Cooperation to the Fullest Extent Possible? An Infinitely Repeated Games Experiment

Andrew Kloosterman*
ask5b@virginia.edu

University of Virginia

P.O. Box 400182
Charlottesville, VA 22904-4182

JEL Classification: C73, C92
Keywords: Repeated Games, Cooperation, Basin of Attraction, Risk Dominance

October 24th, 2019

Abstract

Laboratory experiments on the infinitely repeated prisoner's dilemma find little cooperation when the discount factor is near the theoretical cutoff discount factor for which cooperation can be supported in equilibrium. The explanation is that non-cooperation is the best response to most beliefs about the other player's strategy. I study a new game that reverses this prediction to understand whether cooperation can be empirically observed to the fullest extent that theory says is possible. The main finding is that there is still not pervasive cooperation, less than 50% by the end of the experiments in both treatments.

*I wish to thank Charlie Holt and participants at the North American ESA Conference 2018, BEEMA V Conference 2019, University of Virginia Experimental/Theory seminar, and Appalachian State Experimental Economics seminar for insightful comments. Also, thanks to UVA for financial support.

1 Introduction

Arguably the foremost contribution of the theory of infinitely repeated games is to show that cooperation may be supported in equilibrium, even when it is not a Nash equilibrium of the stage game. Indeed, there is often a cutoff discount factor such that cooperative outcomes are possible in equilibria of the infinitely repeated game when the discount factor equals or exceeds the cutoff. On the other hand, repetition of the non-cooperative Nash equilibrium of the stage game is also an equilibrium of the infinitely repeated game. The multiple equilibria provide an obvious empirical question. Can players coordinate on a cooperative equilibrium?

The canonical example in this environment is the infinitely repeated prisoner's dilemma. Recent laboratory evidence from this game has found that the answer to this question is sometimes: cooperation is common when the discount factor far exceeds the theoretical cutoff, but defection is common when the discount factor is close to the cutoff (see the literature review below). The coordination environment introduces strategic uncertainty, because a player does not know which strategy the other player will choose. Under this uncertainty, one equilibrium selection criteria that seems reasonable is the (size of the) *basin of attraction*: the minimal probability that the other player is cooperating for cooperation to be a best response. A closely related idea is *risk dominance*: the strategy that is the best response to the other player cooperating and defecting with equal probability. The idea of both arguments is that cooperation is more likely when it is the best response to more beliefs. Thus, the behavioral prediction in the latter case is that the risk dominant strategy will be selected, and a more continuous argument is made for the former case in which the likelihood of cooperation is predicted to increase as the basin of attraction decreases. The crucial fact in the infinitely repeated prisoner's dilemma is that *not only is defection risk dominant, but the basin of attraction is 1 when the discount factor equals the theoretical cutoff*. It only decreases as the discount factor increases (defining a second larger cutoff where cooperation is risk dominant for all discount factors above this second cutoff). Hence, the basin of attraction can explain the findings in the literature.

However, this crucial fact about the relationship between beliefs and the discount factor is an artifact of the payoffs in the prisoner's dilemma and, more generally, any stage game where there is a strategy that is both the best response to the cooperative strategy and the Nash equilibrium strategy in the stage game (e.g. any stage game with a dominant strategy). While the games that have been predominantly studied in the laboratory all have dominant strategies (e.g. defect in the prisoner's dilemma and zero contribution in public goods provision), not all cooperation games have a dominant strategy (e.g. Cournot oligopoly). In a game with no dominant strategy, it is possible for cooperation to be the best response to most beliefs for all discount factors equal to or exceeding the theoretical cutoff. For such games, perhaps cooperation at the cutoff could be observed empirically. That is the question I investigate here. Can cooperation be sustained to the fullest extent that theory says is possible?

In this paper, I construct a game with this property and investigate cooperation in a laboratory experiment where the discount factor equals the theoretical cutoff. The game modifies the infinitely repeated prisoner's dilemma by taking the dominant strategy defect and separating it into two non-cooperative strategies, one that is the best response to cooperation and one that is the Nash equilibrium. Neither new strategy is dominant, although, as in the prisoner's dilemma, cooperation is still dominated. I consider two sets of parameters where cooperation is risk dominant: Medium Risk where the basin of attraction is 0.268 and Low Risk where the basin of attraction is just 0.048. The two sets of parameters define the two infinitely repeated game treatments, Medium Risk R and Low Risk R, that are the focus of my analysis. I also consider baseline treatments with non-repeated one-shot games, Medium Risk OS and Low Risk OS, and a final control treatment, Low Risk C.

The period 0 cooperation rate for experienced subjects in Medium Risk R is 0.42 and in Low Risk R treatment is 0.35.[1] Both rates are greater than the rates in the baseline treatments, 0.03 in Medium Risk OS and 0.21 in Low Risk OS, but still not even remotely close to full cooperation. Ultimately, making cooperation the best response to most (in fact, almost all in Low Risk R) beliefs is not enough to achieve much cooperation, let alone full cooperation. In the final section of the results, I argue that it seems that most subjects are attempting to maximize profits, but fail to fully translate the stage game presented in the instructions into the infinitely repeated game that they play.

Dal Bó and Fréchette (2018) survey cooperation in laboratory experiments on infinitely repeated games. They consider 32 treatments from 15 papers that investigate the infinitely repeated prisoner's dilemma including Blonski and Spagnolo (2015), the first to apply the risk dominance criterion, and Dal Bó and Fréchette (2011), who generalized the argument to the basin of attraction.[2] The period 0 cooperation rate (in match 7, so subjects have some experience) is only about 0.2 for the two treatments where the discount factor is closest to the cutoff, and less than 0.4 for several other treatments where it is reasonably close (Dal Bó and Fréchette (2018), Figure 4). In an infinitely repeated public goods provision game, Lugovsky et al (2017) find average period 0 contributions of just 17.3% to the public good for experienced subjects (last 5 matches) in the treatment where the discount factor is close to the theoretical cutoff. Engle-Warnick and Slonim (2004) do find high levels of trust in an infinitely repeated trust game with a discount factor close to the cutoff, although the trust game is sequential so

[1] Precisely, these numbers are the proportion of choices that were cooperate in the first period of each of the last 10 infinitely repeated games the subjects played.

[2] Though published in 2015, Blonski and Spagnolo's original paper dates back to 2003.

trust by the first-mover signals cooperative behavior to the second-mover and the coordination problem studied here is not really an issue.

In the context of games without a dominant strategy, infinitely repeated Cournot oligopoly has been studied in Holt (1985), but the grid of strategies that he considered was quite large and so the coordination problem was exacerbated because there were many cooperative strategies (of course, cooperative to different degrees). He found little cooperation, likely at least partly because subjects had a hard time coordinating on which cooperative strategy to choose.[3] In fact, this result is the reason to develop my game with just 3 strategies as a minimal modification of the prisoner's dilemma to give cooperation the best shot possible.

2 Theory

The theory section has two parts. First, I recap known theory for the infinitely repeated prisoner's dilemma. Second, I introduce my modified game and generalize the theory from the first section.

2.1 The Prisoner's Dilemma

The (normalized) prisoner's dilemma has two players; two strategies for each player, cooperate (C) and defect (D); and payoff matrix

	C	D
C	$(1,1)$	$(-l,1+g)$
D	$(1+g,-l)$	$(0,0)$

The payoff matrix has two parameters: $g > 0$ is the gain obtained from defection against cooperation and $l > 0$ is the corresponding loss. I make one more typical assumption: $g - l < 1$ so (C,C) maximizes the sum of payoffs. The strategy D is strictly dominant and the unique Nash equilibrium is (D,D).

In the infinitely repeated prisoner's dilemma, the prisoner's dilemma is played each period $t = 0, 1, 2, \ldots, \infty$ and players discount future payoffs at the common rate δ. It is well-known that (C,C) is possible in a subgame perfect equilibrium (hereafter equilibrium) if and only if $\delta \geq \delta^*$ where δ^* is the smallest discount factor for which grim trigger is an equilibrium.[4] Following grim trigger gives each player 1 every period, while deviating to D gives the deviant $1+g$ in the current period but triggers mutual defection and 0 in all future periods. So following grim

[3] Van Huyck et al (2002) consider a 5 strategy game with no dominant strategy that slightly resembles my 3 strategy game. But, they do not consider a discount factor close to the cutoff, and the repeated game consists of an infinitely repeated game followed by two more periods.

[4] Grim trigger is the strategy C until any non-(C,C) outcome and then D forever after.

trigger is optimal if $1/(1-\delta) \geq (1+g)$ which solved at equality determines the cutoff discount factor:

$$\delta^* = \frac{g}{1+g}$$

The strategy defect in every period (hereafter always defect) is also an equilibrium for every discount factor. Indeed, there are multiple equilibria when $\delta \geq \delta^*$ so the game is a coordination game.

Will players cooperate in this coordination game? It is typically assumed that cooperation is conditional: players who cooperate continue to do so as long as the other player cooperates. The question then becomes will players cooperate in period 0? This is because behavior in periods $t > 0$ is determined by the period 0 outcome, cooperate forever if both cooperated and otherwise defect forever. To illustrate the tradeoffs, Blonski and Spagnolo (2015) and Dal Bó and Fréchette (2011) analyze the 2×2 static game where the strategy choices are the infinitely repeated game strategies grim trigger, representing cooperate in period 0, and always defect, representing defect in period 0. Denoting the strategies by GT and AD respectively, the payoffs of the static coordination game in general (on the left) and when $\delta = \delta^* = g/(1+g)$ (on the right) are

Coordination Game: $\delta \geq \delta^*$

	GT	AD
GT	$\left(\frac{1}{1-\delta}, \frac{1}{1-\delta}\right)$	$(-l, 1+g)$
AD	$(1+g, -l)$	$(0,0)$

Coordination Game: $\delta = \delta^*$

	GT	AD
GT	$(1+g, 1+g)$	$(-l, 1+g)$
AD	$(1+g, -l)$	$(0,0)$

There are two Nash equilibria, (GT, GT) and (AD, AD). A *basin of attraction* theory can be useful to predict selection. In this theory, a player believes the other player will play GT with probability q and AD with probability $1-q$. The player chooses the strategy that is the best response to their beliefs. There is a cutoff q^* such that GT is the best response if and only if $q \geq q^*$. This cutoff q^* is called the (size of the) *basin of attraction.* The essence of the basin of attraction is that more cooperation is predicted when more beliefs select GT as the best response. That is, more cooperation when q^* is smaller. Another way to think about q^* is as a measure of risk regarding the uncertainty about the other player's strategy. GT is *risk dominant* (Harsanyi and Selten (1988)) when $q^* \leq 1/2$. That is, it is the best response to an opponent who chooses randomly and uniformly between the two strategies. Setting the expected payoff from GT equal to the expected payoff from AD determines the cutoff:

$$q^* = \frac{(1-\delta)l}{1-(1-\delta)(1+g-l)}$$

If $\delta = \delta^* = g/(1+g)$, then $q^* = 1$. That is, a player must be 100% confident that the other player will choose GT for GT to be a best response. This can be clearly seen by looking at the right panel of the coordination game and noting that GT is a Nash equilibrium of this game, but it is weakly dominated by AD. Additionally, q^* is a continuous variable, so for $\delta \approx \delta^*$ (larger, but only slightly so), $q^* \approx 1$ implying that GT is a best response for very few beliefs in this case too. In summary, when experiments find little cooperation for discount factors near δ^*, it is not too surprising given that a player must be almost 100% confident that the other player is playing cooperatively for cooperation to be a best response.

I conclude this section with a remark on the simplification to just two strategies that form the coordination game. Behavior is likely more variable, but these two strategies capture the idea that period 0 choices effectively determine the future path of play. Even when players choose other conditionally cooperative and non-cooperative strategies, as long as theses strategies are not wildly different, the period 0 choices approximately determine the future path of play and so payoffs will be highly dependent on the period 0 choices. Hence, payoffs should be well-approximated by the simplified coordination game I have presented. This is an assumption, and I will test its empirical validity in the results of the experiment.

2.2 The Experimental Game

The reason that $\delta = \delta^*$ implies $q^* = 1$ in the infinitely repeated prisoner's dilemma is that defect is both the best response to cooperation and the Nash equilibrium strategy in the (non-repeated) prisoner's dilemma. To see why, recall that δ^* is determined by equating the payoffs of GT and the best response to C followed by the Nash equilibrium forever after. Because D is both the best response and Nash equilibrium, the latter payoff is the payoff of AD against an opponent who chooses GT. The former payoff is the payoff of GT against an opponent who chooses GT. In other words, the payoffs of GT and AD against an opponent who chooses GT are equal (e.g. see the right panel of the coordination game where these payoffs are both $1 + g$). On the other hand, GT against an opponent that plays AD yields a lower payoff than AD against an opponent that plays AD, because D is the Nash equilibrium strategy so while the payoffs in all $t \geq 1$ are equal, not playing a best response in period 0 lowers the total payoff of GT.

But this is just a feature of the prisoner's dilemma, and, more generally, any game with this property of a single strategy playing these two separate roles. This includes every game with a dominant strategy in the stage game. This experiment separates the two roles of defect into two separate strategies, Nash (N) and defect (D). Both strategies should be interpreted as non-cooperative strategies where Nash is the Nash equilibrium strategy and defect is the best response to cooperation. The payoff matrix is

	C	N	D
C	$(1,1)$	$(-l,0)$	$(-l,1+g)$
N	$(0,-l)$	$(0,0)$	$(0,-l)$
D	$(1+g,-l)$	$(-l,0)$	$(-l,-l)$

While many of the payoffs are pinned down as a modification of the prisoner's dilemma, there are additional payoffs in the 3×3 matrix, and why they were chosen this way will be made clear as the analysis progresses. The unique Nash equilibrium is (N,N).[5]

In the infinitely repeated game, (C,C) may be possible in an equilibrium following the same analysis as for the infinitely repeated prisoner's dilemma. Here, grim trigger is the strategy C until any non-(C,C) outcome and then N forever after. Then, (C,C) is possible if and only if grim trigger is an equilibrium.[6] The incentive constraints for grim trigger are the same and so the cutoff discount factor is the same in this new game:

$$\delta^* = \frac{g}{(1+g)}$$

The difference is the static coordination game. Mirroring the above analysis, grim trigger and always Nash represent two repeated game strategies corresponding to C in period 0 and N in period 0 respectively. A non-cooperative repeated game strategy corresponding to D in period 0 must be included too. I assume this strategy is defect trigger, D until any non-(C,C) outcome and then N forever after (and thus always prescribes D in period 0 and then N in all $t>0$ along the path of play). Denoting the strategies by GT, AN, and DT, the 3×3 coordination game is

Coordination Game: $\delta \geq \delta^*$

	GT	AN	DT
GT	$\left(\frac{1}{1-\delta},\frac{1}{1-\delta}\right)$	$(-l,0)$	$(-l,1+g)$
AN	$(0,-l)$	$(0,0)$	$(0,-l)$
DT	$(1+g,-l)$	$(-l,0)$	$(-l,-l)$

Coordination Game: $\delta = \delta^*$

	GT	AN	DT
GT	$(1+g,1+g)$	$(-l,0)$	$(-l,1+g)$
AN	$(0,-l)$	$(0,0)$	$(0,-l)$
DT	$(1+g,-l)$	$(-l,0)$	$(-l,-l)$

For a general 3×3 game, beliefs are represented by 2 parameters: the probability of choosing the first strategy and the probability of choosing the second strategy (with the remaining probability being the probability of choosing the third strategy). But in this game, beliefs are represented by just 1 parameter, because the payoffs from the other player choosing AN and DT are identical. Hence, it suffices to assume a player believes the other player plays GT with

[5] This can be found by iterative deletion of strictly dominated strategies. C is strictly dominated by a combination of N and D that puts sufficient weight on D. Then N dominates D.

[6] The formal reason is that 0 is still the minmax payoff so it is sufficient to assume any deviation from C is followed by N forever after. The best deviation against C is D.

probability q and any arbitrary mixture over AN or DT with probability $1 - q$. In fact, this simplification is the reason to choose the other payoffs in the 3×3 stage game payoff matrix as such so that behavioral predictions are based on a single parameter. As above, GT is a best response if and only if $q \geq q^*$ where q^* is determined by setting the expected payoffs of GT and AN equal:

$$q^* = \frac{(1-\delta)l}{1 + (1-\delta)l}$$

The key feature here is that if $\delta = \delta^*$, then $q^* = l/(1+g+l) < 1$. Indeed, q^* can be arbitrarily close to 0 when $l \approx 0$. Cooperation can be a best response to most, and in fact arbitrarily close to all, beliefs. In other words, the strategic risk that potentially deters cooperation in the prisoner's dilemma can be entirely reversed to deter non-cooperation in this game.

Finally, there is a potentially important difference between this game and the prisoner's dilemma. GT is not a unique best response in the static coordination game when $\delta = \delta^*$. The payoffs from GT and DT are identical and so both are best responses when $q \geq q^*$. It seems reasonable to hypothesize that GT may be selected (e.g. DT is not an equilibrium strategy, is Pareto dominated by GT, and simply hurts the other player without increasing one's own payoff), but this is an empirical issue, and it is addressed with a control experiment.

3 Experimental Design

There are five treatments in the experiment, a 2×2 design plus 1 control treatment. The treatment variables for the 2×2 design are the value of the loss l, set to $l = 1.1$ (Medium Risk) and $l = 0.1$ (Low Risk), and the horizon of the game, infinitely repeated (R) or played once as a one-shot game (OS). In all treatments, $g = l + 0.9$ which minimizes q^* to give cooperation the best chance.[7] In the infinitely repeated game treatments, $\delta = \delta^*$. The control treatment (Low Risk C) consists of playing the static coordination game with the $l = 0.1$ parameters. Table 1 summarizes.

Table 1: Treatments

Treatment	l	g	Horizon	δ	q^*	# of Sessions	# of Subjects
Medium Risk R	1.1	2	∞	2/3	0.268	6	76
Medium Risk OS	1.1	2	1	-	-	3	32
Low Risk R	0.1	1	∞	1/2	0.048	6	80
Low Risk OS	0.1	1	1	-	-	3	38
Low Risk C	0.1	1	1	-	-	3	40

[7] For a fixed l, minimizing $q^* = l/(1 + g + l)$ requires maximizing g. Technically, given the constraint that $g - l < 1$, there is no maximum. I set g to $l + 0.9$ to achieve an approximate minimum for q^*.

To avoid negative payoffs and mitigate any framing effects, the actual payoffs given to subjects were the affine transformation or the normalized game $20 * payoff + 32$ yielding the payoff matrices

Medium Risk

	C	N	D
C	$(52, 52)$	$(10, 32)$	$(10, 92)$
N	$(32, 10)$	$(32, 32)$	$(32, 10)$
D	$(92, 10)$	$(10, 32)$	$(10, 10)$

Low Risk

	C	N	D
C	$(52, 52)$	$(30, 32)$	$(30, 72)$
N	$(32, 30)$	$(32, 32)$	$(32, 30)$
D	$(72, 30)$	$(30, 32)$	$(30, 30)$

The corresponding payoffs in the static coordination games are

Medium Risk Coordination Game

	GT	AN	DT
GT	$(156, 156)$	$(74, 96)$	$(74, 156)$
AN	$(96, 74)$	$(96, 96)$	$(96, 74)$
DT	$(156, 74)$	$(74, 96)$	$(74, 74)$

Low Risk Coordination Game

	GT	AN	DT
GT	$(104, 104)$	$(62, 64)$	$(62, 104)$
AN	$(64, 62)$	$(64, 64)$	$(64, 62)$
DT	$(104, 62)$	$(62, 64)$	$(62, 62)$

266 subjects from the undergraduate population at the University of Virginia were recruited from the VeconLab and Darden BRAD Lab pools. All treatments were run between subjects. There were 6 sessions of Medium Risk R and Low Risk R and 3 sessions of Medium Risk OS, Low Risk OS and Low Risk C. No subject participated in more than 1 session. The infinite horizon was implemented as an indefinitely repeated game with continuation probability equal to the discount factor.[8]

After reading the instructions and taking a short quiz to test understanding of them, subjects played the game 25 times, each of which I will call a match. They were randomly rematched with another subject in the room between each match. At the end of each period in each match, subjects observed the outcome and received their payoffs. The instructions were neutrally framed and the experiment was run with z-tree (Fischbacher 2007). Instructions for Treatment Medium Risk R and Medium Risk OS are appended (and instructions for the Low Risk treatments differ only in the payoffs). Earnings were cumulative and converted to dollars at the rate of 150 points per dollar in Medium Risk R, 100 points per dollar in Low Risk R and Low Risk C, and 50 points per dollar in Medium Risk OS and Low Risk OS. This calibration means that expected earnings would be \$26 if the outcome was (C, C) in every period and \$16 if the outcome was (N, N) in every period for every treatment.[9] Average earnings were \$19.51 for sessions that

[8]This is the most common way to implement an infinite horizon. The games are equivalent if subjects are risk neutral, and, for concreteness, I will continue to use the term infinitely repeated throughout with the understanding that the games in the experiment are actually indefinitely repeated.

[9]These are exact earnings for the non-repeated games, but expected earnings for the infinitely repeated games because the total number of periods is random.

lasted about 1 hour (for the infinitely repeated game treatments) or 30 minutes (for the one-shot and coordination game treatments).

The unbalanced design, 6 sessions for some treatments and 3 for others, was because I am almost exclusively interested in the infinitely repeated games. The inclusion of one-shot games is simply to show that, to the extent that there is cooperation in the infinitely repeated game, it is not an artifact of the underlying stage game but rather an artifact of repetition. There are other potential baselines (e.g. finitely repeated games that are repeated $1/(1-\delta)$ times), but the one-shot game is the simplest.

The values of l were chosen with results from experiments on the infinitely repeated prisoner's dilemma in mind. Little cooperation is found when it is not risk dominant, $q^* > 1/2$, but some experiments document significant cooperation for $q^* \approx 1/4$ and so Medium Risk takes this as a starting point (e.g. Dal Bó and Fréchette (2011) have a treatment with $q^* = 0.27$ and find a period 0 cooperation rate of 0.61). That cooperation is risk dominant is indeed the reason to call it Medium Risk, and not, say, High Risk. After running the sessions for Medium Risk, I determined that more extreme parameters may provide a better chance for cooperation and so then ran the Low Risk treatments where cooperation is a best response to nearly all beliefs one could have.

4 Results

The results section proceeds in three parts. First, I analyze period 0 choices to measure cooperation. Second, I analyze choices from all periods to assess all behavior, and specifically, the extent to which behavior is consistent with the static coordination game. Third, I discuss the results and potential mechanisms that could explain them. Statistical significance is calculated with non-parametric tests on data averaged at the session level. Significance is noted with stars on inequalities corresponding to the standard levels ($*$ if $p < .01$, $**$ if $p < .05$, and $***$ if $p < .01$). As a check of robustness, I evaluate significance with dummy regressions and standard errors clustered at the session level in the appendix. There are a few minor differences which I note in the main text.

4.1 Period 0 Choices

This section analyzes period 0 choices for the four main treatments. This is the first choice in each match of the infinitely repeated game and the first, and only, choice in each match of the one-shot game. Focusing on period 0 abstracts away from the effects of outcomes of previous periods on choices to provide a measure of cooperation. Figures 1 and 2 present the proportion

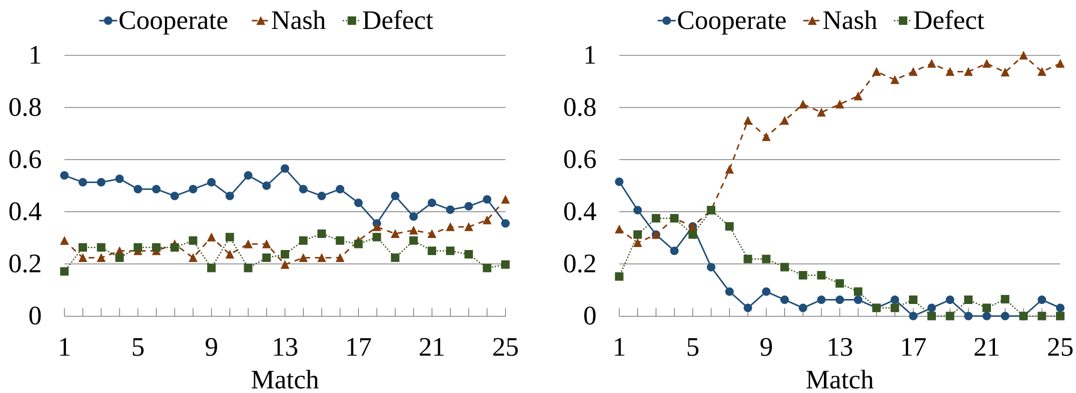

(a) Medium Risk R

(b) Medium Risk OS

Figure 1: Period 0 Choices By Match, Medium Risk Treatments

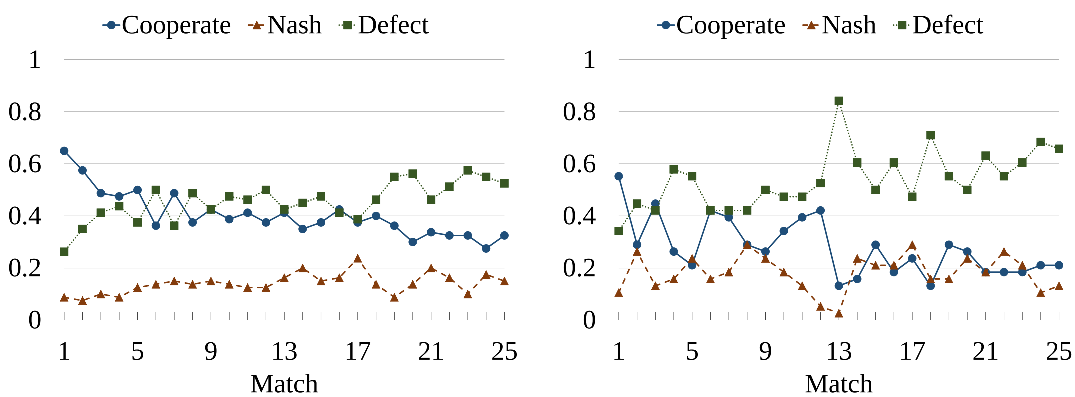

(a) Low Risk R

(b) Low Risk OS

Figure 2: Period 0 Choices By Match, Low Risk Treatments

of each choice by match.

Figure 1 presents a clear picture for the Medium Risk treatments. Behavior in Match 1 is similar between the two treatments, but quickly converges to the Nash equilibrium in the one-shot game while some cooperation is maintained in the infinitely repeated game. Still, the level of cooperation in the infinitely repeated game is well below 100% (it is not even 50% by the end of the experiments).[10] Focusing on experienced subjects, the cooperation rate in the last 10 matches in Medium Risk R is 0.42. Statistically, that rate is significantly larger than the corresponding cooperation rate of 0.03 in Medium Risk OS ($0.42 >^{**} 0.03$, Wilcoxon-Mann-Whitney test), but the main takeaway is that 0.42 is not remotely close to 1.

Figure 2 presents a similar picture for Low Risk R. As for Medium Risk R, there is cooperation

[10] Cooperation slightly decreases over time, but it's not converging to 0, which is clear from the data that is disaggregated by session. See Figure 3.

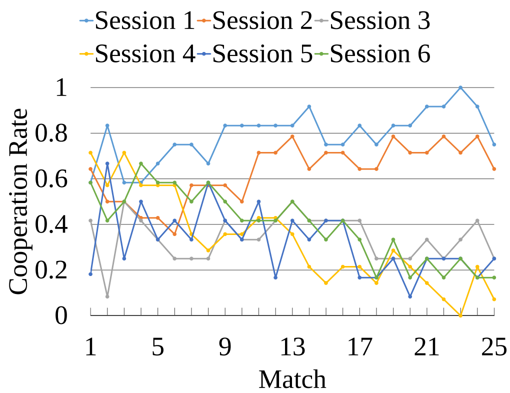

(a) Medium Risk R

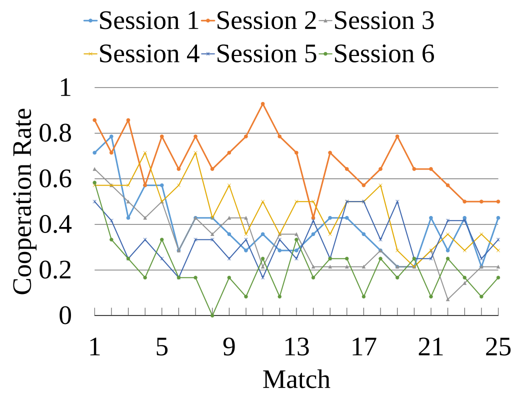

(b) Low Risk R

Figure 3: Period 0 Cooperation Rates By Match and Session

in the infinitely repeated game, but again, far less than 100% (again, less than 50% by the end of the experiments). There is a difference between the one-shot games. In Low Risk R, choices do not converge to the Nash equilibrium.[11] The cooperation rate in the last 10 matches in Low Risk R is 0.35, which is larger than the cooperation rate of 0.21 in Low Risk OS. The difference is not significant ($0.35 > 0.21$ with $p > .1$, Wilcoxon-Mann-Whitney test) although this seems to be due to non-zero cooperation in Low Risk OS.[12] Again, the main takeaway is that 0.35 is not remotely close to 1.

As mentioned, cooperation in the two infinitely repeated game treatments is very similar. For the last 10 matches, the difference in cooperation rates is just 0.07 and it is not significant ($0.42 > 0.35$ with $p > .1$, Wilcoxon-Mann-Whitney test).

The low levels of cooperation in the infinitely repeated game treatments is somewhat surprising, particularly in Low Risk R, as the experiments are designed so that cooperation is risk dominant for both sets of parameters and nearly a best response to every belief in Low Risk R. In fact, cooperation is a best response in all 12 sessions of the infinitely repeated games to the empirical distribution of choices in the first 10 matches. To illustrate, Figure 3 presents the cooperation rates by match for each of the 6 sessions of each infinitely repeated game treatment. Recall that $q^* = 0.268$ in Medium Risk R and $q^* = 0.048$ in Low Risk R so cooperation is a best response to beliefs that put probability at least 0.268 and 0.048 on the other player cooperating in the respective treatments. The actual cooperation rates in the first 10 matches for the six sessions of Medium Risk R are 0.33, 0.40, 0.51, 0.51, 0.53, and 0.71, and of Low Risk R are 0.23, 0.32, 0.46, 0.49, 0.56, and 0.74.

[11]This could be because N is only the best response to beliefs that put very high probability on the other player choosing N in Low Risk R.
[12]This difference is significant at the 10% level for the corresponding dummy regression in Appendix A.

All rates exceed q^*, which means that if beliefs were formed by experience in the early matches, cooperation would have emerged. But it only emerges in 2 of the 6 sessions of Medium Risk R, and perhaps 1 of the sessions of Low Risk R. Nevertheless, cooperation rates do change over time, particularly in Medium Risk R, suggesting that learning is occurring. To assess learning, I estimate a probit regression with the choice to cooperate in period 0 of match t as the dependent variable and potential drivers of this choice as independent variables. These drivers are initial cooperativeness measured by a dummy variable equal to 1 if the subject cooperated in period 0 of match 1 (*Coop. in Match 1*), experience from previous matches measured by a dummy variable equal to 1 if the other player cooperated in period 0 of match $t-1$ (*Other Coop. Last Match*), and experience about how long matches last measured by the number of periods in match $t-1$ (*Last Match Length*).

All three variables are predicted to have positive coefficients. The first two are quite intuitive, while the last one is expected to be positive, because subjects learn about the expected length of matches and know that cooperation is more valuable in longer matches. These variables have all been found to be significant drivers of learning in the infinitely repeated prisoner's dilemma (e.g. Dal Bó and Fréchette (2011)). The regression includes data from all matches greater than 1 and standard errors are clustered at the session level. The results are in Table 2.

Table 2: Learning Model

	Medium Risk R	Low Risk R
Coop. in Match 1	0.756***	0.863***
	(0.186)	(0.169)
Other Coop. Last Match	0.570***	0.314***
	(0.167)	(0.0946)
Last Match Length	0.0183***	0.0791***
	(0.00926)	(0.0144)
constant	-0.832***	-1.159***
	(0.158)	(0.169)
N	1824	1920

Probit: dependent variable 1 if cooperate

Period 0 choices for matches greater than 1

s.e. clustered by session in parentheses

$*$ if $p < .1$, $**$ if $p < .05$, $***$ if $p < .01$

The regressions indicate that all three of these variables influence the decision to cooperate in the expected direction. Match 1 cooperation is highly predictive of future cooperation and both cooperation by the other subject in the previous match and more periods in the previous match lead to more cooperation too. The positive coefficients on *Other Coop. Last Match* indicate that subjects are learning about the composition of cooperators in the room and responding with more cooperation when they meet a cooperative other player. This may seem to be at odds with the finding that period 0 cooperation is rarely emerging even though it is the best response to the empirical distribution of choices. However, it is properly interpreted as only more cooperation after experiencing cooperation than after experiencing non-cooperation and does not indicate whether the level of cooperation is increasing or decreasing. Indeed, it could be that cooperation decreases more after non-cooperation than it increases after cooperation, and this seems to be the case here.

4.2 Strategies and Payoffs

Perhaps the reason for low cooperation rates in period 0 is the restriction to the three strategies (grim trigger, always Nash, and defect trigger) under which q^* was calculated, and that actually strategies were consequentially different to the point where q^* would be meaningless. In this section, I analyze behavior and payoffs in all periods to argue that this does not seem to be the case.

The main behavioral assumption behind the restriction to three strategies is that cooperation is conditionally cooperative. To analyze this assumption, I investigate cooperation in periods $t > 0$ conditional on the period $t - 1$ outcome. Here, I consider the last 10 matches.[13] Table 3 presents the cooperation rates after each possible $t-1$ outcome. The number at the intersection of row x and column y corresponds to the cooperation rate for subjects who chose x while their opponent chose y in the previous period.

Table 3: Cooperation Rate By Last Outcome

	Medium Risk R				Low Risk R		
	C	N	D		C	N	D
C	0.90	0.52	0.05	C	0.62	0.51	0.30
N	0.35	0.08	0.06	N	0.49	0.22	0.16
D	0.18	0.09	0.11	D	0.32	0.10	0.17

Matches 16-25, Periods $t > 0$

If subjects strictly played the three strategies considered in the theory section then the cooperation rate would be 1 after the outcome (C, C) and 0 after all other outcomes. Not surprisingly, the table indicates play is somewhat more varied. However, cooperation does seem to

[13] I consider robustness checks for only $t = 1$ choices and all matches in the appendix. The results are qualitatively similar.

be largely conditional as the largest cooperation rate is after the outcome (C, C) in the previous period. Statistically, the cooperation rate after (C, C) is larger than the cooperation rate after all other outcomes for Medium Risk R ($0.9 >^{***} 0.52, 0.05, 0.35, 0.08, 0.06, 0.18, 0.09, 0.11$, Wilcoxon-Mann-Whitney tests) and larger than the cooperation rate after 6 other outcomes for Low Risk R ($0.62 >^{***} 0.30, 0.22, 0.16, 0.10, 0.17$; $0.62 >^{**} 0.32$, Wilcoxon-Mann-Whitney tests).[14]

In Medium Risk R, the rates are actually fairly close to the theoretical point predictions except after the outcomes (C, N) and (N, C) where there is intermediate cooperation. In these two cases, there is miscoordination, but no one benefited at the expense of the other which probably explains the intermediate cooperation rate. However, this difference does not explain low period 0 cooperation rates as the value of cooperate is only increased by the possibility of coordination in later periods. In Low Risk R, the same result holds qualitatively, but the differences between the rates among the various previous period outcomes is not nearly as large.

The point of the theoretical simplification is that period 0 choices determine what will happen in the entire infinitely repeated game. In particular, each period 0 outcome is associated with a payoff for the entire game which is what simplifies it into the static coordination game. Then, the decision of what strategy to choose in the coordination game is characterized by q^*. Even though behavior in periods greater than 0 is somewhat more varied, if each period 0 outcome is associated with a payoff for the entire game that is close to the payoffs in the simplified coordination game, then q^* will still approximately characterize the period 0 decision.

To illustrate how close payoffs are, Table 4 presents the predicted and actual average payoffs conditional on the period 0 outcome for the last 10 matches.[15] The number at the intersection of row x and column y corresponds to predicted and actual average infinitely repeated game payoffs when the subject chose x in period 0 while the other player chose y in period 0.

The main takeaway is that the actual payoffs in Table 4 are fairly close to the predicted payoffs. They are within 10% for all nine period 0 outcomes in Medium Risk R. They are within 10% for all the period 0 outcomes except (C, N), (C, D) and (N, C) for Low Risk R where they are between 10% and 20% larger than predicted. But the latter finding only makes C and, to an extent N, more attractive. Statistically, there are a number of differences ($147 <^{**} 161$, $94 <^{*} 96$, $146 <^{**} 151$, $73 <^{**} 78$ for Medium Risk R, $116 <^{**} 128$, $78 >^{**} 66$, $84 >^{**} 73$, $79 >^{**} 68$, $123 >^{**} 115$, $60 >^{**} 58$ for Low Risk R, Wilcoxon signed rank tests).[16] For the

[14]The comparison of 0.62 to 0.49 is significant at the 5% level for the dummy regression in Appendix A.

[15]The predicted payoffs would be the coordination game payoffs from Section 3 if the average number of periods was exactly equal to the expected number of periods. The random draws and low number of data in some cells led to some fairly large differences, and thus I have decided to report the predicted payoffs which come from strategies in the coordination game but with the actual average number of periods. For example, the average number of periods for matches in Medium Risk R where the period 0 outcome was (C, N) was 2.65 so the predicted payoff is 10 in period 0 and then 32 for 1.65 periods or $10 + 1.65(32) = 63$.

[16]The dummy regressions in Appendix A show a few differences in significance. Mostly, these are marginally significant with the non-parametric test and significant with the parametric test, or vice versa.

significant differences, payoffs are less for one period 0 outcome for each of the three possible choices in Medium Risk R, and greater for one to two period 0 outcomes for each of the three possible choices in Low Risk R indicating that there is no choice that is systematically better or worse than predicted. Overall, it seems that the period 0 choices determine infinitely repeated game payoffs in a way that is pretty similar to how payoffs are determined by the three strategies alone and so q^* is indeed a relevant parameter.

Table 4: Payoffs By Period 0 Outcome

Medium Risk R Predicted

	C	N	D
C	161	63	69
N	85	96	100
D	151	78	57

Medium Risk R Actual

	C	N	D
C	147	67	67
N	92	94	98
D	146	73	56

Low Risk R Predicted

	C	N	D
C	128	66	73
N	68	51	61
D	115	59	58

Low Risk R Actual

	C	N	D
C	116	78	84
N	79	53	63
D	123	63	60

Matches 16-25

4.3 Discussion

The results indicate that the static coordination game is a reasonably good simplification, and so cooperation in period 0 should be approximately a best response for any player who believes the other player cooperates with probability at least q^*. Puzzlingly, even though q^* is small, almost 0 in Low Risk R, little cooperation is observed.

The first potential explanation I explore is that it is just best response behavior. Defect is also a best response in the coordination game and so perhaps subjects perfectly understand that they are playing (something akin to) the coordination game but choose defect as a best response anyways. Indeed, the majority of period 0 non-cooperation in Low Risk R is defection. However, the control treatment Low Risk C provides evidence against this explanation. This is, in fact, the only purpose of this treatment, and thus why I have waited until now to discuss it. Figure 4 presents the proportion of each choice by match for this treatment. For the last 10 matches, the

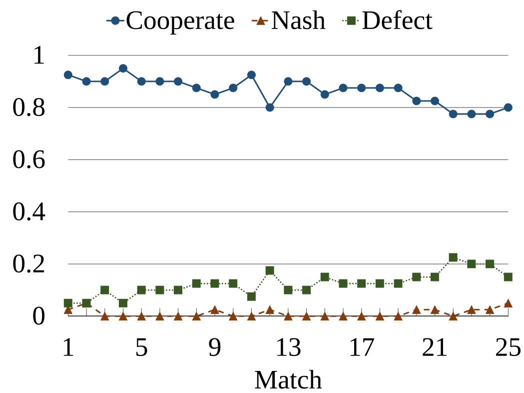

Figure 4: Period 0 Choices By Match, Low Risk C

cooperation rate is .83 in this treatment. That is, a large majority of subjects choose cooperate as their best response.

A second potential explanation is risk aversion. A risk averse subject would be more likely to choose Nash in period 0 (where payoffs are constant across the other player's strategies). One way to think of this is that q^* was developed in a model where agents are risk neutral, and q^* would be larger in any model where agents are risk averse. A number of subjects choose Nash in Medium Risk R, but as just noted, very few do in Low Risk R where the result is most perplexing.

A third potential explanation for low cooperation is non-standard preferences in which emotions or the other player's payoff enter into the subject's utility function. At first glance, this seems less likely because intuition suggests non-standard preferences should only increase cooperation. For example, if cooperation is expected to be reciprocated, then there is a stronger incentive to cooperate. However, there are models in which cooperation is decreased. Perhaps most plausibly, combining inequality aversion (Fehr and Schmidt (1999)) with strategic uncertainty could imply that GT is not a best response in the coordination game. For example, a player who is strongly averse to receiving less than the other player (and puts significant probability on the other player choosing DT) will not want to play GT because there is the possibility to receive only 74 while the other receives 156.

So to provide some empirical evidence to address this possibility, I turn to responses from a post-experiment survey where subjects were asked to explain their strategy as best as they could. The details of the survey are provided in the appendix. I focus on Low Risk R, because that is where results are most puzzling, and non-cooperative subjects as this is the behavior I am trying to explain. There are 35 subjects in Low Risk R who chose defect or Nash in period 0 for at least 7 out of the last 10 matches (an ad hoc approximation of subjects who have a strategy).

Only 7 of them mention anything related to other's payoff or emotions. Alternatively, 17 of them mention trying to maximize their own payoff.[17] It seems that non-standard preferences, while possibly a part of the explanation, are not the main driver of low cooperation rates.

A fourth potential explanation is that people are trying to maximize payoffs, but just fail to do so because the environment is complicated. The 17 mentions in the survey of such motivations indicate evidence in favor of this explanation. And there are two specific complications of this environment that could decrease cooperation, both coming from the fact that subjects only receive the stage game payoffs in the instructions. First, they must calculate their expected payoffs in the coordination game. The main component of this calculation is the expected number of periods. The positive coefficient on *Other Coop. Last Match* suggests that this is learned from experience rather than known. The geometric distribution that the number of periods is drawn from is heavily right-skewed, which means subjects experience many very short matches and only a few very long matches. To the extent that this leads to overweighting short matches, subjects may underestimate the expected number of periods and this makes non-cooperation more attractive. Second, cooperation is a dominated strategy in the stage game and so subjects may be dismissing it for this reason. Indeed, some of the 17 subjects who mentioned that they were trying to maximize earnings indicated exactly this (e.g. "I always started with a dominant strategy of picking C to try to maximize the amount of points I would get." or "... I was never going to pick A because it simply was always less points than C").[18]

Of the four explanations here, the last one holds up best upon weighing the evidence. However, it is probably the case that there is no one single reason that cooperation cannot be supported to the fullest extent possible. Better analyzing these reasons (and potentially others), though beyond the scope of this paper due to the need for new cleverly designed treatments, is a promising avenue for future research.

5 Conclusion

In this paper, I investigate an infinitely repeated game where cooperation is the best response to most beliefs about the strategy of the other player. Even when almost all beliefs support cooperation as a best response, cooperation rates are well below 100%. Indeed, by the end of sessions the average rates fall below 50%. This indicates that there is more than strategic uncertainty determining behavior in infinitely repeated games.

The paper opens several avenues for future research. First, it is possible that cooperation emerges in a different game than this one. By choosing the game to minimally modify the

[17] The rest either have non-intelligible responses or do not mention motivations at all. Additionally, none of the 35 subjects mentioned risk as a motivating factor in their survey response so this provides additional evidence that risk is not a factor.

[18] Option C is defect and option A is cooperate in the neutral framing.

prisoner's dilemma, and therefore minimize the coordination problem, it seems that this game is as ideally suited as any for cooperation, but perhaps not. Second, it would certainly be worthwhile to delve further into why the infinitely repeated game exhibits little cooperation. This could perhaps be accomplished with more experiments where the game is presented to subjects as something between Low Risk R and Low Risk C (such as forcing them to choose one of the three strategies in the theory, but otherwise letting the game play out across the infinite horizon). Overall, this paper presents a new game that further opens the door to new questions with the ultimate goal of determining when and why people might cooperate in infinitely repeated games.

Appendix A: Additional Statistics and Tables

First, I calculate significance with dummy regressions and standard errors clustered at the session level. For example, the first comparison below takes all period 0 cooperation in the last ten matches in Medium Risk R and Medium Risk OS as the data set and then runs a probit regression with a dummy equal to 1 if cooperation is chosen as the dependent variable and a dummy equal to 1 if the treatment is Medium Risk R. Standard errors are clustered at the session level in the regression, and significance is the p-value of the estimated coefficient for the Medium Risk R dummy.

For ease of comparison to the non-parametric results, statistics are characterized as non-significant ($p > .1$), marginally significant ($.05 < p < .1$), or significant ($p < .05$). I highlight in bold the cases where the statistic is non-significant with one test and significant with the other to illustrate the biggest differences in agreement. I highlight in italics the other cases where on test is marginally significant and the other is either non-significant or significat to illustrate smaller differences in agreement.

The first set of these regressions are probit regressions where the dependent variable is 1 if cooperate and 0 otherwise. The independent variable is another dummy variable equal to 1 for one of the two values of the variable of interest (a treatment dummy for period 0 cooperation and a last outcome was (C, C) dummy for period t cooperation).

For period 0 cooperation:

- Cooperation is larger in Medium Risk R than Medium Risk OS ($.42 >^{***} .03$).
- *Cooperation is marginally larger in Low Risk R than Low Risk OS* ($.35 >^{*} .21$).
- Cooperation is larger, but not significantly, in Medium Risk R than Low Risk R ($.42 > .35$).

For period t cooperation, by last outcome:

- Cooperation is larger after (C, C) than all 8 other outcomes in Medium Risk R ($.9 >^{***} .52, .05, .35, .08, .06, .18, .09, .11$).
- Cooperation is larger after (C, C) than 5 other outcomes in Low Risk R at the 1% level ($.62 >^{***} .30, .22, .16, .10, .17$).
- Cooperation is larger after (C, C) than 1 other outcome in Low Risk R at the 5% level ($.62 >^{**} .32$).
- **Cooperation is larger after (C, C) than 1 more outcome in Low Risk R at the 5% level** ($.62 >^{**} .49$).

- Cooperation is larger, but not significantly, after (C, C) than the last other outcome in Low Risk R ($.62 > .51$).

The final regressions are OLS regressions where the dependent variable is the difference between the total payoff a subject earned in the match and the expected payoff based on the period 0 outcome if the players played the three strategies in the theory section exclusively. The independent variable is just a constant (so the coefficient is the average difference) and it is tested whether it equals 0 or not.

- The difference is significant at the 1% level for one comparison in Medium Risk R ($147 <^{***} 161$).
- The difference is significant at the 5% level for one comparison in Medium Risk R ($146 <^{**} 151$).
- *The difference is significant at the 5% level for one more comparison in Medium Risk R* ($94 <^{**} 96$).
- *The difference is marginally significant at the 10% level for one comparison in Medium Risk R* ($73 <^{*} 78$).
- The difference is significant at the 1% level for one comparison in Low Risk R ($79 >^{***} 68$).
- **The difference is significant at the 5% level for one comparison in Low Risk R** ($63 >^{**} 59$).
- *The difference is marginally significant at the 10% level for four comparisons in Low Risk R* ($116 <^{***} 128$, $78 >^{*} 66$, $84 >^{*} 73$, $123 >^{*} 115$).
- **The difference is positive, but not significantly so, for 1 comparison in Low Risk R** ($.60 > .58$).
- The last 7 differences are not significant.

Now, I recreate Table 3 with just period 1 data (to abstract from the effects of outcomes in period $t - 2$ or earlier) and all matches (to note that little is changed by including decisions made when subjects were inexperienced).

Table 3′: Cooperation Rate By Last Outcome

Medium Risk R					Low Risk R			
	C	N	D			C	N	D
C	.91	.55	.04		C	.64	.54	.30
N	.47	.22	.04		N	.5	.36	.18
D	.26	.07	.13		D	.33	.07	.17

Matches 16-25, Period $t = 1$

Table 3″: Cooperation Rate By Last Outcome

Medium Risk R					Low Risk R			
	C	N	D			C	N	D
C	.89	.49	.14		C	.64	.47	.29
N	.38	.08	.06		N	.46	.24	.18
D	.19	.12	.16		D	.36	.27	.22

Matches 1-25, Periods $t > 0$

Table 3‴: Cooperation Rate By Last Outcome

Medium Risk R					Low Risk R			
	C	N	D			C	N	D
C	.87	.51	.10		C	.63	.45	.29
N	.49	.20	.05		N	.5	.32	.21
D	.25	.13	.14		D	.34	.27	.22

Matches 1-25, Period $t = 1$

Now, I do the same for Table 4 for all matches.

Table 4′: Payoffs By Period 0 Outcome

Medium Risk R Predicted

	C	N	D
C	163	81	89
N	105	106	100
D	170	81	88

Medium Risk R Actual

	C	N	D
C	149	86	88
N	112	103	97
D	164	78	82

Low Risk R Predicted

	C	N	D
C	115	64	66
N	66	60	62
D	108	60	62

Low Risk R Actual

	C	N	D
C	107	76	75
N	75	62	65
D	115	65	66

Matches 1-25

Appendix B: Instructions

Medium Risk R

Welcome. This is an experiment in decision making. Various research foundations and institutions have provided funding for this experiment and you will have the opportunity to make a considerable amount of money which will be paid to you at the end. Make sure you pay close attention to the instructions because the choices you make will influence the amount of money you will take home with you today. Please ask questions if any instructions are unclear.

The Choice Problem

In this experiment, you will engage in the following two-person choice problem with one other participant from this room. You will each choose between three options, which we have labeled A, B, and C. The combination of your choice and the other participant's choice determine how many points you earn and how many points the other participant earns. The following table describes these points:

Your Choice: Row – Other's Choice: Column	A	B	C
A	(52, 52)	(10, 32)	(10, 92)
B	(32, 10)	(32, 32)	(32, 10)
C	(92, 10)	(10, 32)	(10, 10)

To understand this table, think of yourself as choosing a row and the other participant as choosing a column. Then points are described by the cell at the intersection of the chosen row and column. The first number is the number of points you earn and the second number is the number of points the other participant earns. For example, if you choose A and the other participant chooses B, then you will earn 10 points and the other participant will earn 32 points. This can be seen in the table by looking at the cell where row A and column B intersect.

Procedures

The experiment will consist of 25 identical Matches that each proceed as follows. At the beginning of each Match you will be randomly matched with another participant from this room. You will then engage in the choice problem described above with this other participant a randomly determined number of times, each of which we will call a Round. In particular, at the end of each Round, the computer is programmed to select to continue to another Round with 2/3 chance and to end the Match with the remaining 1/3 chance. Let us stress that you engage in the choice problem with the same other participant in each Round of a Match, but then get matched to a different participant when a new Match begins.

Your screen will be laid out as follows. In the upper left corner you will see which of the 25 Matches you are currently in. At the top in the middle you will see what Round you are in.

Below this you will have the table (from above) to remind you about how points are earned. Below the table you make your Choice: A, B, or C, for the current Round by clicking on the corresponding button.

You can also see the outcomes and points earned for all past choice problems. On the left side of your screen, the previous Rounds of the current Match are displayed. On the right side of your screen you can enter the Match number (then click Check) of any previous Match to see that Match (this option is only available before you make your choice each Round). A second screen will show up after each Round as well. It will tell you the outcome and points earned for that Round and whether the Match will continue or end.

At the end of the experiment, we will give you 1 dollar for every 150 points you have earned. You will also get 6 dollars for participating.

Medium Risk OS

Welcome. This is an experiment in decision making. Various research foundations and institutions have provided funding for this experiment and you will have the opportunity to make a considerable amount of money which will be paid to you at the end. Make sure you pay close attention to the instructions because the choices you make will influence the amount of money you will take home with you today. Please ask questions if any instructions are unclear.

The Choice Problem

In this experiment, you will engage in the following two-person choice problem with one other participant from this room. You will each choose between three options, which we have labeled A, B, and C. The combination of your choice and the other participant's choice determine how many points you earn and how many points the other participant earns. The following table describes these points:

Your Choice: Row – Other's Choice: Column	A	B	C
A	(52, 52)	(10, 32)	(10, 92)
B	(32, 10)	(32, 32)	(32, 10)
C	(92, 10)	(10, 32)	(10, 10)

To understand this table, think of yourself as choosing a row and the other participant as choosing a column. Then points are described by the cell at the intersection of the chosen row and column. The first number is the number of points you earn and the second number is the number of points the other participant earns. For example, if you choose A and the other participant chooses B, then you will earn 10 points and the other participant will earn 32 points. This can be seen in the table by looking at the cell where row A and column B intersect.

Procedures

The experiment will consist of engaging in the choice problem above 25 times, each of which we will call a Match. You will be randomly matched with another participant from this room for each Match.

Your screen will be laid out as follows. In the upper left corner you will see which of the 25 Matches you are currently in. In the middle, you will have the table (from above) to remind you about how points are earned. Below the table you make your Choice: A, B, or C, for the current Match by clicking on the corresponding button.

You can also see the outcomes and points earned for all past Matches. This will be on the left side of your screen. A second screen will show up after each Match as well. It will tell you the outcome and points earned for that Match. At the end of the experiment, we will give you 1 dollar for every 50 points you have earned. You will also get 6 dollars for participating.

Appendix C: Survey Responses

At the end of the experiment, the subjects completed a short unincentivized survey. The survey collected demographic information and then asked three open-ended questions. In this appendix, I focus on the first of these open-ended questions: "At the start of each Match, did you have a plan for how you would behave or did you just make decisions as you went along? If you had a plan, please try to describe it as concisely as you can." In this section, I provide the responses from 35 subjects who chose defect or Nash in period 0 of at least 7 out of the 10 last matches in Low Risk R. This is ad hoc selection of those who seem to have employed some sort of strategy in the treatment where the results are most surprising. Data for all subjects (and the answers to other questions) is available upon request.

At the beginning of each response, I note the defect or Nash rate, and I write SP (for social preference) if they mention other's preferences or emotions and PM (for point maximization) if they mention maximizing points.

Defect in at least 7 out of the 10 last matches

1. (rate=.7) SP. I was being cooperative at first i assumed most people would choose A so we could both receive the same points but that wasn't the case since a lot acted selfish since the beginning. It led me to react the same or choose C in which we'd get the same lower number.

2. (rate=.7) PM. At the beginning, I had no set plan. I was just waiting to see the other person's first move and then I decided what to do. For example, if they chose, they were more likely to rely on my goodwill and keep choosing A (as I initially did). After a round, I would change my answer to C even though the other person continued using A. That would give me more points. After a while, I started going for C from the get go. Choosing B was neutral but choosing C was neutral or high reward. A meant that there was a chance we would divide it 5050 but the odds of that were low and choosing A was dangerous at best.

3. (rate=1) PM. I pretty much always chose C first because worst case I made 30 points but I usually made 72

4. (rate=.7) Yes. I planned to choose A first since if my partner and I both chose A, we would earn the most. But, if my partner didn't choose A too, I would choose C or B after that round.

5. (rate=.8) PM. Initially, I tried to start with A because it was the most mutually beneficial if the match partner also selected A. However, most people started with C, so that was not a very profitable choice. I started with C because there were other people that started with

A several times, so which made that option somewhat lucrative. After choosing C, I chose B for every round because on the whole it was the most likely to be profitable, as each possible number of points was the same, so the result was not dependent on the choices of others.

6. (rate=.7) I tended to start by picking C and then would make the rest of the decisions as i went along

7. (rate=.8) PM. I started with C everytime because it seemed that everyone was choosing A first and that would get me the most possible points. After that I figured my partner would try to choose C to get 72 points so I chose B to get 32 points. After that it was pretty much a mind game, there were some people that used patterns so I could identify those and use them.

8. (rate=1) Planned. Used B first and guessed what would be their nect choice based on their first one and their reaction to mine.

9. (rate=1) PM. I had a plan. I would choose C for every first round. If I get 72 points, i would put C for the next round as well. If i get 30 points, i would change my choice to B for the next round.

10. (rate=1) PM. Select C at the start of each Match to maximize personal potential profit. If matched partner also selected C, move to B to maximize personal potential profit in light of their strategy.

11. (rate=.9) SP, PM. I wanted to do A each time to collude but people were not catching on so I started doing C. Also, I would try to lower their winnings rather than make it equal

12. (rate=.7) I started with C often.

13. (rate=1) PM. Choosing C had the chance of yielding the most points so I consistently chose that option

14. (rate=.8) Made decisions as I went along

15. (rate=.9) SP. Yes, I tried to stragetize going between a and c so that my partner and i could earn the max number

16. (rate=.7) SP. Ideally, I would have only chose option A, with the thought that the other player would also chose A, creating a 50-50 split. However, as the matches went along more and more players continued to chose C. I figured I was constantly earning 30 points if I chose A or C, so screw the other person, I chose C from then on. Occasionally the other

play would chose A when I now chose C, and so if another round took place I would chose A again to either give them the higher reward the next match or we would 50-50 split.

17. (rate=1) SP, PM. I always started with a dominant strategy of picking C to try to maximize the amount of points I would get. This strategy paid off some of the time, but I also ensured that the other person did not get more points than me in the first round.
18. (rate=.7) PM. The only plan that I had was that at the beginning of most of the matches, I chose C because most of the time when I did that, I got 72 points. Other than that, I had no plan as the match continued.
19. (rate=1) PM. Go with C basically always, occassionally I went with B but it would not have made a difference in my final earnings if I had just gone with C everytime rather than B. The extra 2 cents is'nt worth it, and if the other person is putting A anyways, I would rather make 72 than 52.
20. (rate=.7) I didn't have a plan but I did have preferences and made observations as to which choices would be best in certain situations regarding how the other player acted.
21. (rate=1) PM. I decided to stick with C for most answers because there was a chance that I could get 72 points. And if not, I would at least get 30.
22. (rate=1) Start with C, but switch to B if it continued. After that, it depended on how the match went.
23. (rate=.8) I just made decisions as I went along. I first I was confused by the concept but after several minutes I got the hang of it.
24. (rate=1) PM. Yes I had a plan. I figured the expected value of choosing the C option would always be better than the other 2 options. I was never going to pick A because it simply was always less points than C.
25. (rate=.7) I did not plan the initial choice at the start of each Match, but did make choices regarding subsequent choices as the rounds went on.
26. (rate=.7) PM. At first, I planned on starting with A but then more and more people I faced started with C so I switched my method to starting with B. I tried to start with C a few times to maximize my earnings, hoping that some others (like I did at first) would start with A
27. (rate=.8) PM. Yes. I went for C almost every time at first in the later matches in the hopes that my partner would pick A and I would get 72 points.

28. (rate=1) PM. I made decisions as I went along in the game. However, after the first three or so matches, I realized that there is barely any benefit in me choosing the option A since 2/3 of the times I would be earning less points as my opponent. Knowing this, I went along in the game by always choosing the option C for the first round since it not only has the chance to win me 72 points but also the chance to either win 30 or 32 points which isn't bad. If my opponent chose C, then I would then choose B for the next round assuming that my opponent would continue choosing C as his/her answer: this way, I would earn 32 points instead of a 30. If my opponent were to choose B for the first round, I would do the same and choose B in the next round to guarantee myself at least 32 points.

29. (rate=.8) SP. I wanted to start by choosing A, and hoped that my partner would do the same, so as to keep getting 52. When I got burned the first time, I always either did B or C first, never A

<u>Nash in at least 7 out of the 10 last matches</u>

1. (rate=1) SP. I was going to pick the selection that would benefit both parties to the max

2. (rate=.7) PM. I planned to start with B because that would enable me to have the highest points if somebody chose C. After that I would choose C, hoping that the person would choose A.

3. (rate=.8) I did not have a plan at the beginning, but then came up a plan as the experiments went. In my plan, I firstly chose B for at least two Rounds to see how the other participant behaved. If the other participant chose C or B in both Rounds, I would then continued to choose B, if the other participant chose A, I would then chose C.

4. (rate=1) I select B. If the other chooses C, I continue selecting B. Otherwise, I select A next round since I feel I can trust the other.

5. (rate=.7) At the beginning I chose C every time because I anticpated people to be more trusting at that point.

6. (rate=.7) At the start, I was thinking that picking A as my first choice would be my best outcome because I assumed the other player would pick A as well so we can both earn the same amount. I was also thinking that if the other person had picked C, then the next round they could pick A and I would pick C. If we had switched back and forth, it would've been the same amount, but I was assuming too much.

References

[1] Blonski, Matthias and Spagnolo, Giancarlo. Prisoners' Other Dilemma. *International Journal of Game Theory* 2015. 44(1), pp. 61-81.

[2] Dal Bó, Pedro and Fréchette, Guillaume. The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence. *American Economic Review* 2011. 101(1), pp. 411-429.

[3] Dal Bó, Pedro and Fréchette, Guillaume. On the Determinants of Cooperation in Infinitely Repeated Games: A Survey. *Journal of Economic Literature* 2018. 56(1), pp. 60-114.

[4] Engle-Warnick, Jim and Slonim, Robert. The Evolution of Strategies in a Repeated Trust Game. *Journal of Economic Behavior and Organization* 2004. 55, pp. 553-573.

[5] Fehr, Ernst and Schmidt, Klaus M. A Theory of Fairness, Competition, and Cooperation. *The Quarterly Journal of Economics* 1999. 114(3), pp. 817-868.

[6] Fischbacher, Urs. z-Tree: Zurich Toolbox for Ready-made Economic Experiments. *Experimental Economics* 2007. 10(2), pp. 171-178.

[7] Harsanyi, John, and Selten, Reinhard. A General Theory of Equilibrium Selection in Games. 1988. MIT Press.

[8] Holt, Charles. An Experimental Test of the Consistent-Conjectures Hypothesis. *American Economic Review* 1985. 75, pp. 315-325.

[9] Lugovsky, Volodymyr, Puzzello, Daniela, Sorensen, Andrea, Walker, James, and Arlington Williams. An Experimental Study of Finitely and Infinitely Repeated Linear Public Goods Games. *Games and Economic Behavior* 2017. 102, pp. 286-302.

[10] Van Huyck, John, Wildenthal, John, and Raymond Battalio. Tacit Cooperation, Strategic Uncertainty, and Coordination Failure: Evidence from Repeated Dominance Solvable Games. *Games and Economic Behavior* 2002. 38, pp.156-175.