds, regardless of the partner's play during this time, while 2TFT only returns to C if the partner played C in both of the "punishment rounds." Additionally, T2 begins its punishment if either player defects, whereas 2TFT responds only to the partner's defection.

²² Subjects' free-response descriptions of their strategies are reproduced in the online Appendix.

²³ More generally, the average length of eight interactions imposes restrictions on our ability to estimate the extent to which subjects condition on long histories. This constraint reflects a tradeoff between our interest in the way subjects use past history and our desire to limit sessions to 90 minutes (to avoid fatiguing the subjects) while

As strategies that are both lenient and forgiving, we include TFT variants that switch to defection only after the other player chooses D multiple times in a row, considering TF2T (the player plays D if the partner's last two moves were both D) and TF3T (the player plays D if the partner's last three moves were D). For strategies that are lenient but not forgiving, we include Grim variants that wait for multiple rounds of D (by either player) before switching permanently to defection, considering Grim2 (the player waits for two consecutive rounds in which either player played D) and Grim3 (the player waits for three consecutive D rounds). We include three strategies that punish twice (intermediate to TFT's one round of punishment and Grim's unending punishment) but can be implemented by conditioning only on the last three rounds. These are T2 and 2TFT, discussed above, and 2TF2T ("2 Tits for 2 Tats"), which waits for the partner to play D twice in a row, and then punishes by playing D twice in a row. Because we do not include strategies that punish for a finite number of rounds greater than two, our estimated share of "Grim" strategies may include some subjects who use such strategies with more than two rounds of punishment.

Other subjects indicated that they used strategies that tried to take advantage of the leniency of others by defecting initially and then switching to cooperation. Thus, we consider "exploitative" versions of our main cooperative strategies that defect on the first move and then return to the strategy as normally specified: D-TFT²⁴, D-TF2T, D-TF3T, D-Grim2, and D-Grim3.²⁵ Because TF2T appears prevalent in many treatments, we also look at whether subjects used the strategy that alternates between D and C (DC-Alt), as this strategy exploits the leniency and forgiveness of TF2T. Lastly, some subjects reported playing strategies that give the first impression of being cooperative and then switch to defection, hoping the partner will assume the subsequent Ds are due to error. Therefore we include a strategy that plays C in the first round and D thereafter (C-to-ALLD). Each strategy is described verbally in Table 2; complete descriptions are given in the online Appendix.

To assess the prevalence of each strategy in our data, we follow Dal Bó and Frechette (2011) and suppose that each subject chooses a fixed strategy at the beginning of the session (or alternatively for the last four interactions, when we restrict attention to those observations),²⁶ and moreover that in addition to the extrinsically imposed execution error, subjects make mistakes when choosing their intended action, so every sequence of choices (e.g., of intended actions) has positive probability.²⁷ More specifically, we suppose that if subject i uses strategy s , her chosen action in round r of interaction k is C if $s_{ikr}(s) + \gamma \varepsilon_{ikr} \geq 0$, where $s_{ikr}(s) = 1$ if strategy s says to play C in round r of interaction k given the history to that point, and $s_{ikr}(s) = -1$ if s says to play D. Here ε_{ikr} is an error term that is independent

allowing them to play enough interactions to have some chance to learn. Fortunately, the average length of eight was enough to provide convincing evidence that subjects can use forgiving strategies with memory greater than one.

²⁴ Boyd and Lorberbaum (1987) call this strategy "Suspicious Tit for Tat."

²⁵ ALLD and D-Grim are identical except when the other player plays C in the first round, and you mistakenly also play C in the first round: here D-Grim cooperates while ALLD defects. Thus, we do not include D-Grim in our analysis as we do not have a sufficient number of observations per subject to differentiate between the two.

²⁶ We found that conducting the MLE supposing that subjects pick a fixed strategy at the beginning of each interaction, as opposed to using the same strategy throughout the session, gave qualitatively similar results to those presented below.

²⁷ Recall that we, unlike our subjects, observe the intended actions as well as the implemented ones. We use this more informative data in our estimates.

TABLE 2—DESCRIPTIONS OF THE 20 STRATEGIES CONSIDERED

Strategy	Abbreviation	Description
Always Cooperate	ALLC	Always play C
Tit-for-Tat	TFT	Play C unless partner played D last round
Tit-for-2-Tats	TF2T	Play C unless partner played D in both of the last 2 rounds
Tit-for-3-Tats	TF3T	Play C unless partner played D in all of the last 3 rounds
2-Tits-for-1-Tat	2TFT	Play C unless partner played D in either of the last 2 rounds (2 rounds of punishment if partner plays D)
2-Tits-for-2-Tats	2TF2T	Play C unless partner played 2 consecutive Ds in the last 3 rounds (2 rounds of punishment if partner plays D twice in a row)
T2	T2	Play C until either player plays D, then play D twice and return to C (regardless of all actions during the punishment rounds)
Grim	Grim	Play C until either player plays D, then play D forever
Lenient Grim 2	Grim2	Play C until 2 consecutive rounds occur in which either player played D, then play D forever
Lenient Grim 3	Grim3	Play C until 3 consecutive rounds occur in which either player played D, then play D forever
Perfect Tit-for-Tat / Win-Stay-Lose-Shift	PTFT	Play C if both players chose the same move last round, otherwise play D
Perfect Tit-for-Tat with 2 rounds of punishment	2PTFT	Play C if both players played C in the last 2 rounds, both players played D in the last 2 rounds, or both players played D 2 rounds ago and C last round. Otherwise play D
Always Defect	ALLD	Always play D
False cooperator	C-to-ALLD	Play C in the first round, then D forever
Exploitative Tit-for-Tat	D-TFT	Play D in the first round, then play TFT
Exploitative Tit-for-2-Tats	D-TF2T	Play D in the first round, then play TF2T
Exploitative Tit-for-3-Tats	D-TF3T	Play D in the first round, then play TF3T
Exploitative Grim2	D-Grim2	Play D in the first round, then play Grim2
Exploitative Grim3	D-Grim3	Play D in the first round, then play Grim3
Alternator	DC-Alt	Start with D, then alternate between C and D

across subjects, rounds, interactions, and histories, γ parameterizes the probability of mistakes, and the density of the error term is such that the overall likelihood that subject i uses strategy s is

$$(1) \quad p_i(s) = \Pi_k \Pi_r \left(\frac{1}{1 + \exp(-s_{ikr}(s)/\gamma)} \right)^{y_{ikr}} \left(\frac{1}{1 + \exp(s_{ikr}(s)/\gamma)} \right)^{1-y_{ikr}},$$

where y_{ikr} is 1 if the subject chose C and 0 if the subject chose D.²⁸

To better understand the mechanics of the specification, suppose that an interaction lasts w rounds, that in the first round the subject chose C, the first-round outcome was that the subject played C and her partner played D, and in the second

²⁸ Thus, the probability of an error in implementing one's strategy is $1/(1 + \exp(1/\gamma))$. Note that this represents error in intention, rather than the experimentally imposed error in execution. This formulation assumes that all strategies have an equal rate of implementation error. In the online Appendix we show that the MLE estimates of strategy shares are robust to allowing each strategy have a different value of γ .

round the subject chose D. Then for strategy $s = \text{TFT}$, which plays C in the first round, and plays D in the second round following (C,D), the likelihood of the subject's play is the probability of two "no-error" draws. This is the same probability that we would assign to the overall sequence of the subject's play given the play of the opponent—it makes no difference whether we compute the likelihood round by round or for the whole interaction.

For any given set of strategies S and proportions p , (probability distribution on S) we then derive the likelihood for the entire sample, namely $\sum_{i \in I} \ln(\sum_{s \in S} p(s) p_i(s))$. Note that the specification assumes that all subjects are *ex ante* identical with the same probability distribution over strategies and the same distribution over errors; one could relax this at the cost of adding more parameters. Because p describes a distribution over strategies, this likelihood function implies that in a very large sample we expect fraction $p(s)$ of subjects to use strategy s , though for finite samples there will be a nonzero variance in the population shares. We use maximum likelihood estimation (MLE) to estimate the prevalence of the various strategies, and bootstrapping to associate standard errors with each of our frequency estimates. We construct 100 bootstrap samples for each treatment by randomly sampling the appropriate number of subjects with replacement. We then determine the standard deviation of the MLE estimates for each strategy frequency across the 100 bootstrap samples.

To investigate the validity of this estimation procedure, we tested it on simulated data. For a given strategy frequency distribution, we assigned strategies to 3 groups of 20 computer agents. We then generated a simulated history of play across four interactions by randomly pairing members of each group to play games with representative lengths from the game length sequence used in the experiment ($t_1 = 5$, $t_2 = 11$, $t_3 = 8$, $t_4 = 9$). As in the main experimental treatments, we included a $1/8$ probability of error, and recorded both the intended and actual action of each agent. We generated simulated data in this way using strategy distributions similar to those estimated from the experimental data (see Table 3), and then used the MLE method described above to estimate the strategy frequencies. The MLE results were consistent with the actual strategy frequencies, giving us confidence in the estimation procedure.²⁹

IV. Results

We begin by examining behavior in the four treatments with $E = 1/8$; we will then compare this to the behavior in the controls with $E = 0$ and $E = 1/16$ using $b/c = 4$, which is the ratio most favorable to cooperation.

Examining play in the first round of each interaction, as displayed in Figure 2, suggests that there was some learning, except perhaps when $b/c = 1.5$; this is confirmed by the statistical analysis reported in Appendix B. To reduce the potential effects of learning while striking a balance with the need for data, our analysis will focus on how subjects played in the last four interactions of the session, which is roughly the last third of each session.³⁰

²⁹ See Appendix A for MLE results using simulated data.

³⁰ Our results are not sensitive to this particular cutoff. Using either the last 6 or last 2 interactions instead yields very similar results. See the online Appendix for details.




FIGURE 2. FIRST ROUND COOPERATION OVER THE COURSE OF THE SESSION, BY PAYOFF SPECIFICATION




FIGURE 3. FIRST ROUND AND OVERALL COOPERATION BY PAYOFF SPECIFICATION, AVERAGED OVER THE LAST FOUR INTERACTIONS OF EACH SESSION

Notes: See Table 4 for a list of the cooperation frequencies displayed here. Error bars indicate standard error of the mean, clustered on subject and interaction pair.

QUESTION 1: *Is cooperation more frequent in treatments where there are cooperative equilibria? Is risk dominance of TFT over ALLD a good predictor of cooperation?*

Figure 3 reports both cooperation in the first round of the last four interactions and the average cooperation over the last four interactions as a whole, which can depend on the relationship between the two subjects' strategies and also on possible random errors.³¹ We see that there is markedly less cooperation when $b/c = 1.5$, both in the first round (1.5 versus 2, $p = 0.016$; 1.5 versus 2.5, $p = 0.003$; 1.5 versus 4,

³¹ For each pairwise b/c comparison, we report the results of a logistic regression over first-round/all individual decisions, with a b/c value dummy as the independent variable, clustered on both subject and interaction pair.

TABLE 3—MAXIMUM LIKELIHOOD ESTIMATES USING THE LAST 4 INTERACTIONS OF EACH SESSION

	b/c = 1.5	b/c = 2	b/c = 2.5	b/c = 4
ALLC	0 (0)	0.03 (0.03)	0 (0.02)	0.06* (0.03)
TFT	0.19*** (0.05)	0.06 (0.04)	0.09** (0.04)	0.07** (0.03)
TF2T	0.05 (0.03)	0 (0)	0.17** (0.06)	0.20*** (0.07)
TF3T	0.01 (0.01)	0.03 (0.03)	0.05 (0.05)	0.09** (0.04)
2TFT	0.06 (0.04)	0.07* (0.04)	0.02 (0.02)	0.03 (0.02)
2TF2T	0 (0.02)	0.11** (0.05)	0.11* (0.06)	0.12** (0.05)
Grim	0.14*** (0.04)	0.07 (0.05)	0.11** (0.04)	0.04* (0.02)
Grim2	0.06* (0.03)	0.18*** (0.06)	0.02 (0.03)	0.05* (0.03)
Grim3	0.06 (0.03)	0.28*** (0.08)	0.24*** (0.07)	0.11*** (0.04)
ALLD	0.29*** (0.06)	0.17*** (0.06)	0.14*** (0.04)	0.23*** (0.04)
D-TFT	0.14*** (0.05)	0 (0)	0.05* (0.03)	0 (0)
Gamma	0.46*** (0.02)	0.5*** (0.03)	0.49*** (0.03)	0.43*** (0.02)

Notes: All payoff specifications use error rate $E = 1/8$. Bootstrapped standard errors (shown in parentheses) used to calculate p-values.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

$p = 0.004$) and overall (1.5 versus 2, $p = 0.001$; 1.5 versus 2.5, $p < 0.001$; 1.5 versus 4, $p < 0.001$). Conversely, we see little difference in first-round cooperation between the three treatments with cooperative equilibria (2 versus 2.5, $p = 0.71$; 2 versus 4, $p = 0.83$; 2.5 versus 4, $p = 0.87$); and while there is an increase in overall cooperation going from $b/c = 2$ to $b/c = 2.5$, this increase is smaller than that between $b/c = 1.5$ and $b/c = 2$ and is only marginally significant ($p = 0.058$). Moreover, there is no significant difference in overall cooperation between $b/c = 2.5$ and $b/c = 4$ ($p = 0.73$). The reason there is about the same amount of initial cooperation in $b/c = 2$ and 2.5, yet somewhat more overall cooperation in the latter case, seems related to the fact that subjects are more forgiving in the latter treatment, as seen in the discussion of Questions 3–4. Because the largest difference in cooperation occurs between $b/c = 1/5$ and $b/c = 2$, as opposed to between $b/c = 2$ and $b/c = 2.5$, the data do not show the strong support for risk dominance of TFT as the key determinant of the level of cooperation in games with noise that was seen in studies of games without noise.

QUESTION 2: *Is there substantially more cooperation when cooperation is the outcome of an equilibrium in strategies that base their play on outcomes in previous round?*

Indeed, we see a substantial amount of cooperation when $b/c = 2$ and 2.5, even though cooperative equilibria in these treatments require memory 2 or more.³² To the extent that play resembles an equilibrium of the repeated game, these results are a first sign that the predictions of the memory-1 restriction are not consistent with the data.

QUESTION 3: *What strategies do subjects use in the noisy prisoner's dilemma? Do subjects use PTFT when it is an equilibrium strategy?*

We now report the results of the MLE analysis of strategy choice, examining the last four interactions of each session.³³ We consider 20 strategies in total (Table 2): the fully cooperative strategies ALLC, TFT, TF2T, TF3T, 2TFT, 2TF2T, Grim, Grim2, Grim3, PTFT, 2PTFT, and T2, which always play C against themselves in the absence of errors; the fully noncooperative strategies ALLD and D-TFT, which always play D against themselves in the absence of errors; and the partially cooperative strategies C-to-ALLD, D-TF2T, D-TF3T, D-Grim2, D-Grim3, and DC-Alt, which play a combination of C and D against themselves in the absence of error.³⁴ Of these, only 11 are present at frequencies significantly greater than 0 in at least one payoff specification: the cooperative strategies ALLC, TFT, TF2T, TF3T, 2TFT, 2TF2T, Grim, Grim2, and Grim3, and the noncooperative strategies ALLD and D-TFT.³⁵

Thus, we restrict our attention to these 11 strategies (Table 3). We do not find any evidence of subjects using PTFT in any payoff specification—PTFT never received a positive weight in any of the bootstrapped samples. In the treatments with cooperative equilibria, the most common cooperative strategies TF2T, TF3T, Grim2, and Grim3 are all lenient.

Note that in all treatments the MLE assigns a substantial share to strategies that depend on outcomes from more than one period ago. To provide additional evidence for this conclusion, we use a logistic regression to test whether subjects condition play on their partner's decision two rounds ago. Bias introduced by heterogeneity presents a potential challenge for this approach: the other's play two rounds ago interacts with own play two rounds ago to determine the history in the previous round, so other's play two rounds ago could have a spuriously significant coefficient in a heterogeneous population of subjects all of whom use memory-1

³²This theoretical result is robust to considering memory-1 strategies that forget the state with some nonzero probability. Regardless of the probability of forgetting, TFT and PTFT are not equilibria at $E = 1/8$ and $b/c = 2$ or $b/c = 2.5$. See the online Appendix for further discussion.

³³Analysis of simulated data suggests that MLE on four interactions lasting on average eight rounds detects strategies whose frequencies are five percent or higher, but that lower-frequency strategies may not be detected.

³⁴Following Dal Bó and Frechette, we consider only pure strategies in the MLE estimation. The evolutionary game theory literature suggests the consideration of "generous tit-for-tat" (GTFT) (Nowak and Sigmund 1990), which cooperates with probability strictly between 0 and 1 if the partner defected last period, and otherwise cooperates. Simulations show that this strategy will be identified by our estimation as playing TF3T, but that expanding our MLE procedure to include stochastic strategies can differentiate GTFT from TF3T. Doing so suggests that the majority of leniency observed in our data is not in fact the result of stochastic memory-1 strategies, but is rather due to lenient strategies with longer memories. The online Appendix reports these results, which should be viewed as a first step towards understanding how to test for mixed strategies; we hope to explore the issues posed by mixed strategies in greater detail in future work.

³⁵See Appendix C for the estimates and standard errors for the full set of 20 strategies. At the request of the referees, we explored versions of TFT and PTFT that forget the state with some nonzero probability, versions of TFT and Grim that ignore a defection in the first round, and a family of strategies that cooperate until the fraction of D by their partner passes some threshold. None of these strategies was present at frequencies significantly greater than 0. See the online Appendix for the details of these robustness checks.

strategies. To control for bias introduced by heterogeneity, we include controls for the type of the player making the decision, as in Aoyagi and Frechette (2009).³⁶ We conduct a logistic regression with correlated random effects, regressing own decision in round t against own play in round $t - 1$, other's play in $t - 1$, own play in $t - 2$ and other's play in $t - 2$, and including controls for b/c ratio and own average frequency of first round cooperation and overall cooperation, both over the last four interactions.³⁷ Consistent with the use of longer memories, we find a significant effect of other's play two rounds ago (coeff = 1.01, $p < 0.001$). This result supports the conclusion that many subjects are conditioning on more than only the last round.

QUESTION 4: *How do the strategies used vary with the gains to cooperation?*

The strategies employed by subjects clearly vary according to the gains from cooperation. This can be seen from descriptive statistics analyzing aggregate behavior as well as from the MLE analysis, both of which are summarized in Table 4. Three trends are apparent.

First, cooperation is significantly lower at $b/c = 1.5$ than at the higher b/c ratios, as shown in Table 4 and visualized in Figure 3. Consistent with this observation, Table 4 also shows that the share of the noncooperative strategies ALLD and D-TFT is 43 percent when $b/c = 1.5$, which is substantially and significantly higher than in other b/c conditions ($b/c = 2$, 17 percent, $p < 0.001$; $b/c = 2.5$, 19 percent, $p < 0.001$; $b/c = 4$, 23 percent, $p = 0.001$).³⁸

Second, leniency also increases when moving from $b/c = 1.5$ to the higher b/c ratios. To get a measure of leniency distinct from the MLE estimates, we examine all histories in which both subjects played C in all but the previous round, while in the previous round one subject played D.³⁹ We then ask how frequently the subject who had hitherto cooperated showed leniency by continuing to cooperate despite the partner's defection.⁴⁰ At $b/c = 1.5$, 17 percent of histories show leniency, compared to the significantly higher values of 63 percent at $b/c = 2$ ($b/c = 1.5$ versus $b/c = 2$, $p = 0.001$), 67 percent at $b/c = 2.5$ ($b/c = 1.5$ versus $b/c = 2.5$, $p < 0.001$) and 66 percent at $b/c = 4$ ($b/c = 1.5$ versus $b/c = 4$, $p < 0.001$). No significant difference in leniency exists among the higher b/c ratios ($p > 0.20$ for all comparisons). Thus, leniency increases across the transition from $b/c = 1.5$ to $b/c = 2$. Analyzing strategy frequencies paints a similar picture. The combined frequency of the lenient strategies ALLC, TF2T, TF3T,

³⁶For a general treatment of the topic, see Chamberlain (1980) and Heckman (1981).

³⁷When we simulate data for various combinations of memory-1 strategies, we find that this regression returns a significant coefficient on partner's action in $t - 2$ no more often than predicted by chance. We also find that when simulating only memory-1 strategies, the size of the estimated coefficient of play two rounds ago is at least an order of magnitude smaller than the coefficients for play last round, in contrast to the estimates on the experimental data; see the online Appendix for a detailed discussion of these issues. Note that as the regression conditions on play two rounds ago, it necessarily omits decisions made in the first two rounds of each interaction.

³⁸For each pairwise comparison of aggregated MLE coefficients, we report the results of a two-sample t -test using bootstrapped standard errors of the aggregated coefficients.

³⁹We also include second-round decisions in which the first round's outcome was CD.

⁴⁰For each pairwise b/c comparison of aggregate descriptive statistics, we report the results of a logistic regression over all decisions in the relevant histories, with a b/c dummy as the independent variable, clustered on both subject and interaction pair.

TABLE 4—DESCRIPTIVE STATISTICS OF AGGREGATE BEHAVIOR, AS WELL AS AGGREGATED MLE FREQUENCIES FROM TABLE 3

	b/c = 1.5	b/c = 2	b/c = 2.5	b/c = 4
<i>Descriptive statistics</i>				
Percent C first round	54%	75%	79%	76%
Percent C all rounds	32%	49%	61%	59%
Leniency	29%	63%	67%	66%
Forgiveness	15%	18%	33%	32%
<i>MLE aggregation</i>				
Cooperative strategies	57%	83%	81%	77%
Lenient strategies	18%	62%	60%	63%
Forgiving strategies	31%	29%	44%	57%

Notes: All specifications use $E = 1/8$. The descriptive statistics for leniency and forgiveness are defined in the text. For MLE aggregation, all strategies other than ALLD and D-TFT are cooperative; lenient strategies are TF2T, TF3T, 2TF2T, Grim2, and Grim3; and forgiving strategies are TFT, TF2T, TF3T, 2TFT, and 2TF2T.

2TF2T, Grim2, and Grim3 is 18 percent at $b/c = 1.5$, which is significantly less than at $b/c = 2$ (62 percent, $p < 0.001$), $b/c = 2.5$ (60 percent, $p < 0.001$) or $b/c = 4$ (63 percent, $p < 0.001$).

The gains to cooperation also influence the frequency of forgiveness. Forgiveness is more complicated to define and measure, as it describes a more complex pattern of behavior: to us it means that the players are initially cooperating, that one of them then defects, leading the other player to “punish” the initial defector, and finally that the punishing player relents and returns to cooperation. To develop an operational measure of histories where forgiveness occurs, we first identify all histories in which (i) at least one subject chose C in the first round, (ii) in at least one previous round, the initially cooperative subject chose C while the other subject chose D, and (iii) in the immediately previous round the formerly cooperative subject played D. We then ask how frequently this formerly cooperative subject showed forgiveness by returning to C. For example, if the outcome in the first round is (C, D), and the first player plays D in the second round, and C in the third, we would say that the first player had “forgiven” the second player. We find significantly less forgiveness at $b/c = 1.5$ (15 percent) and $b/c = 2$ (18 percent) compared to $b/c = 2.5$ (33 percent) and $b/c = 4$ (32 percent) (1.5 versus 2.5, $p < 0.001$; 1.5 versus 4, $p < 0.001$; 2 versus 2.5, $p = 0.008$; 2 versus 4, $p = 0.007$). Thus, forgiveness increases significantly when b/c increases from 2 to 2.5. This is confirmed again by examining strategy frequencies. The forgiving strategies ALLC, TFT, TF2T, TF3T, 2TFT, and 2TF2T are less common at $b/c = 1.5$ (31 percent) and $b/c = 2$ (29 percent) than at $b/c = 2.5$ (44 percent) and $b/c = 4$ (57 percent) (1.5 versus 2.5, $p = 0.061$; 1.5 versus 4, $p < 0.001$; 2 versus 2.5, $p = 0.054$; 2 versus 4, $p < 0.001$; 2.5 versus 4, $p = 0.054$).

QUESTION 5: *How do the strategies used vary with the level of noise?*

To explore how play varies with the error rate, we now examine our two additional control treatments using $b/c = 4$ with $E = 1/16$ and $E = 0$, and compare

them to our results using $b/c = 4$ at $E = 1/8$. We begin by asking whether subjects condition on their partner's play two rounds ago in the last four interactions of each session, using a logistic regression with correlated random effects and regressing own decision in round t against other's play in round $t - 2$, own play in $t - 2$, other's play in $t - 1$, own play in $t - 1$, own first-round cooperation frequency in the last four interactions, and own frequency of cooperation in all rounds of the last four interactions. As with the $E = 1/8$ treatments, we find a highly significant and sizable dependence on other's play two rounds ago for $E = 1/16$ (coeff = 1.221, $p < 0.001$); while in the no-error $E = 0$ control, however, we find no significant dependence on other's play two rounds ago (coeff = 0.387, $p = 0.247$). This provides our first direct evidence that the presence of noise plays an important role in strategy selection, promoting more complicated strategies.

Next, we present the MLE results for each strategy in Table 5, as well as aggregated MLE results and descriptive statistics in Table 6. As shown in Table 6, overall cooperation is lower at $E = 1/8$ (59 percent) than in the lower error conditions, and these differences are statistically significant ($E = 1/16$, 82 percent, $p < 0.001$; $E = 0$, 78 percent, $p = 0.002$).⁴¹ Considering cooperation in the first round, there is somewhat less cooperation at $E = 1/8$, but the differences are smaller and either not significant or just marginally so ($E = 1/8$, 76 percent, $E = 1/16$, 87 percent, $E = 0$, 83 percent; $1/8$ versus $1/16$, $p = 0.050$, $1/8$ versus 0 , $p = 0.346$). We find no significant difference at $E = 0$ compared to $E = 1/16$ in either overall cooperation ($p = 0.486$) or first-round cooperation ($p = 0.338$). Complementing this aggregate analysis, the share of the noncooperative strategies ALLD and D-TFT is significantly larger at $E = 1/8$ (23 percent) compared to $E = 1/16$ (11 percent, $p = 0.001$), and larger at $E = 1/8$ than $E = 0$ (16 percent, $p = 0.191$) although the difference between $E = 0$ and $E = 1/8$ is not significant.⁴² There is also no significant difference between $E = 0$ and $E = 1/16$ ($p = 0.297$).

Turning to leniency, we examine cooperation frequency in the subset of histories in which leniency is possible, as described above in response to Question 5. At $E = 0$, 42 percent of the eligible histories show leniency, compared to the significantly higher value of 77 percent at $E = 1/16$ ($p = 0.001$) and the marginally significantly higher value of 66 percent at $E = 1/8$ ($p = 0.052$). We also find that marginally significantly more eligible histories showed leniency at $E = 1/16$ than at $E = 1/8$ ($p = 0.071$). We see similar results when analyzing strategy frequencies. The combined frequency of lenient strategies ALLC, TF2T, TF3T, 2TF2T, Grim2, and Grim3 is significantly lower at $E = 0$ (40 percent) than at $E = 1/16$ (82 percent, $p < 0.001$) or $E = 1/8$ (63 percent, $p < 0.001$). As with the analysis of histories, we also see significantly more leniency at $E = 1/16$ than at $E = 1/8$ ($p < 0.001$).

Considering cooperation frequency in histories with the potential for forgiveness, as described above in the response to Question 5, we also see significantly less forgiveness at $E = 0$ (19 percent) compared to $E = 1/16$ (47 percent, $p < 0.001$) or $E = 1/8$ (32 percent, $p < 0.001$). We find no significant difference

⁴¹ For each pairwise comparison of aggregate descriptive statistics, we report the results of a logistic regression over all decisions in the relevant histories, with an error rate dummy as the independent variable, clustered on both subject and interaction pair.

⁴² For each pairwise comparison of aggregated MLE coefficients, we report the results of a two-sample t -test using bootstrapped standard errors of the aggregated coefficients.

TABLE 5—MAXIMUM LIKELIHOOD ESTIMATES FOR OUR $E = 0$, $E = 1/16$, AND $E = 1/8$ CONDITIONS USING THE LAST FOUR INTERACTIONS OF EACH SESSION

	$E = 0$	$E = 1/16$	$E = 1/8$
ALLC	0.24** (0.10)	0 (0.04)	0.06* (0.03)
TFT	0.14* (0.08)	0.04 (0.04)	0.07** (0.03)
TF2T	0 (0.02)	0.24** (0.10)	0.20*** (0.07)
TF3T	0 (0.04)	0.42*** (0.09)	0.09** (0.04)
2TFT	0.15** (0.07)	0 (0)	0.03 (0.02)
2TF2T	0 (0)	0.08 (0.06)	0.12** (0.05)
Grim	0.15* (0.08)	0.03 (0.02)	0.04* (0.02)
Grim2	0.16* (0.09)	0.09 (0.05)	0.05* (0.03)
Grim3	0 (0.06)	0 (0)	0.11*** (0.04)
ALLD	0.07* (0.04)	0.05 (0.03)	0.23*** (0.04)
D-TFT	0.09** (0.04)	0.05 (0.03)	0 (0)
Gamma	0.35*** (0.03)	0.44*** (0.03)	0.43*** (0.02)

Notes: All specifications use $b/c = 4$. Bootstrapped standard errors (shown in parentheses) used to calculate p -values.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

in forgiveness between $E = 1/16$ than at $E = 1/8$ ($p = 0.298$). Examining the aggregated MLE frequencies, the forgiving strategies ALLC, TFT, TF2T, TF3T, 2TFT, and 2TF2T are less common at $E = 0$ (52 percent) compared to $E = 1/16$ (78 percent, $p = 0.007$) and at $E = 1/8$ (57 percent) compared to $E = 1/16$ ($p = 0.001$). There is no significant difference between forgiving strategies at $E = 0$ and $E = 1/8$ ($p = 0.670$).

In summary, we find that substantial levels of leniency and forgiveness are not unique to the high error rate of $E = 1/8$, but are also present at the lower error rate of $E = 1/16$. When the error rate is zero, leniency is much less frequent. The somewhat greater leniency and forgiveness at $E = 1/16$ compared to $E = 1/8$ is surprising; investigating this issue further is an interesting topic for future study.

To further explore the difference in play between error and no-error games, we reanalyze data from Dal Bó and Frechette (2011) and Dreber et al. (2008) using our strategy set from Table 2. Doing so finds TFT to be the most common cooperative strategy in all but one payoff specification.⁴³ Additionally, the aggregate frequency

⁴³ See Appendix D for MLE results.

TABLE 6—DESCRIPTIVE STATISTICS OF AGGREGATE BEHAVIOR, AS WELL AS AGGREGATED MLE FREQUENCIES

	E = 0	E = 1/16	E = 1/8
<i>Descriptive statistics</i>			
Percent C First Round	83%	87%	76%
Percent C All Rounds	78%	82%	59%
Leniency	42%	77%	66%
Forgiveness	19%	47%	32%
<i>MLE aggregation</i>			
Cooperative strategies	84%	89%	77%
Lenient strategies	40%	82%	63%
Forgiving strategies	53%	78%	57%

Notes: All specifications use $b/c = 4$. The descriptive statistics for leniency and forgiveness are defined in the text. For MLE aggregation, all strategies other than ALLD and D-TFT are cooperative; lenient strategies are TF2T, TF3T, 2TF2T, Grim2, and Grim3; forgiving strategies are TFT, TF2T, TF3T, 2TFT, and 2TF2T.

of strategies with memory at most 1 (namely ALLC, TFT, D-TFT, and ALLD) is 76 percent in the games without noise (including our $E = 0$ control), compared to only 33 percent in our games with noise; this difference is largely driven by lenient strategies, most of which by definition look back more than 1 round, and have an aggregate frequency of 13 percent without noise compared to 57 percent with noise.⁴⁴

The importance of noise for promoting leniency is also reflected in the post-experimental questionnaire. Many subjects reported cooperating following their partner's first defection because they assumed it was due to error.

V. Discussion

To relate play in the experiment to theoretical predictions, we would like to understand the extent to which the observed distribution of play approximates an equilibrium, and to the extent that play is not an equilibrium, what sorts of alternative strategies would perform better. To that end, we used simulations to compute the expected payoff matrix for the strategies that had nonnegligible shares in the MLE estimation, along with a few "exploitative" strategies that struck us as good responses to the commonly used lenient strategies.⁴⁵ The resulting payoff matrices are displayed in Appendix E. We will use this table to compute the expected payoff to each strategy given the estimated frequencies, but first we use it to make some observations about the equilibria of the game. In particular, any strategy that is not a Nash equilibrium in this payoff matrix cannot be a Nash equilibrium in the full

⁴⁴Note that here we do find some evidence of longer-memory strategies in the no-error games, while our test based on adding only partner's play two periods ago found an insignificant effect. This may be in part due to the fact that the MLE includes strategies like "Grim" that have a longer memory, and in part to the fact that the no-noise case provides less information about play at many histories. The results are qualitatively equivalent when restricting our attention to the no-noise payoff specifications where TFT risk-dominates ALLD: 62 percent of strategies use memory of at most 1, and lenient strategies have a weight of 20 percent.

⁴⁵Analytic computations of the payoff for two different strategies playing each other is complicated due to the combination of discounting and noise, especially if the strategies look back more than one round and/or have many implicit "states."

game. The converse is of course false, but we will then check which of the strategies that are equilibria of the payoff matrix are also equilibria of the full game.

Using that calculated payoff matrix, we see that the lenient-and-forgiving strategy TF2T, which was common when $b/c = 2.5$ or 4 , is not an equilibrium in any treatment: it can be invaded by DC-Alt (the strategy that alternates between D and C) in all payoff specifications, as well as by ALLD when $b/c = 1.5$ and by various “exploitative” strategies that start with D in the other treatments.⁴⁶ Note also that TFT is never an equilibrium, although it is fairly common when $b/c = 1.5$: it is invaded by ALLD when $b/c = 1.5$ and by ALLC (!) at other b/c values. This is a reflection of the fact that errors can move TFT into an inefficient 2-cycle.

Of course, ALLD is an equilibrium in every treatment, and as discussed in Section II, PTFT is an equilibrium at $b/c = 4$. Perhaps surprisingly, it turns out that Grim2 is also an equilibrium when $b/c = 4$, even though it is not an equilibrium in the game without errors.⁴⁷ Moreover, the range of error probabilities for which Grim2 is an equilibrium increases with b/c . Intuitively, there is more reason to be lenient when the rewards to cooperation are greater, which is consistent with the way overall leniency in the data increases with b/c .

Of course, these equilibrium calculations do not tell us what strategies can persist in a mixed-strategy equilibrium, and they do not tell us which strategies have good payoffs given the actual distribution of play. Table 7 shows the expected payoff of each strategy given the prevailing strategy frequencies. In the $b/c = 1.5$ treatment, where ALLD is most prevalent, ALLD is also the best response to the prevailing strategy frequencies. Furthermore, the average earnings per round of subjects for whom ALLD is the strategy with the greatest likelihood are significantly higher than other subjects’ earnings (coeff = 0.108, $p = 0.028$). ALLD does about as well as the average subject in the $b/c = 2$ treatment (coeff = -0.039 , $p = 0.672$), and is significantly worse at $b/c = 2.5$ (coeff = -0.400 , $p < 0.001$) and $b/c = 4$ (coeff = -0.914 , $p < 0.001$).⁴⁸ We also see that subjects showed good judgment in avoiding PTFT, which performs very poorly in all treatments.

In the treatments with cooperative equilibria, lenient strategies perform very well. Within each of these treatments, the highest payoff strategy that is played is lenient ($b/c = 2$, Grim2; $b/c = 2.5$, 2TF2T; $b/c = 4$, TF2T). Furthermore, all common lenient strategies (frequency of 10 percent or higher) earn within 1 percent of the highest payoff earned by any strategy played in that treatment, except for 2TF2T at $b/c = 2$, which earns 1.6 percent less than the highest payoff. Various start-with-D strategies would have been the highest earners ($b/c = 2$, D-Grim3; $b/c = 2.5$ and

⁴⁶ This is because the one round of punishment provided by TF2T, multiplied by the increased probability of punishment associated with the first D, is too small to outweigh the short-run gain to deviation. When b/c becomes sufficiently large, it does pay to conform to TF2T at histories where the strategy says to cooperate, but then it is also optimal to play C at histories where TF2T says to play D. Mathematica computations show that TF2T is not an equilibrium in any of our treatments.

⁴⁷ Without errors it would be better to play D in the first round and subsequently play Grim2, as the continuation payoff after one D is the same as after no Ds at all. In the presence of errors, however, the expected continuation payoff to Grim2 is lower in the round following a D, and numerical calculations show that when $b/c = 4$ and $\delta = 7/8$, Grim2 is an equilibrium provided that the error probability is between 0.0332 and 0.2778. Moreover, the exploitative but lenient strategy D-Grim2 is also an equilibrium when $b/c = 4$ and $\delta = 7/8$, although it is not used. See the online Appendix for the Grim2 and D-Grim2 equilibrium calculations.

⁴⁸ We report the results of a linear regression over profit in all rounds of all interactions, with an ALLD dummy as the independent variable, clustered on both subject and interaction pair.

TABLE 7—OBSERVED FREQUENCIES AND RESULTING EXPECTED PAYOFFS FOR EACH STRATEGY

	b/c = 1.5		b/c = 2		b/c = 2.5		b/c = 4	
	Frequency	Expected payoff						
ALLC		-1.25	0.03	6.92		13.27	0.06	28.13
TFT	0.19	2.40	0.06	8.71	0.09	14.64	0.07	29.01
TF2T	0.05	1.53		8.69	0.17	14.65	0.20	29.67
TF3T	0.01	0.90	0.03	8.44	0.05	14.53	0.09	29.56
2TFT	0.06	2.87	0.07	8.59	0.02	13.58	0.03	27.08
2TF2T		1.86	0.11	8.89	0.11	14.72	0.12	29.62
GRIM	0.14	3.02	0.07	8.40	0.11	12.33	0.04	23.99
GRIM2	0.06	2.37	0.18	9.03	0.02	13.98	0.05	27.90
GRIM3	0.06	1.79	0.28	9.02	0.24	14.67	0.11	29.23
ALLD	0.29	3.73	0.17	8.53	0.14	11.33	0.23	21.04
D-TFT	0.15	2.89		9.19	0.05	14.66		28.76
PTFT		0.72		6.34		12.05		25.36
D-TF2T		1.93		9.14		<u>14.87</u>		<u>29.73</u>
D-Grim3		2.34		<u>9.54</u>		14.83		28.92

Note: Highest payoff strategy among those that were used is shown in bold; highest payoff strategy among all strategies considered is underlined.

b/c = 4, D-TF2T), but these strategies were not played. Perhaps this is because these exploitative strategies do not fit well with subjects' intuitions and heuristics about cooperative play, and only outperformed the cooperative strategies by a very small margin (the best cooperative strategy earns 5.3 percent less than the best exploitative strategy at b/c = 2, 0.9 percent less at b/c = 2.5 and 0.2 percent less at b/c = 4). Given the incomplete learning we observe in all treatments, it may therefore not be such a surprise that subjects did not discover the benefit of these exploitative strategies. Furthermore, given the roughly equal payoffs, subjects might reasonably prefer lenient cooperative strategies to those that exploit. Exploring the lack of exploitative strategies is an important direct for future work.

Based on the expected payoffs in Table 7, perhaps the largest surprise is not the success of leniency and forgiveness, but rather the high proportion of subjects playing ALLD, particularly at b/c = 4. The reason that low-performing strategies such as ALLD can persist despite receiving low expected payoffs is probably that the complexity of the environment makes it difficult to learn the optimal response. Even though ALLD is not a best response to what people are really doing, ALLD *is* a best response to a belief that everyone else plays ALLD or any other history-independent strategy, and because of the noisy observation of intended play, subjects who have such false beliefs may not learn that more cooperative strategies yield a higher payoff. Consistent with this, 12 percent of subjects defected in more than 85 percent of all rounds in all interactions at b/c = 2, 9 percent of subjects at b/c = 2.5, and 16 percent of subjects at b/c = 4. This accounts for a substantial fraction of the players classified as ALLD by the MLE, and suggests that these subjects almost never experimented with cooperation, preventing them from learning about its benefits. Furthermore, examining the play of these stubborn defectors, we find no positive correlation between first round cooperation and the previous partner's cooperation in the first round of the previous interaction; in fact, the relationship is negative, although not significant

(coeff = -0.081 , $p = 0.882$).⁴⁹ Thus, meeting a first-round cooperator does not increase these subjects' probability of cooperating in future interactions. This is reminiscent of heterogeneous self-confirming equilibrium (Fudenberg and Levine 1993), and the diversity of strategies is consistent with heterogeneous self-confirming equilibrium in the absence of noise; in the presence of noise, similar situations can persist for a while.⁵⁰ A similar logic applies to Grim, which is a best response to the belief that a substantial fraction of the population plays Grim while the rest plays ALLD—a subject who always uses Grim may not learn about the benefits of being more lenient.

We find no difference in first-round cooperation between $b/c = 2$ and $b/c = 2.5$, and that the increase in overall cooperation as b/c increases from 1.5 to 2 is larger than the increase in moving from $b/c = 2$ to $b/c = 2.5$, even though ALLD risk-dominates TFT at $b/c = 2$ but not $b/c = 2.5$. Thus, the risk-dominance criterion has at best limited predictive power regarding cooperation in games with noise.

To explore possible nonstrategic motivations for leniency and forgiveness, we examined the subjects' social preferences using a post-experimental dictator game and a set of survey questions from social psychology. In Dreber, Fudenberg, and Rand (2010), we show that dictator giving is not correlated with cooperation in histories where there is the possibility of leniency, and not consistently correlated with cooperation in histories with the possibility for forgiveness. Dictator giving is also uncorrelated with both first-round cooperation and overall cooperation in the specifications with cooperative equilibria (where leniency and forgiveness are common). Furthermore, while lenient and forgiving strategies earn high expected monetary payoffs, the Fehr and Schmidt (1999) model of inequity aversion gives little utility to these strategies. We use this, the survey data, and additional analysis to argue that social preferences do not seem to be a key factor in explaining the leniency and forgiveness observed in our experiments.

VI. Conclusion

We conclude that subjects do tend to cooperate in noisy repeated games when there is a cooperative equilibrium, that they cooperate even when there are no cooperative equilibria in memory-1 strategies, and that they cooperate even when TFT is risk-dominated by ALLD. This shows that conclusions based on evolutionary game theory models that incorporate the memory-1 restriction need not apply to play in laboratory experiments, and that subjects can and do use strategies with more complexity. We also see that strategies such as TF2T that involve leniency and forgiveness are both common and rather successful in the sense of obtaining high payoffs given the actual distribution of play, even though it is not an equilibrium for all agents to play TF2T: in an uncertain world, it can be payoff-maximizing to be slow to anger and fast to forgive.

⁴⁹When considering only subjects whose realized actions (as opposed to intended actions) resulted in over 85 percent D, we also find no correlation between the previous partner's first round cooperation and own cooperation in the first round of the present interaction (coeff = -0.350 , $p = 0.828$).

⁵⁰Evolutionary models such as the replicator dynamic, when applied to repeated games by restricting the strategy set, can converge to steady states with multiple strategies present, as in Feldman and Thomas (1987). These polymorphic steady states, however, require that all of the active strategies obtain the same payoff, which is not a good approximation of the situation here.

APPENDIX A—MLE STRATEGY FREQUENCY ESTIMATES USING SIMULATED DATA

TABLE A1—MAXIMUM LIKELIHOOD ESTIMATES FOR SIMULATED HISTORIES

	b/c = 1.5		b/c = 2		b/c = 2.5		b/c = 4	
	Actual	Estimated	Actual	Estimated	Actual	Estimated	Actual	Estimated
ALLC	0	0 (0)	0.03	0.02 (0.02)	0	0 (0)	0.05	0.05** (0.03)
TFT	0.18	0.18*** (0.05)	0.08	0.08*** (0.04)	0.08	0.08*** (0.04)	0.10	0.10*** (0.04)
TF2T	0.03	0.03* (0.02)	0	0 (0)	0.17	0.19*** (0.05)	0.18	0.18*** (0.05)
TF3T	0	0 (0)	0.03	0.02 (0.03)	0.05	0.08** (0.05)	0.08	0.04* (0.03)
2TF2T	0	0 (0)	0.10	0.10*** (0.04)	0.13	0.09** (0.05)	0.12	0.13*** (0.05)
Grim	0.22	0.22*** (0.05)	0.10	0.10*** (0.04)	0.12	0.12*** (0.04)	0.03	0.03* (0.02)
Grim2	0.03	0.03* (0.02)	0.20	0.19*** (0.05)	0.02	0.02 (0.02)	0.07	0.07** (0.03)
Grim3	0.08	0.08*** (0.03)	0.28	0.31*** (0.06)	0.25	0.24*** (0.06)	0.12	0.14*** (0.04)
PTFT	0	0 (0)	0	0 (0)	0	0 (0)	0	0 (0)
2PTFT	0	0 (0)	0	0 (0)	0	0 (0)	0	0 (0)
2TFT	0	0 (0)	0	0 (0)	0	0 (0)	0	0 (0)
T2	0	0 (0)	0	0 (0)	0	0 (0)	0	0 (0)
ALLD	0.28	0.28*** (0.06)	0.17	0.17*** (0.06)	0.13	0.13*** (0.05)	0.25	0.25*** (0.06)
C-to-ALLD	0	0 (0)	0	0 (0)	0	0 (0)	0	0 (0)
D-TFT	0.17	0.17*** (0.05)	0	0 (0)	0.05	0.05* (0.03)	0	0 (0)
D-TF2T	0	0 (0)	0	0 (0)	0	0 (0)	0	0 (0)
D-TF3T	0	0 (0)	0	0 (0)	0	0 (0)	0	0 (0)
D-Grim2	0	0 (0)	0	0 (0)	0	0 (0)	0	0 (0)
D-Grim3	0	0 (0)	0	0 (0)	0	0 (0)	0	0 (0)
DC-Alt	0	0 (0)	0	0 (0)	0	0 (0)	0	0 (0)
Gamma		0.02 (0.01)		0.02 (0.01)		0 (0.01)		0.05*** (0.01)

Notes: For each b/c ratio, the first column shows the actual frequency in the simulated data, and the second column shows the MLE estimate. Bootstrapped standard errors shown in parentheses.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

APPENDIX B—EVIDENCE OF LEARNING

First, we investigate the extent of learning over the session for each payoff specification using $E = 1/8$.⁵¹ We do this by examining cooperation in the first round of each interaction (Figure 1), because this reflects each subject's strategy independent of the play of their current partner. There is no significant relationship between interaction number and first-round cooperation when $b/c = 1.5$ (coeff = 0.006, $p = 0.788$), a significant positive relationship when $b/c = 2$ (coeff = 0.089, $p = 0.001$), and a nonsignificant relationship with a nonetheless rather sizable positive coefficient when $b/c = 2.5$ (coeff = 0.056, $p = 0.100$) and $b/c = 4$ (coeff = 0.034, $p = 0.166$). Examining learning at the individual level, we see a significant positive correlation between first-round cooperation and the previous partner's cooperation in the first round of the previous interaction (coeff = 0.335, $p < 0.001$).⁵² Thus, cooperative partners tend to make one more cooperative, although the effect size is moderate (first-round cooperation after meeting a defector, 65 percent; after meeting a cooperator, 72 percent). In the $E = 1/16$ control, we find a similar relationship between first-round cooperation and the previous partner's first-round decision (coeff = 0.513, $p = 0.047$; after defection, 81 percent; after cooperation, 88 percent), although there is no change in first-round cooperation across interactions (coeff = 0.005, $p = 0.898$). For $E = 0$, the opposite is true: there is no significant relationship between first-round cooperation and the previous partner's opening move (coeff = -0.073 , $p = 0.818$), but there is a significant increase in cooperation across interactions (coeff = 0.129, $p = 0.034$). Thus, there is also evidence of learning in both controls.

⁵¹We report the results from a logistic regression over all individual first round decisions, with the interaction number as the independent variable. To account for the non-independence of observations from a given subject, and from subjects within a given pairing, we clustered on both subject and interaction pair.

⁵²The positive correlation between first-round cooperation and the previous partner's cooperation in the first round of the previous interaction remains significant (coeff = 0.291, $p = 0.001$) when controlling for interaction number and b/c ratio, and we find no significant interaction between either interaction number (coeff = -0.016 , $p = 0.484$) or b/c ratio (coeff = 0.0003, $p = 0.998$). Furthermore, we continue to observe this positive relationship when restricting our analysis to the last four rounds of each interaction (coeff = 0.342, $p = 0.001$). Thus, the effect of meeting cooperative partners does not appear to vary across interaction or payoff specification.

APPENDIX C—MLE STRATEGY FREQUENCIES USING FULL 20 STRATEGY SET

TABLE A2—MAXIMUM LIKELIHOOD ESTIMATES FOR THE LAST FOUR INTERACTIONS OF EACH SESSION, ALL 20 STRATEGIES.

	b/c = 1.5 E = 1/8	b/c = 2 E = 1/8	b/c = 2.5 E = 1/8	b/c = 4 E = 1/8	b/c = 4 E = 1/16	b/c = 4 E = 0
ALLC	0 (0)	0.03 (0.03)	0 (0.02)	0.05* (0.03)	0 (0.05)	0.24** (0.10)
TFT	0.19*** (0.05)	0.07 (0.04)	0.09** (0.04)	0.07** (0.03)	0.04 (0.03)	0.15** (0.07)
TF2T	0.05* (0.03)	0 (0)	0.16** (0.07)	0.19*** (0.06)	0.22** (0.09)	0 (0.03)
TF3T	0.01 (0.01)	0.03 (0.03)	0.05 (0.04)	0.09** (0.04)	0.40*** (0.11)	0 (0.03)
2TFT	0.06 (0.04)	0.07* (0.04)	0.02 (0.02)	0.03 (0.02)	0 (0)	0.16** (0.07)
2TF2T	0 (0.02)	0.11** (0.05)	0.11* (0.06)	0.12** (0.06)	0.08 (0.07)	0 (0)
Grim	0.14*** (0.05)	0.05 (0.05)	0.11** (0.05)	0.02 (0.02)	0.02 (0.02)	0.12 (0.08)
Grim2	0.05* (0.03)	0.16** (0.07)	0.02 (0.03)	0.05* (0.03)	0.09* (0.05)	0.16** (0.08)
Grim3	0.06* (0.03)	0.27*** (0.08)	0.24*** (0.08)	0.11** (0.04)	0 (0)	0 (0.05)
PTFT	0 (0)	0 (0)	0 (0)	0 (0)	0.02 (0.02)	0 (0)
2PTFT	0 (0)	0.03 (0.03)	0 (0)	0 (0)	0.01 (0.02)	0 (0)
T2	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0.01)
ALLD	0.27*** (0.05)	0.17*** (0.06)	0.14*** (0.05)	0.21*** (0.04)	0.05 (0.03)	0.06* (0.03)
C-to-ALLD	0 (0.01)	0.01 (0.02)	0 (0)	0.01 (0.01)	0.02 (0.02)	0.02 (0.02)
D-TFT	0.10*** (0.04)	0 (0)	0.03 (0.03)	0 (0)	0.02 (0.02)	0 (0.02)
D-TF2T	0 (0)	0 (0)	0.01 (0.02)	0 (0)	0.02 (0.01)	0 (0.03)
D-TF3T	0.01 (0.01)	0 (0)	0 (0)	0 (0)	0 (0)	0.04 (0.03)
D-Grim2	0.05* (0.03)	0.01 (0.01)	0 (0.01)	0.01 (0.01)	0 (0.01)	0 (0)
D-Grim3	0 (0)	0 (0)	0.01 (0.01)	0 (0)	0.03 (0.03)	0.04 (0.03)
DC-Alt	0 (0)	0 (0)	0 (0)	0.01 (0.01)	0 (0)	0 (0)
Gamma	0.46*** (0.02)	0.49*** (0.03)	0.49*** (0.03)	0.43*** (0.02)	0.43*** (0.02)	0.34*** (0.02)

Note: Bootstrapped standard errors in parentheses.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

APPENDIX D—MLE STRATEGY FREQUENCIES FOR PREVIOUS EXPERIMENTS
WITHOUT NOISE

TABLE A3—MAXIMUM LIKELIHOOD ESTIMATES FOR DAL BÓ AND FRECHETTE (2011), THE TWO-OPTION CONTROL
AND THE PRISONER'S DILEMMA GAMES FROM DREBER ET AL. (2008)

Payoff δ	Dal Bo and Frechette 2011						Dreber et al. 2008	
	R = 32 0.5	R = 40 0.5	R = 48 0.5	R = 32 0.75	R = 40 0.75	R = 48 0.75	b/c = 1.5 0.75	b/c = 2 0.75
ALLC	0 (0)	0 (0.01)	0.01 (0.02)	0 (0)	0 (0.02)	0.02 (0.04)	0 (0)	0 (0)
TFT	0.07* (0.04)	0.06 (0.04)	0.24*** (0.06)	0.23*** (0.07)	0.21* (0.12)	0.55*** (0.15)	0.15** (0.06)	0.40*** (0.15)
TF2T	0 (0)	0.02 (0.02)	0.16** (0.07)	0.11 (0.07)	0.22*** (0.06)	0 (0)	0 (0)	0 (0)
TF3T	0 (0)	0 (0.01)	0.01 (0.02)	0 (0)	0 (0.02)	0.06 (0.04)	0 (0)	0 (0)
2TFT	0 (0)	0.06 (0.04)	0 (0)	0 (0.02)	0.35*** (0.13)	0.09 (0.09)	0 (0)	0 (0)
2TF2T	0 (0)	0 (0)	0 (0.04)	0 (0.04)	0 (0.06)	0 (0)	0 (0)	0 (0)
Grim	0 (0)	0 (0.02)	0 (0)	0 (0.02)	0.04 (0.06)	0.2 (0.12)	0.07 (0.04)	0.21 (0.16)
Grim2	0 (0)	0.01 (0.01)	0.02 (0.02)	0 (0)	0 (0.01)	0.02 (0.04)	0 (0)	0 (0)
Grim3	0 (0)	0.01 (0.01)	0.02 (0.02)	0 (0)	0 (0.01)	0.06* (0.04)	0 (0)	0 (0)
ALLD	0.91*** (0.04)	0.76*** (0.06)	0.49*** (0.07)	0.66*** (0.07)	0.11** (0.05)	0 (0)	0.64*** (0.1)	0.3*** (0.11)
D-TFT	0.02 (0.02)	0.08* (0.04)	0.04 (0.03)	0 (0)	0.08* (0.04)	0 (0)	0.14 (0.09)	0.09 (0.07)
Gamma	0.34*** (0.04)	0.49*** (0.04)	0.4*** (0.04)	0.45*** (0.04)	0.32*** (0.03)	0.28*** (0.03)	0.36*** (0.03)	0.42*** (0.03)

Notes: For $\delta = 1/2$, the last 16 interactions were analyzed, and for $\delta = 3/4$, the last 8 interactions. Bootstrapped standard errors in parentheses. Note that MLE results for our no-noise treatment are shown in the main text, Table 5.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

APPENDIX E—CALCULATED PAYOFF MATRICES

b/c=1.5	ALLC	TFT	TF2T	TF3T	2TFT	2TF2T	GRIM	GRIM2	GRIM3	ALLD	D-TFT	PTFT	D-TF2T	D-Grim3
ALLC	6.99	5.03	6.79	6.98	3.52	6.62	-4.17	3.05	6.05	-10.99	2.78	0.20	4.53	3.46
TFT	8.31	5.09	7.83	8.22	2.89	7.31	2.01	4.16	6.56	-1.81	2.86	4.70	6.33	4.39
TF2T	7.13	5.03	6.90	7.11	3.08	6.74	0.57	4.83	6.16	-3.97	2.80	2.76	4.72	3.44
TF3T	7.01	5.03	6.80	6.99	3.42	6.64	-0.55	4.41	6.38	-5.61	2.78	1.77	4.55	3.93
2TFT	9.32	5.13	7.90	8.88	3.62	6.89	2.70	5.05	6.03	-0.81	2.93	6.97	6.43	3.23
2TF2T	7.24	5.05	6.92	7.19	3.23	6.65	1.08	5.03	6.31	-3.19	2.81	3.97	4.76	3.72
GRIM	14.44	5.15	7.32	8.99	4.13	6.55	3.27	4.96	6.49	-0.66	2.94	8.45	5.72	4.07
GRIM2	9.63	5.10	6.96	7.59	3.73	6.68	2.15	5.69	6.32	-2.23	2.86	6.80	4.79	3.82
GRIM3	7.63	5.06	6.89	7.13	3.65	6.66	1.13	5.26	6.68	-3.65	2.82	5.31	4.65	4.33
ALLD	18.99	5.22	8.45	10.91	3.71	7.29	3.48	5.83	7.97	1.00	2.97	10.75	6.19	5.57
D-TFT	9.81	5.14	9.12	9.69	1.94	8.34	1.31	5.25	7.70	-0.32	2.90	4.89	7.05	5.82
PTFT	11.53	5.10	7.88	9.29	1.86	6.75	-0.18	2.10	4.23	-5.51	2.86	5.85	6.37	1.08
D-TF2T	8.64	5.06	8.26	8.60	1.70	7.98	-0.54	5.52	7.52	-2.46	2.84	2.57	6.07	4.94
D-Grim3	9.35	5.10	8.31	8.65	2.63	7.90	0.62	6.16	8.06	-2.05	2.84	6.29	6.05	5.81

b/c=2	ALLC	TFT	TF2T	TF3T	2TFT	2TF2T	GRIM	GRIM2	GRIM3	ALLD	D-TFT	PTFT	D-TF2T	D-Grim3
ALLC	14.00	11.37	13.71	13.97	9.35	13.50	-0.93	8.73	12.71	-9.99	8.37	4.95	10.71	9.31
TFT	15.31	10.18	14.53	15.17	6.68	13.70	5.28	8.67	12.51	-0.81	6.64	9.57	12.15	9.01
TF2T	14.13	11.19	13.81	14.10	8.09	13.54	3.83	10.52	12.62	-2.97	8.10	7.57	10.87	8.84
TF3T	14.01	11.34	13.71	13.98	9.01	13.50	2.73	10.14	13.07	-4.61	8.32	6.55	10.72	9.75
2TFT	16.31	9.37	13.86	15.58	7.23	12.30	5.96	9.59	11.12	0.19	5.46	11.91	10.98	6.22
2TF2T	14.23	11.00	13.78	14.17	7.92	13.30	4.36	10.69	12.78	-2.20	7.83	9.06	10.82	9.11
GRIM	21.47	9.03	11.95	14.18	7.68	10.90	6.53	8.78	10.83	0.33	5.23	13.46	8.95	6.73
GRIM2	16.62	9.83	13.08	13.91	7.99	12.68	5.43	11.38	12.20	-1.22	6.69	11.71	9.88	8.61
GRIM3	14.63	10.72	13.51	13.96	8.26	13.18	4.39	10.95	13.36	-2.65	7.59	10.18	10.45	10.12
ALLD	25.99	7.62	11.93	15.23	5.61	10.38	5.33	8.43	11.29	2.00	4.61	15.00	8.93	8.10
D-TFT	16.80	9.36	15.70	16.61	4.29	14.48	3.30	9.52	13.44	0.69	5.82	8.98	12.42	10.46
PTFT	18.52	10.03	13.71	15.57	5.75	12.39	3.10	6.11	8.92	-4.50	6.54	11.70	11.23	4.21
D-TF2T	15.64	10.64	15.10	15.57	5.29	14.67	1.43	10.73	13.87	-1.47	7.35	6.67	12.15	10.31
D-Grim3	16.33	9.91	14.66	15.42	5.51	14.12	2.53	11.38	14.62	-1.05	6.90	10.49	11.66	11.62

b/c=2.5	ALLC	TFT	TF2T	TF3T	2TFT	2TF2T	GRIM	GRIM2	GRIM3	ALLD	D-TFT	PTFT	D-TF2T	D-Grim3
ALLC	20.99	17.71	20.63	20.96	15.19	20.36	2.36	14.41	19.41	-8.99	13.98	9.68	16.88	15.10
TFT	22.31	15.27	21.25	22.11	10.48	20.09	8.56	13.22	18.48	0.19	10.42	14.43	17.95	13.70
TF2T	21.14	17.35	20.71	21.08	13.09	20.36	7.10	16.20	19.05	-1.96	13.42	12.41	17.00	14.23
TF3T	21.00	17.64	20.66	20.96	14.61	20.36	6.01	15.79	19.78	-3.61	13.86	11.31	16.89	15.57
2TFT	23.32	13.60	19.84	22.27	10.87	17.73	9.25	14.13	16.20	1.19	7.99	16.84	15.51	9.21
2TF2T	21.23	16.93	20.63	21.15	12.61	19.95	7.60	16.41	19.21	-1.20	12.85	14.15	16.87	14.52
GRIM	28.45	12.94	16.57	19.36	11.25	15.25	9.80	12.60	15.17	1.33	7.50	18.44	12.17	9.41
GRIM2	23.61	14.55	19.21	20.23	12.26	18.69	8.69	17.05	18.10	-0.23	10.50	16.63	14.95	13.35
GRIM3	21.62	16.38	20.09	20.79	12.86	19.72	7.63	16.63	20.04	-1.65	12.35	15.06	16.23	15.98
ALLD	32.98	10.04	15.40	19.53	7.52	13.49	7.15	11.05	14.60	3.00	6.28	19.25	11.66	10.62
D-TFT	23.80	13.59	22.28	23.53	6.61	20.61	5.26	13.82	19.17	1.69	8.70	13.06	17.78	15.11
PTFT	25.52	14.98	19.53	21.86	9.66	18.03	6.35	10.11	13.60	-3.50	10.23	17.55	16.07	7.42
D-TF2T	22.64	16.21	21.94	22.56	8.88	21.38	3.39	15.97	20.22	-0.47	11.88	10.77	18.22	15.71
D-Grim3	23.35	14.72	21.03	22.19	8.37	20.36	4.50	16.61	21.23	-0.05	10.94	14.69	17.28	17.42

b/c=4	ALLC	TFT	TF2T	TF3T	2TFT	2TF2T	GRIM	GRIM2	GRIM3	ALLD	D-TFT	PTFT	D-TF2T	D-Grim3
ALLC	41.99	36.74	41.40	41.91	32.73	40.97	12.21	31.44	39.43	-6.02	30.75	23.91	35.43	32.53
TFT	43.28	30.54	41.39	42.94	21.86	39.32	18.36	26.85	36.29	3.19	21.74	29.00	35.47	27.67
TF2T	42.11	35.82	41.40	42.01	28.08	40.78	16.95	33.25	38.47	1.04	29.42	26.79	35.44	30.37
TF3T	42.01	36.58	41.42	41.93	31.43	40.96	15.83	32.78	39.79	-0.62	30.52	25.66	35.42	32.98
2TFT	44.30	26.31	37.72	42.35	21.69	33.96	19.02	27.69	31.42	4.20	15.60	31.61	29.17	18.21
2TF2T	42.24	34.77	41.16	42.05	26.64	39.89	17.43	33.47	38.60	1.80	27.87	29.39	35.02	30.67
GRIM	49.44	24.61	30.43	34.91	21.91	28.36	19.63	24.14	28.21	4.32	14.39	33.42	21.80	17.40
GRIM2	44.63	28.74	37.55	39.15	25.06	36.79	18.53	34.13	35.78	2.76	21.94	31.42	30.20	27.70
GRIM3	42.63	33.34	39.90	41.25	26.68	39.27	17.49	33.68	40.11	1.35	26.64	29.70	33.63	33.44
ALLD	53.99	17.23	25.86	32.45	13.22	22.76	12.64	18.89	24.55	6.00	11.25	31.99	19.86	18.19
D-TFT	44.80	26.28	42.03	44.28	13.66	38.96	11.14	26.70	36.40	4.69	17.45	25.37	33.88	28.98
PTFT	46.52	29.77	37.02	40.72	21.30	34.95	16.14	22.07	27.55	-0.51	21.34	35.09	30.65	16.83
D-TF2T	43.61	32.94	42.46	43.49	19.68	41.47	9.31	31.65	39.34	2.54	25.40	23.10	36.46	31.81
D-Grim3	44.34	29.11	40.11	42.47	16.99	39.04	10.42	32.34	40.93	2.96	23.03	27.29	34.07	34.84

FIGURE A1. ROW PLAYER'S PAYOFF IS SHOWN, AVERAGED OVER 105 RANDOMLY SIMULATED GAMES USING $c = 2$, $E = 1/8$, AND $\delta = 7/8$

Note: Best response in each column is shown in bold.

REFERENCES

- Abreu, Dilip, David Pearce, and Ennio Stacchetti.** 1990. "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring." *Econometrica* 58 (5): 1041–63.
- Aoyagi, Masaki, and Guillaume Frechette.** 2009. "Collusion as Public Monitoring Becomes Noisy: Experimental Evidence." *Journal of Economic Theory* 144 (3): 1135–65.
- Aumann, Robert J., and Lloyd S. Shapley.** 1994. "Long-term Competition—A Game Theoretic Analysis." In *Essays in Game Theory: In Honor of Michael Maschler*, edited by N. Megiddo, 1–15. New York: Springer-Verlag.
- Axelrod, Robert.** 1984. *The Evolution of Cooperation*. New York: Basic Books.
- Axelrod, Robert, and William D. Hamilton.** 1981. "The Evolution of Cooperation." *Science* 211 (4489): 1390–96.
- Binmore, Kenneth G., and Larry Samuelson.** 1992. "Evolutionary Stability in Repeated Games Played by Finite Automata." *Journal of Economic Theory* 57 (2): 278–305.
- Blonski, Matthias, and Giancarlo Spagnolo.** 2004. "Prisoners' Other Dilemma." Unpublished.
- Blonski, Matthias, Peter Ockenfels, and Giancarlo Spagnolo.** 2011. "Equilibrium Selection in the Repeated Prisoner's Dilemma: Axiomatic Approach and Experimental Evidence." *American Economic Journal: Microeconomics* 3 (3): 164–92.
- Boyd, Robert.** 1989. "Mistakes Allow Evolutionary Stability in the Repeated Prisoner's Dilemma." *Journal of Theoretical Biology* 136 (1): 47–56.
- Boyd, Robert, and Jeffrey P. Lorberbaum.** 1987. "No Pure Strategy is Evolutionarily Stable in the Repeated Prisoner's Dilemma Game." *Nature* 327: 58–59.
- Camera, Gabriele, Marco Casari, and Maria Bigoni.** 2010. "Cooperative Strategies in Groups of Strangers: An Experiment." Purdue University Krannert School of Management Working Paper 1237.
- Chamberlain, Gary.** 1980. "Analysis of Covariance with Qualitative Data." *Review of Economic Studies* 47 (1): 225–38.
- Dal Bó, Pedro.** 2005. "Cooperation under the Shadow of the Future: Experimental Evidence from Infinitely Repeated Games." *American Economic Review* 95 (5): 1591–1604.
- Dal Bó, Pedro, and Guillaume R. Frechette.** 2011. "The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence." *American Economic Review* 101 (1): 411–29.
- Dreber, Anna, Drew Fudenberg, and David G. Rand.** 2010. "Who Cooperates in Repeated Games?" Unpublished.
- Dreber, Anna, David G. Rand, Drew Fudenberg, and Martin A. Nowak.** 2008. "Winners Don't Punish." *Nature* 452: 348–351.
- Duffy, John, and Jack Ochs.** 2009. "Cooperative Behavior and the Frequency of Social Interaction." *Games and Economic Behavior* 66 (2): 785–812.
- Engle-Warnick, Jim, and Robert L. Slonim.** 2006. "Inferring Repeated-Game Strategies from Actions: Evidence from Trust Game Experiments." *Economic Theory* 28 (3): 603–32.
- Fehr, Ernst, and Klaus M. Schmidt.** 1999. "A Theory of Fairness, Competition, and Cooperation." *Quarterly Journal of Economics* 114 (3): 817–68.
- Feinberg, Robert M., and Thomas A. Husted.** 1993. "An Experimental Test of Discount-Rate Effects on Collusive Behaviour in Duopoly Markets." *Journal of Industrial Economics* 41 (2): 153–60.
- Feldman, Marcus W., and Ewart A.C. Thomas.** 1987. "Behavior-Dependent Contexts for Repeated Plays of the Prisoner's Dilemma II: Dynamical Aspects of the Evolution of Cooperation." *Journal of Theoretical Biology* 128 (3): 297–315.
- Fischbacher, Urs.** 2007. "Z-Tree: Zurich Toolbox for Ready-Made Economic Experiments." *Experimental Economics* 10 (2): 171–78.
- Friedman, James W.** 1971. "A Non-cooperative Equilibrium for Supergames." *Review of Economic Studies* 38 (113): 1–12.
- Fudenberg, Drew, and David K. Levine.** 1993. "Self-Confirming Equilibrium." *Econometrica* 61 (3): 523–45.
- Fudenberg, Drew, David I. Levine, and Eric Maskin.** 1994. "The Folk Theorem with Imperfect Public Information." *Econometrica* 62 (5): 997–1039.
- Fudenberg, Drew, and Eric Maskin.** 1986. "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information." *Econometrica* 54 (3): 533–54.
- Fudenberg, Drew, and Eric Maskin.** 1990. "Evolution and Cooperation in Noisy Repeated Games." *American Economic Review* 80 (2): 274–79.
- Fudenberg, Drew, and Eric Maskin.** 1994. "Evolution and Noisy Repeated Games." Unpublished.
- Green, Edward J., and Robert H. Porter.** 1984. "Noncooperative Collusion under Imperfect Price Information." *Econometrica* 52 (1): 87–100.

- Heckman, James J.** 1981. "The Incidental Parameters Problem and the Problem of Initial Condition in Estimating a Discrete Time-Discrete Data Stochastic Process." In *Structural Analysis of Discrete Data with Econometric Applications*, edited by Charles F. Manski and Daniel L. McFadden, 179–95. Cambridge, MA: MIT Press.
- Imhof, Lorens A., Drew Fudenberg, and Martin A. Nowak.** 2007. "Tit-for-Tat or Win-Stay, Lose-Shift?" *Journal of Theoretical Biology* 247 (3): 574–80.
- Murnighan, J. Keith, and Alvin E. Roth.** 1983. "Expecting Continued Play in Prisoner's Dilemma Games: A Test of Several Models." *Journal of Conflict Resolution* 27 (2): 279–300.
- Nowak, Martin A., and Karl Sigmund.** 1990. "The Evolution of Stochastic Strategies in the Prisoner's Dilemma." *Acta Applicandae Mathematicae* 20 (3): 247–65.
- Nowak, Martin A., and Karl Sigmund.** 1993. "A Strategy of Win-Stay, Lose-Shift that Outperforms Tit-for-Tat in the Prisoner's Dilemma Game." *Nature* 364: 56–58.
- Nowak, Martin A., Akira Sasaki, Christine Taylor, and Drew Fudenberg.** 2004. "Emergence of Cooperation and Evolutionary Stability in Finite Populations." *Nature* 428: 646–50.
- Posch, Martin.** 1999. "Win-Stay, Lose-Shift Strategies for Repeated Games—Memory Length, Aspiration Levels and Noise." *Journal of Theoretical Biology* 198 (2): 183–95.
- Roth, Alvin E., and J. Keith Murnighan.** 1978. "Equilibrium Behavior and Repeated Play of the Prisoner's Dilemma." *Journal of Mathematical Psychology* 17 (2): 189–98.
- Wedekind, Claus, and Manfred Milinski.** 1996. "Human Cooperation in the Simultaneous and the Alternating Prisoner's Dilemma: Pavlov versus Generous Tit-for-Tat." *Proceedings of the National Academy of Sciences of the United States of America* 93 (7): 2686–89.

Online Appendix
Slow to Anger and Fast to Forgive:
Cooperation in an Uncertain World

Appendix 0-A: Sample instructions

Instructions:

Thank you for participating in this experiment.

Please read the following instructions carefully. If you have any questions, do not hesitate to ask us. Aside from this, no communication is allowed during the experiment.

This experiment is about decision making. You will be randomly matched with other people in the room. None of you will ever know the identity of the others. Everyone will receive a fixed show-up amount of \$10 for participating in the experiment. In addition, you will be able to earn more money based on the decisions you and others make in the experiment. Everything will be paid to you in cash immediately after the experiment.

You will interact numerous times with different people. Based on the choices made by you and the other participants over the course of these interactions, you will receive between \$0 and \$30, in addition to the \$10 show-up amount.

You begin the session with 50 units in your account. Units are then added and/or subtracted to that amount over the course of the session as described below. At the end of the session, the total number of units in your account will be converted into cash at an exchange rate of 30 units = \$1.

The Session:

The session is divided into a series of interactions between you and other participants in the room.

In each interaction, you play a random number of rounds with another person. In each round you and the person you are interacting with can choose one of two options. Once the interaction ends, you get randomly re-matched with another person in the room to play another interaction.

The setup will now be explained in more detail.

The round

In each round of the experiment, the same two possible options are available to both you and the other person you interact with: A or B.

The payoffs of the options (in units)

Option	You will get	The other person will get
A:	-2	+8
B:	0	0

If your move is A then you will get -2 units, and the other person will get +8 units.

If your move is B then you will get 0 units, and the other person will get 0 units.

Calculation of your income in each round:

Your income in each round is the sum of two components:

- the number of units you get from the move you played
- the number of units you get from the move played by the other person.

Your round-total income for each possible action by you and the other player is thus

		Other person	
		A	B
You	A	+6	-2
	B	+8	0

For example:

If you play A and the other person plays A, you would both get +6 units.

If you play A and the other person plays B, you would get -2 units, and they would get +8 units.

If you play B and the other person plays A, you would get +8 units, and they would get -2 units.

If you play B and the other person plays B, you would both get 0 units.

Your income for each round will be calculated and presented to you on your computer screen.

The total number of units you have at the end of the session will determine how much money you earn, at an exchange rate of 30 units = \$1.

Each round you must enter your choice within 30 seconds, or a random choice will be made.

A chance that the your choice is changed

There is a $7/8$ probability that the move you choose actually occurs. But with probability $1/8$, your move is changed to the opposite of what you picked. That is:

When you choose A, there is a $7/8$ chance that you will actually play A, and $1/8$ chance that instead you play B. The same is true for the other player.

When you choose B, there is a $7/8$ chance that you will actually play B, and $1/8$ chance that instead you play A. The same is true for the other player.

Both players are informed of the moves which actually occur. Neither player is informed of the move chosen by the other. Thus with $1/8$ probability, an error in execution occurs, and you never know whether the other person's action was what they chose, or an error.

For example, if you choose A and the other player chooses B then:

- With probability $(7/8)*(7/8)=0.766$, no changes occur. You will both be told that your move is A and the other person's move is B. You will get -2 units, and the other player will get +8 units.
- With probability $(7/8)*(1/8)=0.109$, the other person's move is changed. You will both be told that your move is A and the other person's move is A. You both will get +6 units.
- With probability $(1/8)*(7/8)=0.109$, your move is changed. You will both be told that your move is B and the other person's move is B. You will both get +0 units.
- With probability $(1/8)*(1/8)=0.016$, both your move and the other person's moves are changed. You will both be told that your move is B and the other person's move is A. You will get +8 units and the other person will get -2 units.

Random number of rounds in each interaction

After each round, there is a $7/8$ probability of another round, and $1/8$ probability that the interaction will end. Successive rounds will occur with probability $7/8$ each time, until the interaction ends (with probability $1/8$ after each round). Once the interaction ends, you will be randomly re-matched with a different person in the room for another interaction. Each interaction has the same setup. You will play a number of such interactions with different people.

You will not be paired twice with the same person during the session, or with a person that was previously paired with someone that was paired with you, or with someone that was paired with someone that was paired with someone that was paired with you, and so on. Thus, the pairing is done in such a way that the decisions you make in one interaction cannot affect the decisions of the people you will be paired with later in the session.

Summary

To summarize, every interaction you have with another person in the experiment includes a random number of rounds. After every round, there is a $7/8$ probability of another round. There will be a number of such interactions, and your behavior has no effect on the number of rounds or the number of interactions.

There is a $1/8$ probability that the option you choose will not happen and the opposite option occurs instead, and the same is true for the person you interact with. You will be told which moves actually occur, but you will not know what move the other person actually chose.

At the beginning of the session, you have 50 units in your account. At the end of the session, you will receive \$1 for every 30 units in your account.

You will now take a very short quiz to make sure you understand the setup.

The session will then begin with one practice round. This round will not count towards your final payoff.

Appendix O-B – Demographic statistics by session

b/c	error	Female	Economics	Age
1.5	1/8	52%	21%	20.9
1.5	1/8	48%	18%	21.0
1.5	1/8	50%	10%	19.6
<i>Average</i>		<i>50%</i>	<i>17%</i>	<i>20.5</i>
2	1/8	45%	14%	20.5
2	1/8	43%	10%	20.4
<i>Average</i>		<i>44%</i>	<i>12%</i>	<i>20.5</i>
2.5	1/8	38%	13%	21.3
2.5	1/8	64%	9%	19.5
2.5	1/8	52%	17%	20.7
<i>Average</i>		<i>52%</i>	<i>13%</i>	<i>20.4</i>
4	1/8	62%	17%	22.8
4	1/8	69%	6%	22.1
4	1/8	61%	13%	21.8
4	1/8	42%	30%	20.3
<i>Average</i>		<i>59%</i>	<i>16%</i>	<i>21.7</i>
4	1/16	35%	17%	22.7
4	1/16	44%	19%	21.0
4	1/16	69%	21%	21.2
<i>Average</i>		<i>47%</i>	<i>19%</i>	<i>21.6</i>
4	0	41%	6%	24.9
4	0	44%	6%	22.8
4	0	36%	8%	21.0
<i>Average</i>		<i>41%</i>	<i>7%</i>	<i>23.2</i>

Appendix O-C – Sequence of game lengths

b/c	e	t0 [†]	t1	t2	t3	t4	t5	t6	t7	t8	t9	t10	t11	t12	t13	t14	t15
1.5	1/8	3	8	7	10	7	8	9	5	11	9	8	7	8	16		
1.5	1/8	3	8	7	10	7	8	9	5	11	9	8					
1.5	1/8	3	8	7	10	7	8	9	5	11	9	8					
2	1/8	3	8	7	10	7	8	9	5	11	9						
2	1/8	3	8	7	10	7	8	9	5	11	9	8	7	8	16	4	
2.5	1/8	3	8	7	10	7	9	10	5	11							
2.5	1/8	3	8	7	10	7	9	10	5	11	9	8	7				
2.5	1/8	3	8	7	10	7	8	9	5	11	9	8	7	8	11		
4	1/8	3	8	7	10	1	8	9	5	11	9	8	7	8			
4	1/8	3	8	7	10	8	9	6	11	7							
4	1/8	3	8	7	10	7	8	9	5	11	9	8	7	8	16	4	9
4	1/8	3	8	7	10	7	8	9	5	11	9	8					
4	1/16	3	8	7	10	7	8	9	5	11	9						
4	1/16	3	8	7	10	7	8	9	5	11	9	8	7	8			
4	1/16	3	8	7	10	7	8	9	5								
4	0	3	8	7	10	7	8	9	5	11	9						
4	0	3	8	7	10	7	8	9									
4	0	3	8	7	10	7	8	9	5	11	9						

[†] t0 was a practice round that did not count toward the players' earnings.

The sequence of games in each session is shown in the preceding table. The starting place in the sequence of random game lengths that was used in the experiment was picked by the programmer, and the sequence following the chosen starting place had an unusually low number of short games. Although the average probability over all rounds of the game continuing was not significantly different from 7/8 (t-test: $p > 0.10$ for all sessions), the overall distribution of game lengths differed significantly from what would be expected using a geometric distribution (Chi² goodness of fit test, $p < 0.05$ in all sessions).

This raises the concern that over the course of the session, subjects might have come to believe that the game was more likely to end in later rounds than in earlier ones, and adjusted their play accordingly. Particularly, we might expect that the tendency for cooperation to decrease over the course of an interaction would be greater in later interactions; or put differently, that the relationship between round number and cooperation will become more negative as interaction number increases.

It is not clear why such an effect would alter our findings, but nonetheless we check for evidence of this occurring. To do so, we ran a logistic regression with robust standard errors clustered on subject and session, including controls for b/c ratio and error rate. In addition to round number and interaction number as independent variables, we also add a [round X interaction] term. A significant negative coefficient on the [round X interaction] term would indicate that in later interactions, subjects are less likely to cooperate in later rounds, suggesting that after subjects have had time to learn that games are more likely to end in later rounds, they become more likely to defect in those later rounds. However, we find no evidence of a such relationship between round number and interaction number, (the coefficient for the [round X interaction] term in the regression is not significantly different from 0), either when considering all histories (coeff=-0.003, p=0.330), only considering histories where in the previous round both players cooperated (coeff=-0.004, p=0.418), only considering histories with the possibility of leniency (coeff=-0.004, p=0.698) or only considering histories with the possibility of forgiveness (coeff=-0.001, p=0.832). This suggests that increasing experience with the game length distribution did not affect subjects' probability to cooperate in later rounds, and in particular did not affect their levels of leniency or forgiveness.

A second complication with the game lengths is that due to technical difficulties with the computer software, the actual sequence of game lengths deviated somewhat from the pre-generated sequence in 4 out of the 18 sessions. We did not inform subjects of this error (which we were unaware of at the time) and they were most likely not aware of it either; no subject commented on the issue, and our experience is that subjects in our experiments do communicate with us when they are aware of software errors.

Appendix O-D – Equilibrium calculations

O-D.1 Equilibrium calculation for TFT with error

If both players use TFT then all histories fall into one of 4 phases defined by play in the previous round, while what happened 2 rounds ago doesn't matter either for current play or for continuation:

- P1: CC yesterday. Here the strategy says to play C.
- P2: CD yesterday. Here the strategy says to play D.
- P3: DC yesterday. Here the strategy says to play C.
- P4: DD yesterday. Here the strategy says to play D

Regardless of the current phase, next round's phase is P1 if both play C; P2 if player 1 plays C while player 2 plays D; P3 if player 1 plays D while player 2 plays C; and P4 if both play D.

Let the payoffs from following TFT in these phases be v_1, v_2, v_3, v_4 respectively, and let the error rate and discount factor be e and δ respectively. Then

$$v_1 = (1 - e)^2(b - c + \delta v_1) + e(1 - e)(b - c + \delta(v_2 + v_3)) + e^2 \delta v_4$$

$$v_2 = e(1 - e)(b - c + \delta(v_1 + v_4)) + (1 - e)^2(b + \delta v_3) + e^2(-c + \delta v_2)$$

$$v_3 = e(1 - e)(b - c + \delta(v_1 + v_4)) + (1 - e)^2(-c + \delta v_2) + e^2(b + \delta v_3)$$

$$v_4 = e^2(b - c + \delta v_1) + e(1 - e)(b - c + \delta(v_2 + v_3)) + (1 - e)^2 \delta v_4.$$

For example, consider v_1 . Here both players intend to play C. Thus with probability $(1-e)^2$, no errors occur, both players play C and remain in phase 1 for the next round; and therefore player 1 earns $(b-c)$ now plus the value of v_1 discounted by δ . With probability $e(1-e)$, player 1 (only) makes an error and accidentally plays D, shifting to phase 3 for the next round; therefore player 1 earns b today plus the value of v_3 discounted by δ . Also with probability $e(1-e)$, player 2 (only) makes an error, exploiting player 1 and shifting to phase 2 for the next round; therefore player 1 earns $-c$ today plus the value of v_2 . Together this results in the 2nd term $e(1-e)(b-c+\delta(v_2+v_3))$. Finally, with probability e^2 both players error and play D, shifting to phase 4 in the next round; here both earn 0 today, plus the value of v_4 discounted by δ .

It is enough to consider deviations in P1 and P2 to show that TFT is never an equilibrium in the presence of noise; for TFT to be an equilibrium, it is necessary (but not sufficient) to have

$$v_1 \geq e(1-e)(b-c+\delta v_1) + (1-e)^2(b+\delta v_3) + e^2(-c+\delta v_2) + e(1-e)\delta v_4$$

and

$$v_2 \geq (1-e)^2(b-c+\delta v_1) + e(1-e)(b+\delta v_3) + e(1-e)(-c+\delta v_2) + e^2\delta v_4.$$

However, these two conditions are mutually exclusive. Thus TFT is never an equilibrium.

53

O-D.2 Equilibrium calculation for PTFT with error

If both players use PTFT then all histories fall into one of 2 phases:

- P1: CC or DD yesterday. Here the strategy says to play C, and what happened 2 days ago doesn't matter either for current play or for continuation. Next round's phase is P1 if both play C or both play D, else P2.
- P2: CD or DC yesterday. Here the strategy says to play D, and what happened 2 days ago again doesn't matter either for current play or for continuation. Next round's phase is P1 if both play C or both play D, else P2.

Let the payoffs from following PTFT in these phases be v_1, v_2

respectively, and let the error rate and discount factor be e and δ respectively. Then

$$v_1 = (1-e)^2(b-c) + e(1-e)(b-c) + \delta((1-e)^2 + e^2)v_1 + 2e(1-e)v_2$$

⁵³ Moreover, a forgetful version of TFT that forgets the current state and picks a new state randomly with some non-zero probability is also never an equilibrium. For TFT, forgetting is equivalent to increasing the error rate e , as follows. From the point of view of an individual TFT player, there are two states: s1 (the other played C last round) and s2 (the other played D last round). In state s1, TFT plays C with probability $1-e$ and play D with probability e ; in state s2, TFT plays C with probability e and play D with probability $1-e$. Now imagine that players have a probability $2p$ of forgetting the state and randomly picking a new state, such that with probability p a player switches state. When the state is forgotten in s1 (with probability p), the player gets 'confused' and switches to state s2, and therefore intends to defect. Thus a TFT player in state s1 plays C with probability $(1-p)(1-e)+ep$ and plays D with probability $p(1-e)+(1-p)e$. This is equivalent to non-forgetful TFT with error rate $e'=e+p-2ep$. The same is true when considering a player that forgets in state s2, who plays C with probability $p(1-e)+(1-p)e$ and plays D with probability $(1-p)(1-e)+ep$. As shown above, TFT is never an equilibrium regardless of the value of e . Therefore forgetful TFT is never an equilibrium.

$$v_2 = (1 - e)^2 \delta v_1 + e^2(b - c + \delta v_1) + e(1 - e)((b - c) + 2\delta v_2).$$

Following PTFT is clearly optimal in P2, so it is enough to check for a one-stage deviation in P1. Thus PTFT is an equilibrium if

$$v_1 \geq (1 - e)^2(b + \delta v_2) + e^2(-c + \delta v_2) + e(1 - e)(b - c + 2\delta v_1)$$


For $c=2$ and $\delta=7/8$, we evaluate this expression for relevant values of b/c and e .

e	b/c	v_1	v_2	Phase 1 deviation payoff	Is PTFT an equilibrium?
1/8	1.5	5.85	5.1	6.98	No
1/8	2	11.7	10.2	12.46	No
1/8	2.5	17.55	15.3	17.94	No
1/8	4	35.11	30.61	34.39	Yes
1/16	4	40.69	35.44	38.93	Yes
0	4	48	42	44.75	Yes


To more fully explore that range of e values over which PTFT is an equilibrium, we plot the payoff advantage of deviating in P1 as a function of e , for $\delta = 7/8$, $c = 2$ and various b . We see that the for payoff specifications where PTFT is not an equilibrium at $e = 1/8$, increasing e does not lead to PTFT becoming an equilibrium.⁵⁴

⁵⁴ A forgetful version of PTFT that forgets the current state and picks a new state randomly with some non-zero probability is also not an equilibrium at $b/c=1.5$, 2 or 2.5 , as forgetting for PTFT is equivalent to increasing the error rate e , for similar reasons as for TFT. As shown above, increasing the error rate e cannot make PTFT an equilibrium for $b/c=1.5$, $b/c=2$ or $b/c=2.5$ with $\delta=7/8$, and therefore forgetful PTFT is not an equilibrium for those value.


Advantage of Phase 1 Deviation, b.c. 1.5




Advantage of Phase 1 Deviation, b.c. 2



Advantage of Phase 1 Deviation, b.c. 2.5



Advantage of Phase 1 Deviation, b.c. 4



O-D.3 Payoffs of 2x2 game between TFT and ALLD and between Grim and ALLD with error probability $e=1/8$ and $\delta=7/8$

TFT vs. ALLD

Grim vs. ALLD

b/c = 1.5

	TFT	ALLD
TFT	5.09	-1.81
ALLD	5.23	1

b/c = 1.5

	Grim	ALLD
Grim	3.27	-0.66
ALLD	3.49	1

b/c = 2

	TFT	ALLD
TFT	10.18	-0.81
ALLD	7.625	2

b/c = 2

	Grim	ALLD
Grim	6.54	0.34
ALLD	5.32	2

b/c = 2.5

	TFT	ALLD
TFT	15.27	0.19
ALLD	10.03	3

b/c = 2.5

	Grim	ALLD
Grim	9.82	1.34
ALLD	7.15	3

b/c = 4

	TFT	ALLD
TFT	30.55	3.19
ALLD	17.25	6

b/c = 4

	Grim	ALLD
Grim	19.63	4.34
ALLD	12.64	6

O-D.4 Payoffs of 2x2 game between TFT and ALLD and between Grim and ALLD with error probability $e=1/16$ and $\delta=7/8$

TFT vs. ALLD

b/c = 1.5

	TFT	ALLD
TFT	5.87	-2.02
ALLD	4.27	0.50

b/c = 2

	TFT	ALLD
TFT	11.73	-1.52
ALLD	6.03	1.00

b/c = 2.5

	TFT	ALLD
TFT	17.60	-1.02
ALLD	7.79	1.50

b/c = 4

	TFT	ALLD
TFT	35.20	0.48
ALLD	13.06	3.00

Grim vs. ALLD

b/c = 1.5

	Grim	ALLD
Grim	4.29	-1.34
ALLD	3.27	0.50

b/c = 2

	Grim	ALLD
Grim	8.58	-0.84
ALLD	4.69	1.00

b/c = 2.5

	Grim	ALLD
Grim	12.87	-0.34
ALLD	6.11	1.50

b/c = 4

	Grim	ALLD
Grim	25.73	1.16
ALLD	10.38	3.00

O-D.5 Equilibrium calculation for Grim with error

If both players use Grim then all histories fall into one of 2 phases:

- P1: No D yesterday. Here the strategy says to play C, and what happened 2 days ago doesn't matter either for current play or for continuation. Next round's phase is P1 if both play C else P2
- P2: at least one D in the last round: play D. This phase is absorbing.

Let the payoffs from following Grim in these phases be v_1, v_2 respectively, and let the error rate and discount factor be e and δ respectively.

$$\text{Then } v_1 = (1 - e)^2(b - c) + e(1 - e)(b - c) + \delta[(1 - e)^2v_1 + (2e - e^2)v_2],$$

$$v_2 = e^2(b - c) + e(1 - e)(b - c) + \delta v_2 \rightarrow v_2 = (e^2(b - c) + e(1 - e)(b - c)) / (1 - \delta)$$

Following Grim is clearly optimal in P2, so it is enough to check for a one-stage deviation in P1. Thus Grim is an equilibrium if

$$v_1 \geq (1 - e)^2(b + \delta v_2) + e(1 - e)(b - c + \delta v_1) + e(1 - e)(\delta v_2) + e^2(-c + \delta v_2)$$

For $c=2$ and $\delta=7/8$, we evaluate this expression for relevant values of b/c and e .

e	b/c	v_1	v_2	Phase 1 deviation payoff	Is Grim an equilibrium?
1/8	1.5	3.27	1	3.47	No
1/8	2	6.54	2	5.43	Yes
1/8	2.5	9.82	3	7.40	Yes
1/8	4	19.63	6	13.30	Yes
1/16	4	25.73	3	11.17	Yes
0	4	48	0	8	Yes

O-D.6 Equilibrium calculation for Grim2 with error

If both players use Grim2 then all histories fall into one of 3 phases:

- P1: No D yesterday. Here the strategy says to play C, and what happened 2 days ago doesn't matter either for current play or for continuation. Next round's phase is P1 if both play C else P2
- P2: D yesterday but none the day before: play C. Next round's phase is P1 if both play C else P3.
- P3: at least one D in each of the last two rounds: play D. This phase is absorbing.

Let the payoffs from following Grim2 in these phases be v_1, v_2, v_3 respectively, and let the error rate and discount factor be e and δ respectively.

$$\text{Then } v_1 = (1 - e)^2(b - c) + e(1 - e)(b - c) + \delta[(1 - e)^2v_1 + (2e - e^2)v_2],$$

$$v_2 = (1 - e)^2 b - c + \delta v_1 + e(1 - e)(b - c) + (2e - e^2)\delta v_3]$$

$$v_3 = e^2(b - c) + e(1 - e)(b - c) + \delta v_3 \rightarrow v_3 = (e^2(b - c) + e(1 - e)(b - c))/(1 - \delta)$$

Following Grim2 is clearly optimal in P3, so it is enough to check for one-stage deviations in P1 and P2. Thus Grim2 is an equilibrium if

$$v_1 \geq (1 - e)^2(b + \delta v_2) + e(1 - e)(b - c + \delta v_1) + e(1 - e)(\delta v_2) + e^2(-c + \delta v_2) \text{ and}$$


$$v_2 \geq (1 - e)^2(b + \delta v_3) + e(1 - e)((b - c) + \delta v_1) + e(1 - e)(\delta v_3) + e^2(-c + \delta v_3)$$

For $c=2$ and $\delta=7/8$, we evaluate this expression for relevant values of b/c and e .


e	b/c	v_1	v_2	v_3	Phase 1 deviation payoff	Phase 2 deviation payoff	Is Grim2 an equilibrium?
1/8	1.5	5.69	4.89	1	6.73	3.70	No
1/8	2	11.38	9.78	2	11.96	5.90	No
1/8	2.5	17.07	14.68	3	17.20	8.10	No
1/8	4	34.14	29.35	6	32.89	14.69	Yes
1/16	4	41.85	38.12	3	40.92	11.99	Yes
0	4	48	48	0	50	8	No

To more fully explore that range of e values over which Grim2 is an equilibrium when $b=8$, $c=2$ and $\delta=7/8$, we plot the payoff advantage of deviating in each state as a function of e . Numerical calculation shows that Grim2 is an equilibrium when $0.033 < e < 0.278$.

Advantage of Phase 1 Deviation




Advantage of Phase 2 Deviation







Appendix O-E – Strategy definitions

Here we define each strategy included in our analysis. Each phase is represented by a circle, with the strategy's move in that phase shown in the center of the circle, and transitions out of the phase indicated with arrows.


We begin with the strategies where transitions between phases depend only on the partner's move in the previous round: TFT, TF2T, TF3T, 2TFT, 2TF2T, D-TFT, D-TF2T and D-TF3T. For clarity we indicate only the partner's move with each transition arrow. All strategies begin in the leftmost phase.










Next we define the strategies where transitions depend on both players' actions in the previous round: Grim, Grim2, Grim3, D-Grim2, D-Grim3, PTFT, 2PTFT and T2. The last round histories associated with each transition are indicated by the pair $X_i X_j$ where X_i is the strategy's move last round and X_j is the partner's move last round (i.e. CD represents histories where the strategy played C last round while the partner played D).

When only one transition out of a phase exists, irrespective of either player's action, the transition is labeled "All".





Lastly, we define the strategies whose transitions do not depend on previous histories of play: ALLC, ALLD, C-to-ALLD and DC-Alt.



Appendix O-F Robustness of strategy analysis

In the main text, we analyze the last 4 interactions of each session to minimize the effects of learning. Here we replicate our main analyses considering instead the last 2 or last 6 interactions, and find little difference, as shown in the table below. Regardless of the cutoff, we find that cooperation and leniency increase substantially going from $b/c=1.5$ to $b/c=2$, while forgiveness is changes little between $b/c=1.5$ and $b/c=2$, and then steadily increases from $b/c=2$ to $b/c=4$.

Last 2 Interactions	b/c=1.5	b/c=2	b/c=2.5	b/c=4
Descriptive statistics				
%C First Round	53%	79%	78%	77%
%C All Rounds	33%	45%	61%	64%
Leniency	30%	64%	67%	69%
Forgiveness	18%	16%	27%	45%
MLE aggregation				
Cooperative strategies	59%	84%	83%	78%
Lenient strategies	21%	53%	62%	61%
Forgiving strategies	27%	24%	41%	58%

Last 4 Interactions	b/c=1.5	b/c=2	b/c=2.5	b/c=4
Descriptive statistics				
%C First Round	54%	75%	79%	76%
%C All Rounds	32%	49%	61%	59%
Leniency	29%	63%	67%	66%
Forgiveness	15%	18%	33%	32%
MLE aggregation				
Cooperative strategies	57%	83%	81%	77%
Lenient strategies	18%	62%	60%	63%
Forgiving strategies	31%	31%	44%	57%

Last 6 Interactions	b/c=1.5	b/c=2	b/c=2.5	b/c=4
Descriptive statistics				
%C First Round	54%	75%	78%	75%
%C All Rounds	32%	49%	60%	58%
Leniency	30%	59%	68%	64%
Forgiveness	14%	19%	31%	32%
MLE aggregation				
Cooperative strategies	59%	82%	83%	77%
Lenient strategies	23%	65%	63%	64%
Forgiving strategies	36%	37%	46%	57%

Our MLE estimation procedure assumes that the probability of mental error in strategy implementation, parameterized by γ , is equal across strategies. It is possible,

however, that some strategies are more difficult to implement than others and therefore γ may vary across strategies. Here we replicate our MLE estimates from Table 3, now allowing each strategy to have a different γ . We find little difference.

	b/c=1.5	b/c=2	b/c=2.5	b/c=4
	E=1/8	E=1/8	E=1/8	E=1/8
ALLC	0	0.0195	0.0486	0.062
TFT	0.1894	0.0668	0.0864	0.0601
TF2T	0.0473	0.019	0.2109	0.2137
TF3T	0.0122	0.0186	0.0343	0.0616
2TFT	0.069	0.0558	0.0234	0.0381
2TF2T	0.0048	0.0903	0.0887	0.1195
Grim	0.1551	0.1324	0.1054	0.0297
Grim2	0.0567	0.1122	0.0362	0.0902
Grim3	0.0429	0.3147	0.1685	0.1188
ALLD	0.2362	0.1408	0.1396	0.1669
D-TFT	0.1865	0.0299	0.0579	0.0396
γ -ALLC	0.6517	0.3617	0.0438	0.2101
γ -TFT	0.4615	0.505	0.4034	0.3392
γ -TF2T	0.4329	3.8609	0.6046	0.4275
γ -TF3T	0.2886	0	0.4152	0.296
γ -2TFT	0.7757	0.4042	0.6942	0.5697
γ -2TF2T	0.4689	0.4188	0.419	0.3959
γ -Grim	0.4602	0.6609	0.6798	0.4077
γ -Grim2	0.4536	0.4062	0.6554	1.2917
γ -Grim3	0.4474	0.5662	0.408	0.4535
γ -ALLD	0.3165	0.2935	0.356	0.2086
γ -D-TFT	0.6247	1.4884	0.8278	10
Cooperative strategies	58%	83%	80%	79%
Lenient strategies	16%	57%	59%	67%
Forgiving strategies	32%	27%	49%	56%

We now show that our MLE estimates are robust to including two alternative classes of simple strategies. Firstly, we consider forgetful memory-1 strategies (we noted above that these strategies are not equilibria at $b/c=2$ or $b/c=2.5$, but subjects might be playing them anyway). To include F-TFT and F-PTFT in the MLE, we take advantage of the fact that forgetting is equivalent to experiencing a higher error rate. Thus we add an additional parameter γ_F to the MLE which represents the additional probability of mental error for forgetful players (relative to non-forgetful players). The MLE terms for F-TFT and F-PTFT are therefore

$$p_i(s) = \prod_k \prod_r \left(\frac{1}{1 + \exp(-s_{ikr}(s) / (\gamma + \gamma_F))} \right)^{y_{ikr}} \left(\frac{1}{1 + \exp(s_{ikr}(s) / \gamma + \gamma_F)} \right)^{1-y_{ikr}}$$

where $p_i(s)$ is the likelihood of strategy s given the history of subject i , y_{ikr} is the actual decision of subject i in round r of interaction k ($0=D$, $1=C$), s_{ikr} is the predicted move of strategy s ($-1=D$, $1=C$) and γ parameterizes the mental error rate of non-forgetful strategies. A referee also suggested that subjects might be playing simple strategies which simply ignore their partner's first move before triggering (i.e. always cooperate in the first 2 interactions). To test for this, we also include the strategy C-TFT with always plays C for the first 2 periods then switches to TFT, and C-Grim with always plays C for the first 2 periods then switches to Grim. As in the main text, we estimate the frequency of each strategy using MLE and determine whether a strategy is present at frequency significantly greater than 0 using a t-test with bootstrapped standard errors.

As can be seen in the following table, we find little of any of these strategies. Bootstrapping standard errors shows that none are present at levels significantly greater than 0 ($p > 0.05$ for all 4 strategies in all 4 payoff specifications).

	b/c=1.5	b/c=2	b/c=2.5	b/c=4
ALLC	0	0.02	0.01	0.06
TFT	0.15	0.04	0.08	0.06
TF2T	0.05	0	0.15	0.17
TF3T	0.01	0.01	0.05	0.08 [†]
2TF2T	0	0.11	0.11	0.12
Grim	0.13	0.02	0.07	0.01
Grim2	0.05	0.15	0.02	0.01
Grim3	0.05	0.28	0.23	0.12
PTFT	0	0	0	0
2PTFT	0	0.03	0	0
2TFT	0.03	0.07	0.02	0.03
T2	0	0	0	0
ALLD	0.27	0.17	0.14	0.18
C-to-ALLD	0	0.02	0	0.01
D-TFT	0.09	0	0.02	0
D-TF2T	0	0	0.02	0
D-TF3T	0.01	0	0	0
D-Grim2	0.05	0	0	0
D-Grim3	0	0	0.01	0
DC-Alt	0	0	0	0
C-TFT	0.03	0.03	0	0.02
C-Grim	0.01	0.04	0.03	0.03
F-TFT	0.07	0.03	0.06	0.04
F-PTFT	0	0	0	0.05
<i>Gamma</i>	<i>0.43</i>	<i>0.48</i>	<i>0.46</i>	<i>0.37</i>
<i>Gamma_F</i>	<i>0.6</i>	<i>2.17</i>	<i>1.09</i>	<i>1.92</i>

A referee also suggested the strategy that cooperates until the fraction of D moves by the partner passes some threshold, at which point it switches permanently to defection. To test for this possibility, we re-analyze the data including this family of strategies. We include 9 strategies which stop cooperating once the fraction of Ds by the partner is greater than 10%, 20%, 30%, 40%, 50%, 60%, 70%, 80% or 90%. We find that none of these strategies are present at a frequency significantly greater than 0 in any payoff specification ($p > 0.05$ for all). The MLE results are shown in the following table:

	b/c=1.5	b/c=2	b/c=2.5	b/c=4
ALLC	0	0.012	0	0.04
TFT	0.1648	0.065	0.088	0.064
TF2T	0.0459	0	0.1719	0.194
TF3T	0.0099	0.038	0.0383	0.083
2TFT	0.0033	0.097	0.1038	0.093
2TF2T	0.126	0.066	0.1147	0.028
Grim	0.0497	0.187	0.0093	0.023
Grim2	0.0321	0.245	0.1997	0.065
Grim3	0.054	0.07	0	0.013
ALLD	0.2848	0.173	0.1407	0.227
D-TFT	0.1321	0	0.0478	0
FracD-10%	0	0	0	0.029
FracD-20%	0	0	0	0
FracD-30%	0.0273	0	0	0.016
FracD-40%	0.0224	0	0	0
FracD-50%	0	0.029	0.0859	0
FracD-60%	0.0446	0	0	0.077
FracD-70%	0.0031	0	0	0.048
FracD-80%	0	0	0	0
FracD-90%	0	0.017	0	0
<i>Gamma</i>	<i>0.453</i>	<i>0.497</i>	<i>0.49</i>	<i>0.423</i>

Appendix O-G MLE estimates including stochastically forgiving strategies

Our main analysis considers only deterministic strategies. Here we extend the MLE formulation to include mixed strategies. The original formulation is

$$p_i(s) = \prod_k \prod_r \left(\frac{1}{1 + \exp(-s_{ikr}(s) / \gamma)} \right)^{y_{ikr}} \left(\frac{1}{1 + \exp(s_{ikr}(s) / \gamma)} \right)^{1-y_{ikr}}$$

where $p_i(s)$ is the likelihood of strategy s given the history of subject i , y_{ikr} is the actual decision of subject i in round r of interaction k ($0=D$, $1=C$), s_{ikr} is the predicted move of strategy s ($-1=D$, $1=C$) and γ parameterizes the mental error rate. We replace this with a new formulation

$$p_i(s) = \prod_k \prod_r \left[s_{ikr} \left(\frac{1}{1 + \exp(-1 / \gamma)} \right) + (1 - s_{ikr}) \left(\frac{1}{1 + \exp(1 / \gamma)} \right) \right]^{y_{ikr}} \cdot \left[(1 - s_{ikr}) \left(\frac{1}{1 + \exp(-1 / \gamma)} \right) + s_{ikr} \left(\frac{1}{1 + \exp(1 / \gamma)} \right) \right]^{1-y_{ikr}}$$

where s_{ikr} now represents the probability that strategy s cooperates given the history preceding round r of interaction k (0 to 1).

We use this new formulation to consider stochastic conditional strategies. In particular we explore a family of ‘generous’ TFT (GTFT) strategies which have received significant attention in the evolutionary game theory literature (e.g. Nowak and Sigmund 1990). These strategies are stochastically forgiving. Like TFT, GTFT always responds to C with C. In response to an opponent’s D, however, GTFT stochastically cooperates with probability q .

First we analyze simulated data to see whether the MLE can differentiate between a GTFT that forgives defection 20% of the time and TF3T. We simulate a session with 10 ALLD players, 10 TF3T players, and 10 GTFT-2 players. We simulate 4 interactions, each lasting 8 rounds, and find that the MLE correctly assigns 1/3 to ALLD, TF3T and GTFT-2. Thus the MLE is able to distinguish between memory-1 stochastic forgiveness and longer deterministic forgiveness.

We now reanalyze our data using the 11 strategies from the main text Table 3 plus 9 GTFT variants- those which forgive defection with 10% probability (GTFT-1), 20% probability (GTFT-2), ... 90% probability (GTFT-9). We find a somewhat sizable fraction of people playing stochastically forgiving memory 1 strategies (between 10 and 19% in the treatments with cooperative equilibria). However, the inclusion of these strategies doesn't undermine the importance of lenient strategies with more than 1 period of memory – the longer memory lenient strategies are much more common:

	b/c=1.5	b/c=2	b/c=2.5	b/c=4
GTFT-X	11%	12%	10%	19%
TF2T+TF3T+2TF2T+Grim2+Grim3	16%	57%	51%	43%

Thus we conclude that exploring stochastic strategies is an interesting avenue for future research, but that our main findings related to leniency are robust to including stochastic forgiveness.

Appendix O-H Logistic regression analyzing dependence on play 2 rounds ago

In the main text, we use a logistic regression to provide evidence that subjects' decisions are influenced by the partner's move two rounds ago as well as decisions in the previous round. A potential concern with this methodology lies in the possibility for such correlations to occur in heterogeneous populations of subjects with different strategies each of which conditions on at most play in the previous period. This is because the partner's decision two periods ago can give information about a subject's type even if it does not directly influence her decision. For example, consider a population of $\frac{1}{2}$ ALLD and $\frac{1}{2}$ TFT, and imagine that in the previous round both subject played C. If Player A played C two rounds ago then Player A is much more likely to be a TFT player than an ALLD player. This is because last round Player A played C in response to her partner's C two rounds ago, and this is consistent with TFT but not ALLD. Therefore Player A is also likely to play C now. If Player A's partner played D two rounds ago, however, then Player A is equally likely to be TFT or ALLD, because both strategies intend to play D in response to D. Therefore Player A is equally likely to play C or D in current round. And so if Player A's partner played C two rounds ago, she is more likely to play C now than if her partner played D last round, even though her strategy does not look back two periods.

Including controls for player type can help address this issue. For example, in a population of ALLD and TFT players, first round cooperation can do a good job of cleanly differentiating types. First round cooperation does not differentiate between TFT and GTFT, however, but GTFT players will have higher overall cooperation. Thus we include controls for both first round cooperation and overall cooperation in our regression.

To explore how pervasive of a problem bias stemming from heterogeneity might be for our analysis, and how effective our controls are for mitigating it, we conduct various simulations. We consider 5 different populations:


- i. $\frac{1}{3}$ ALLD, $\frac{1}{3}$ TFT, $\frac{1}{3}$ GTFT-5
- ii. $\frac{1}{3}$ ALLD, $\frac{1}{3}$ TFT, $\frac{1}{3}$ F-TFT (forgets the state with 10% probability)
- iii. $\frac{1}{2}$ ALLD, $\frac{1}{2}$ TFT

- iv. 1/2 75%-TFT 25%-CoinFlip, 1/2 25%-TFT 75%-CoinFlip (on each decision, these strategies randomly choose to either play according to TFT or select a move at random (i.e. CoinFlip) – one half the population plays TFT 75% of the time, the other half picks randomly 75% the time)
- v. 1/2 ALLD, 1/2 TF2T

Agents in populations (i), (ii), (iii) and (iv) do not condition on play two periods ago, while 1/2 of the agents in population (v) do. For each population, we simulate 3 sessions of 30 players each, playing 4 interactions each last 8 periods (for maximum comparability to our data). For each population we conduct 25 simulation replicates, and for each replicate we perform the logistic regression with correlated random effects described in the main text (own decision as a function of partner's move 2 rounds ago, own move 2 rounds ago, other's move last round, own move last round, own frequency of first round cooperation, and own frequency of overall cooperation). Across the 100 simulations using only memory 1 strategies, we find 4 instances in which other's play two rounds ago is significant at the $p < 0.05$ level, consistent with what is expected due to chance. In contrast, we consistently find a highly significant effect of play two periods ago for every replicate of (v), where agents are in fact conditioning on play two rounds ago.

Examining the size of the coefficient for partner's move 2 rounds ago, not just the p-value, gives further insight into the relative importance of earlier play. We find that in all 4 scenarios where agents are not actually conditioning on play 2 rounds ago, the coefficient for partner's move 2 rounds ago is at least an order of magnitude smaller than the coefficients for play 1 round ago (usually several orders of magnitude smaller). In the scenario where many agents are actually conditioning on play 2 rounds ago, conversely, all coefficients are of the same order of magnitude. In our analysis of the data from our experiments, we find the latter result rather than the former: a large coefficient on partner's move 2 rounds ago, on the same order of magnitude as the other coefficients. These results suggest that we are in fact picking up subjects explicitly conditioning on the outcome 2 periods ago, rather than only finding spurious correlations due to heterogeneity and stochastic strategies. Graphically displaying these relationships shows

a substantial and consistent effect of play two rounds ago, across all states in the previous round:



Appendix O-I Self-reported strategy descriptions

The post-experimental questionnaire included a free-response question in which subjects had the option of describing the strategy they used in the game. The responses of those subject who answered are reproduced here.

$b/c=1.5$, $E=1/8$

- I almost always started by choosing B and continued to play B for the subsequent rounds
- I chose B almost all the time because you are guaranteed not to lose any points with B
- I chose A to start with. If the other person also chose A, I switched to B the next round. If they too chose B, I stuck with it. If they chose A, I would make a 50/50 decision between A and B
- Once B was played by either player (including by me if my choice was switched) I played B for the rest of the interaction
- I chose B. I continued to chose B unless I felt the daring to press A. I chose B 98
- I'd start with B and continue if he played B. if he did A, I'd switch
- I started out choosing B, but if someone chose A twice I'd play A (I think). If someone was playing straight Bs, I'd occasionally play an A in an effort to guilt them into playing an A (at which point I'd have switched back to straight Bs)
- sometimes I started with B but rarely
- i started by choosing B. if the other person played A in the next round, i started playing A until the other person picked B in a later round. If, after that round, they chose A again, i continued to play A
- i chose B every time unless for two rounds in a row the other person played A
- i chose B most (if not all my rounds) because i wanted to maximize my points and minimize my losses
- random
- i chose A once then chose B the rest of the time
- i played A until they played B twice. Then i would switch to B. if they dont switch back to B, i would keep playing B or possibly do a one-off and switch to A (only if i have positive (or more than 2) points)
- i played B every chance i had because i know that the conservative strategy can win. However, i think i played A twice over the entire experiment because i was bored of pressing B
- always play A. switch to B if other player plays B more than 2 times in a row
- i basically always chose B
- always choose B. you cannot lose
- if it became clear we were going tit for tat with each other I'd try and break out of the cycle by playing A even if his or her last move were B

- i start by choosing B. if the other person plays A i will then choose A the next round. If the other person plays A twice in a row or every other time i then play A

b/c=2, E=1/8

- I chose B every time because it had the least risk
- start with A. always chose A unless the person goes B twice in a row
- you want to convey trustworthiness so when/if your response gets changed/changes you will remain with both As subsequently. I chose A until the other person played B 2x, unless it came in the middle of a string of As. I would return to A if my partner did
- I started with A. if I got B back, I would usually keep playing B unless they switched to A. if I got A back, I would usually play A, but slip in a few B's. hoping they would appear accidental.
- I played A until I thought the person was a B idiot, which was based on frequency not total amount of B's or A's chosen
- basically, I start with A and assume A from my opponent until otherwise noted. If opponent played B. I would switch to B until /unless my opponent went back to A indicating the B was a random switch by the comp from A.
- random plays: until the other player plays B. Then, play B afterwards.
- usually, I started by choosing B, then switched to A if the person played A twice in a row
- if the other person gave me 4 points and lost 2 points in the previous round I would try to do the same. if the other person consistently chose B I didn't want to chose A because I would only lose points and gain none. I would chose B after a round where I played A and the other played B, because in the previous round I gave up points to the other person so I expect them to do the same for me in this round
- AAABAABABB. always 3 As in the beginning unless the person played all Bs, in which case I did all Bs
- never played A
- along with A, build trust/rapport that would lead to best outcome for both
- wanted to avoid a B vs. B stalemate
- if I felt the round was ending and there was no incentive to further build the relationship, I might pick B. the other person might have chosen A, but his choice switched by chance. So to reciprocate his intentions, I might choose A
- I tried to use the tit for tat strategy to earn the most points in the long run. the first time the other player played B, I gave them the benefit of the doubt in that the computer had chosen for them.
- A makes the most sense. Both have incentive to gain whereas B only one person has incentive or both get nothing.
- lull them into thinking you were playing fairly (how barbaric)
- when I saw a string of "A" I felt I could trust them. String of "B"s, I didn't
- I would most likely play A for a few rounds as long as the other person was doing the same. Every now and then I would throw in a B to maximize my points. I would also play B if the other person had previously played B. I would never

choose B twice in a row because that would most likely lead to both persons choosing B for the remainder of the experiment. I would start each trial by choosing A. I would choose A until I decided to switch to gain extra points, or until the other person chose B. if that happened I would choose B the next round. If I chose B and the other person chose A, then I would always choose A the next round. rare exceptions: i would play B again only sometimes if the trial had been going for a while and i thought it would be the last round. if i ever chose B twice in a row, then i would choose B for the rest of the experiment unless it was really early and the other person played A twice

$b/c=2.5$, $E=1/8$

- my strategy was as follows: -start with A -after one B by other players, assume it was an error and continue with As -after two consecutive Bs by other player, assume he is a selfish jerk and stay with Bs for rest of interaction. Only when errors mixed up my responses and confused things did I have to become creative.
- I want to build trust with the other person. I wanted to give the other person the benefit of the doubt that his or her answer was changed by the computer.
- i didn't want to get taken advantage of.
- i wouldn't choose B after AA. After AB, I figured it could have been the glitch that made them possibly choose B 7/8 of the time - i wouldn't choose B, I'd give the other person one more chance. I always chose A first, and if my action was B it was because the 7/8 glitch made me or the other person continually played B.
- i never chose A, i always chose B.
- benefit of the doubt: maybe he meant to do A so he doesn't yet deserve retaliation
- i would play B after AA if i had a feeling the computer would change it to A.
- there was a possibility that the other person was picked A but B was what the result. everybody wanted to make as much money as they could and if they all picked B every round, nobody would get anything so it would be important to start the first few rounds with A that way I could let the other person know that if we cooperate, we can go a long way
- B could have been randomly thrown out and not their choice. to see if they randomly selected B and would throw more A choices and work cooperatively, I might choose A.
- choose A initially to earn trust so that I can eventually use B to earn more later.
- didn't really think about other person, but felt taking risk to gain money was better than getting \$0
- basically I took into consideration karma. I would click A again just to see if it wasn't the 1/8 probability that made the other person choose B, and knowing this determined the rest of my responses.
- collusion in the long run in game theory can lead to a greater maximum of welfare instead of Nash equilibrium
- I figured that if the other person had been choosing A, they would choose it again and I'd get 3 units or my 1/8th probability of my choice coming out the opposite had not yet occurred so I took my chances that it would w/ A.

- if there was no reason to believe the other person chose B on purpose/would continue to choose B, I would choose A.
- I would cooperate as long as I thought the other person was. If I had done several (2-3) rounds of A I would then do B because it would look more like it was unintentional.
- It depended a lot on the person I played. one person chose B every time so I had to do the same in order to not get taken advantage of. Other times I would mostly press A and throw in a few Bs if the person was cooperative. X .
- I pretty much always played A and assumed if the person played B it was an error. If a person played B twice in a row I would sometimes play B in the next round.
- Gain trust by starting with A for 4 rounds then change to B.
- If they play B more than once, attribute it to intention not error and respond accordingly.
- I would usually choose B during the first round, then A for a few rounds, and then B once or twice later on.
- I would always start with A, and play A the second round. If the person played B both times, I switched to B. If the person played A at least 1 of those items, I would continue with A until the 6th or 7th round. then I'd play B.
- I punished for multiple B choices, not when it was just one and could have been accidental.
- I just wanted to make the most money per round possible.
- sometimes I chose A to establish it as a pattern so deviations were more likely to be interpreted as error.
- If they had more than 2 rounds with B, I went with B too, as little chance this was due to error. Or i went with B if their choice of B was greater than 1/8.
- played A consistently until other person played B three times, then switched to B until end of round.
- Played for a number of periods and realized that the average number of periods was between 7-10, so played B to maximize profit closer to the end.
- i acted fully in self-interest. However this is a PD so my best interest goes alongside what of my "partner". the interpretation of a B depends on how many rounds of A-A have occurred. If only 1, i will think it is more probable that it was intentional. If we have gone through A-A 5 times or more, i feel it is very unlikely it was intentional.
- i gave very little thought to the feelings of the other player at all. I played B every time to maximize my points and assumed they were doing the same so that any playing of A was unintentional.
- I rarely chose A b/c it earned you least points and seemed most high-risk. risk averse: playing B nets me more points and rarely costs me.
- I gave everyone the benefit of a doubt 1 round before during and after they played B. defensively if after 3 of their B rounds in a row I would [illeg] B as well.
- by choosing A, both players were in a mutually beneficial state, but only as long as they both kept choosing.
- I tried tit for 2 tat.

- some people played B all the time so I didn't want to fall into that perpetual cycle of losing. if the person had played B before, I probably didn't think it was chance.
- mostly tit for tat strategy with the understanding that a defect might have been by error and also hoping the other would assume my defection was an error when the other was a cooperator.

b/c=4, E=1/8

- I appreciated their choosing A (I would always assume they chose it). hoped a B was a 1 time thing on their part and so didn't want to sour future relations.
- I selected B at times because I doubted the trustworthiness of the other player. [if?] I decided to trust the other person would consistently select A
- depends on how many times the player has picked B.
- I thought their B may have been the 1/8 that the answer was switched. If so, they might come back to A in this round and we'd be even again. if they kept playing B, I'd play B so I wouldn't lose more points and so they wouldn't get more points
- If I were continuously playing A then they played B out of nowhere I'd play B the next round out of revenge. No need to be greedy. if they were to continue to play B then I would to, that way no one wins or gains.
- I would choose A unless the other person picked B. If the other person picked B, the next round I would also pick B. if the person then picked A I would think that it was an error and then continue to click A. if the person pressed B, i would continue to choose B. if I had accidentally picked B, i would pick A and make it up to the other person. then if they had chosen B i would assume they were checking me and continue to choose A.
- I never chose B if the previous round was both As. B could have been the error so I was willing to give the other person benefit of the doubt.
- if I chose A, and the other person chose B, the next time I chose B as a defense mechanism so I didn't lose anymore points
- in most of the interactions, if the other person did not play A at first, they eventually played A consistently after seeing that I had played A for the start two or more times in a row. Therefore, a mutual strategy seemed to be derived where both players would play A every time in order to maximize the number of units gained in the whole interaction (mutual altruistic response). Additionally, whenever this strategy was employed and the screen showed one time that the other person chose B, i automatically attributed it to error and continued to play A for the following rounds.
- trying to get equilibrium of all-As
- it felt terrible being betrayed after a long "A" streak. assume at first that their B could be a mistake. .
- if both played A, 12 total points earned. That's why I played it.
- keep the other person honest. Compensate for errors.
- when someone, myself or the other got B I stuck with B straight through.
- after I got a B from someone I always gave them the benefit of the doubt that it was due to error. After that if I got another B I would usually give B as well.

- there was a chance it was the computer if they chose A & I chose A, then we could both get six points, which outweighs the risk of losing 2 point, which is something like 7 cents. If they were consistently choosing B (for 3 rounds) then I would switch to B as well. However, I did so knowing we could both be better off with As. I always started with A for at least the 1st 2 rounds to see how much the other person wanted to cooperate. In rounds that we both began choosing B I never made as much money.
- I always played A unless they had played B.
- as the number of rounds increased, I was more likely to play B in the later rounds (expected number of rounds per interaction=8).
- B was the dominant strategy but gave better payoffs. I always chose B.
- since this game was completely anonymous and there was a chance that my choice would be changed, I felt no incentive to be nice and choose A .
- I started with A to demonstrate good faith and hoped that the other person would too so we could establish picking A. I had originally been skeptical and started with B, but then we both ended up choosing B for each round.
- if you are perceived as trustworthy, the interaction benefits both parties.

b/c=4, E=1/16

- Random chance might have produced their Bs so I want to try to possibly salvage the relationship. I would simply ignore occasional B's in a long string of A's
- The only way to earn money was to work together most of the time. I wasn't sure if B was the players intended decision.
- It could have been an error that B was played.
- Depends on if B was just one occurrence, or if they had played lots of Bs before
- I started with B in several rounds and stuck with it. The last few rounds I started with A.
- I would continue for a string of As, but if the other player played a B 2\3 times recently I would switch to Bs.
- didn't choose B at all during all interactions
- play A. If they go two Bs in a row play B (once or twice) then get back to A's
- if we both chose A, we typically stuck with that for the remainder of tee interaction.
- play A until other person chooses B X2, then play B and see what the outcome is. Play A next round etc.
- I never chose B on purpose, and to avoid getting into a locked system of B's I never played B.
- I chose A to begin with and then about every fourth time chose B to get more earnings. I then went back to A so as to blame the previous round on the 1\16th chance of A change. If the other person played B twice in response, I chose B the rest of the time.
- Always stat with B. If other player is B, stay B until other player chooses A. If other player is A, stay A. If switch to B, give benefit of doubt due to 1\16 switch. If repeats, go to B.
- I always put A regardless of any choice given to me.

- I chose A for the whole interaction except for one B or if the other played B many times.
- I chose A until a few rounds passed. I then would throw in a B to make it seem as though the computer randomly chose.
- I played A as much as possible and then switched to B if they picked B twice in a row. Occasionally I'd try picking an A twice if we played for a long time at 0 to see if they wanted to actually switch to A.
- I started by choosing A, and continued to do so. If the other player chose B, I would choose B the next round. Then I would go back to A. I never chose B unprovoked, and if the computer generated that response for me, I would choose A the next round.
- I started by choosing A. Then I chose A always unless the other player chose B twice. If so, I chose B next round. If they chose B in that round, I chose B next. If they chose A in that round, I chose A next.
- I started choosing A until the other person chose B two times. Then I switched to B until the other person chose A. I immediately returned to choosing A as soon as the other person chose A. (A merciful strategy)
- I started play A 3 times to attempt to signal that I wanted to collaborate. If the other did not switch from B, I would change to B to attempt to change his/her play to A.
- If we got on a spree of both A's I'd pick B to get the +8 and then switch back. It eventually didn't really affect my earnings considerably.
- Started choosing B and then switched to choosing A or B once I ended up with 0 continuously.

b/c=4, E=0

- If they were helping me, I was willing to help them, 1\16 of a dollar isn't a lot to lose. If they chose B first, I felt like they were being greedy and I wanted to show them that strategy wouldn't pay off in the long run
- Always B.
- I hate when things don't make sense. To choose B in after we both chose A would hurt me because the other player would stop choosing A and I'd earn no points. Whether my motives are self-serving or altruistic, A is the most logical. If they've previously chosen B many times, I don't like them, and I don't want them to earn any points
- first I chose A. Then I started with B after seeing the first interaction, I would choose either A or B.
- going for B was not helping with long term sustained benefit.
- If the other person did A I would follow with A
- Both players win if both pick A, no one likes negative points
- If it were one of the middle rounds, I would definitely choose A. If one of the later rounds, most likely B.
- Played A every time. Only had 1 person one time choose B. Played A next round to give them a chance. I think if they had chosen B again I might have considered switching to B but my plan was to always play A.

- I chose A to take more money, collectively, from the university I would prefer another student have it that the university.
- Choosing anything other than A 100% of the time is just mean and vindictive. Its just a study why not help people get rich?
- I played A 3 to 4 times after which I only played B
- I started by choosing A and switched to B if the other player played B in the previous round, but would remain on B as long as the other player did too.
- I started by choosing B. then if the other person chose A in the next round I would choose A. If the other person chose B too, then I would continue playing B.