The Effects of Time Preferences on Cooperation: Experimental Evidence from Infinitely Repeated Games*

Jeongbin Kim†

California Institute of Technology

May 2020

Abstract

This paper studies the effects of time preferences on cooperation in an infinitely repeated prisoner's dilemma game experiment. Subjects play repeated games in the lab all at once, but stage game payoffs are paid over an extended period of time. Changing the time window of stage game payoffs (weekly or monthly) varies discount factors and a delay for the first-stage game payoffs eliminates/weakens present bias. First, subjects with weekly payments cooperate more than subjects with monthly payments—higher discount factors promote greater cooperation. Second, the rate of cooperation is higher when there is a delay—present bias reduces cooperation.

JEL codes: C91, D91

Keywords: Repeated game experiment, Discount factor, Present bias, Cooperation

*I am very grateful to my advisor Pedro Dal Bó for his guidance, financial support, and encouragement. Tom Palfrey's thoughtful guidance and comments are greatly appreciated. I thank Wooyoung Lim, Louis Putterman, Charlie Sprenger, and Sevgi Yuksel for helpful comments. I especially thank Geoffroy de Clippel for inspiring the design of this experiment.

†Postdoctoral scholar in economics, Division of the Humanities and Social Sciences. Email: jbkim@caltech.edu

1 Introduction

The theoretical contributions on infinitely repeated games have shown that time preferences are essential in determining equilibrium behavior. Punishments for opportunistic behavior can be credible if players are patient enough. Most works assume exponential discounting (i.e., dynamic consistency) and use the discount factor to characterize the set of equilibrium outcomes.[1] To test the predictions for repeated games, there have been many experiments in the laboratory.

A random termination rule is the most frequently used way of implementing repeated games in the lab (Roth and Murnighan, 1978).[2] With plausible assumptions such as risk neutrality, the equivalence between random termination and exponential discounting allows for the assessment of the effect of discount factors on behavior. Dal Bó and Fréchette (2018) show in their meta-analysis that the higher the continuation probability is, the higher the rate of cooperation in infinitely repeated prisoner's dilemma (PD) games. That is, a random termination rule allows us to observe behaviors that are consistent with the theoretical predictions.

This paper extends the scope of the literature by raising an important question: how can we test the effect of subjects' underlying "time" preferences (discount factor and present bias) on cooperative behavior in repeated game (and other dynamic games) experiments? The recent developments in empirical and theoretical investigations of time preferences lead us to the two directions. First, although a random termination rule is theoretically equivalent to exponential discounting, assessing randomness of termination in a repeated game experiment may be more suitable in the domain of risk preferences rather than time preferences. Moreover, recent experiments provide evidence that there exist clear differences between risk and time preferences. For instance, Andreoni and Sprenger(2012b) show that subjects perceive uncertain outcomes differently from delayed outcomes.[3] Taken together, even with the fact that behavior under random termination is consistent with theory, it is still an open question of whether and how time preferences affect behavior in repeated games. To address this challenge, we need to introduce "time" into a framework of repeated game experiments.

Second, many experiments measure individual's time preferences by using delayed payments and find that a non-negligible proportion of subjects are present biased (for instance,

[1]For instance, see Fudenberg and Maskin (1986) and Abreu et al. (1990).

[2]The termination of a repeated game at the end of each stage game is stochastically determined by the continuation probability that is commonly known to subjects.

[3]For theoretical investigations of the relationship between risk and time preferences, see Chakraborty et al. (Forthcoming).

see Andreoni and Sprenger (2012a) and Halevy (2015)). On the theoretical side, there exist continuing efforts to study dynamic inconsistency in dynamic games. Chade et al. (2008) is the first study that introduces a model of quasi-hyperbolic discounting (β-δ) into repeated games.[4] One of their main results is a characterization for the repeated PD game such that with δ (β) fixed, the set of equilibrium outcomes expands as β (δ) increases. As aforementioned, however, the equivalence between random termination and exponential discounting makes it hard to admit any deviation from exponential discounting in the existing experimental framework. Therefore, the second challenge is to have a design that allows us to investigate the effect of present bias on behavior in a repeated game experiment.

The main contribution of this paper is to introduce a novel experimental approach that can accommodate the two challenges above. In the lab, subjects play 20 repeated games at once and one repeated game is randomly selected for payment at the end of the experiment. For the selected repeated game, subjects receive stage game payoffs over a long period of time. There are two treatments regarding the time windows of payments. In the weekly treatment, subjects receive stage game payoffs at weekly intervals—payoffs for the first stage game on the same day of the session, the second stage game payoffs in one week, and so on. In the monthly treatment, subjects receive stage game payoffs at monthly intervals.

Changing the time window of stage game payoffs induces the treatment level variations in subjects' discount factors with which they evaluate the stream of payoffs. That is, subjects in the weekly treatment, on average, are designated to induce higher discount factors than subjects in the monthly treatment in assessing the repeated game payoffs. This enables us to directly examine the causal effects of discount factors on cooperation.

In addition, we have the delay-monthly treatment in which there is one month delay for the initial stage game payoffs. Subjects receive the first stage game payoffs in one month from the day of the session, the second stage game payoffs in two months, and so on. While the time window of payoffs is fixed as monthly intervals, the difference between the monthly and the delay-monthly treatments is that in the delay-monthly treatment, *all payoffs* are made in the future so that present bias is in principle eliminated. Comparing behavior across these two treatments makes it possible to examine the causal effects of present bias on cooperation. To make intertemporal payments credible, all payments are made through Venmo, a mobile app. The advantages of using Venmo for payment are (1) to balance the transaction costs between immediate and future payoffs and (2) to alleviate uncertainty that

[4] Back to Phelps and Pollak (1968) and Laibson (1997), under quasi-hyperbolic discounting, all future rewards, compared to an immediate reward, are discounted further by β, on top of exponential discounting with a factor, δ. If β equals 1, it nests exponential discounting.

subject may have about the future payment.

Our main focus in the analysis is on how subjects behave in the initial stage game. We find unambiguous evidence of the effects of time preferences on cooperation, consistent with the predictions in theory. First, the rate of cooperation is higher in the weekly treatment than in the monthly treatment. This implies that the higher discount factors are, the higher the rate of cooperation. Second, subjects cooperate significantly more in the delay-monthly treatment than do subjects in the monthly treatment. This finding confirms the negative effect of present bias on cooperation. Furthermore, in both comparisons, as subjects gain experience, the differences in cooperation rates across these treatment comparisons become more greater, supporting the argument that learning matters for cooperation in repeated games.

This paper is closely related to the three strands of literature. First, the literature of time preferences has given exclusive attention to individual decision making.[5] Researchers find that time preferences measured by delayed payment are related to important intertemporal behavior of individuals such as financial decision-making (Meier and Sprenger, 2010) and to the demand for commitment devices by individuals who have dynamic inconsistency (Augenblick et al., 2015; Gine et al., 2019). This paper is the first attempt to convey the methodology of delayed payment developed in individual decision making into dynamic strategic interactions. Examining the role of time preferences in dynamic games is particularly important because many instances in the real world involve repeated interactions among multiple players, e.g., bargaining within/between organizations such as households and firms.

Second, the literature of repeated game experiments provide a fundamental test ground with this paper. For instance, Dal Bó (2005) and Dal Bó and Fréchette (2011) show that different random termination probabilities affect cooperative behavior in repeated prisoner's dilemma game experiments in a way that is consistent with theoretical predictions.[6] Based on the random termination framework, various studies look at the relation of personal characteristics to cooperation in repeated game experiments. Davis et al. (2016) measure and relate subjects' discount factors to their behavior in repeated games and find no evidence of robust correlations. Sabater-Grande and Georgantzis (2002), Proto et al. (2019), and Davis et al. (2016) examine whether risk aversion affects cooperation. While Proto et al.

[5]See Frederick et al. (2002) for a survey of the literature. For recent developments in measuring time preferences, see Andersen et al. (2008), Andreoni and Sprenger (2012a), and Halevy (2015).

[6]See Dal Bó and Fréchette (2018) for further discussion of infinitely repeated game experiments in the lab.

(2019) and Davis et al. (2016) do not find a significant relationship between risk aversion and cooperation, risk aversion is negatively correlated with cooperation in Sabater-Grande and Georgantzis (2002), in which subjects are assigned into groups based on their risk aversion. Proto et al. (2019) study the relationship between intelligence and cooperation and find that, only for a high continuation probability, a group of subjects with higher IQ test scores cooperate more than a group of subjects with lower IQ test scores as they gain experience. Dreber et al. (2014) show that there is no relationship between giving behavior in a dictator game and cooperation when cooperation can be supported as an equilibrium outcome. A careful explanation for weak results in the literature is that establishing the correlations between individual characteristics and behaviors may not be obvious in strategic interactions. One reason is due to the fact that beliefs about others' preferences may also matter for behavior in strategic interactions, and in consequence, even finding such correlations does not necessarily imply causal effects. In contrast, the experimental framework developed in this paper enables us to directly test the causal effect of time preferences on behavior in repeated games.

Last but not least, this paper is also related to the recent theoretical developments in dynamic games. There exist small, but growing endeavors of investigating the consequence of dynamic inconsistency in dynamic games. For instance, Obara and Park (2017) study the repeated games in which there is a punishment strategy harsher than a revision to Nash equilibrium. They show that present and future bias lead to different shapes of the worst punishment strategy. Schweighofer-Kodritsch (2018) analyzes a general framework that admits dynamic inconsistency in the Rubinstein bargaining model. He shows how different types of deviations from dynamic consistency are related to (non-) unique equilibrium in bargaining outcomes. The experimental framework introduced in this paper can be applied to such new models to shed light on how different aspects of time preferences affect behavior in various dynamic games.

2 Experimental Design

2.1 Infinitely Repeated Prisoner's Dilemma

An infinitely repeated prisoner's dilemma (PD) game is implemented in the lab by using the random termination method. A repeated game (or a match) may be comprised of a sequence of stage games (or rounds). After each round, with a 75% probability, the match will continue for at least another round, and with a 25% probability, the match will be terminated after that round.[7] Table 1 presents the payoffs (in dollars) for each round. Subjects play 20 matches and are randomly matched with another subject across matches. One of the matches is randomly selected for payment at the end of the experiment.

Table 1: Stage game payoffs

	The other's choice	
Your choice	1	2
1	\$4.00, \$4.00	\$1.00, \$5.00
2	\$5.00, \$1.00	\$2.00, \$2.00

Our novel experimental design, which allows for subjects' time preferences to play an important role in determining their cooperative behavior, is highlighted as follows. First, for the selected match, we pay subjects for their round payoffs over a long period of time, and this makes subjects evaluate a stream of intertemporal payoffs with their time preferences. Second, by exogenously changing the time window of the round payoffs, subjects use different discounting across treatments. Table 2 summarizes our experimental design with three treatments.

Table 2: Treatments information

Treatment	Time window	Delay	Prob. of continuation	# of sessions
Weekly	1 week	No	0.75	4
Monthly	1 month	No	0.75	4
Delay-Monthly	1 month	1 month	0.75	4

In the first treatment, *weekly*, subjects receive their stage game payoffs once per week: the first-stage game payoff right after the session, the second-stage game payoff (if available) one week after the session, and so on. Although all decisions are made at the same time in the lab, the fact that the round payoffs will be given over time outside the lab makes it possible for subjects to perceive their payoffs after round 1 as the payment in the future. This implies that subjects in the *weekly* treatment apply their weekly discounting in assessing the continuation payoffs.

In the second treatment, *monthly*, everything holds the same as in the *weekly* treatment except that the time window for the round payoffs is once per month: the first-stage game payoff occurs right after the session, the second-stage game payoff (if available) one month

[7]These probabilities are commonly known to the subjects.

after the session, and so on. The only difference between the *weekly* and the *monthly* treatments is the extent to which discount factors (weekly vs. monthly) are used in determining subjects' behavior, and the exogenous variation in discount factors enables us to look at the effect of discount factors on cooperation.

The third treatment, *delay-monthly*, employs a front-end delay of one month for all round payoffs that are paid at monthly intervals: the first-stage game payoff occurs one month from the session, the second-stage game payoff (if available) two months from the session, and so on. In the literature on measuring time preferences, the front-end delay is used to eliminate or at least weaken present bias, and one month of delay is long enough for such a purpose. For instance, if subjects exhibit quasi-hyperbolic discounting (Laibson, 1997), any delay completely eliminates present bias.[8] This implies that with no front-end delay in payment, subjects in the *monthly* treatment are more likely to have present bias than subjects in the *delay-monthly* treatment, while their discount factors are controlled for. Therefore, the comparison between the *monthly* and *delay-monthly* treatments allows us to investigate the effect of present bias on cooperation.

It is important to note that random termination under our experimental design is robust to subjects' risk attitude. Sherstyuk et al. (2013) show that under the standard laboratory setting where subjects receive all earnings together at the end of the experiment, the equivalence between random termination and exponential discounting holds only if subjects are risk neutral. However, if payoffs are time separable across stage games like in our experiment, random termination induces the same incentive as theory models regardless of subjects' risk attitude.

2.2 Time Preferences Elicitation

Before subjects play repeated games, we measure subjects' time preferences.[9] Subjects are asked to make decisions for 8 blocks, in each of which there are 2,000 questions (or rows). In each row, subjects are asked to choose between option A (sooner payment) and option B (later payment).

We use the Becker-DeGroot-Marschak (BDM) mechanism in eliciting switch points between sooner and later payments. That is, subjects are asked to state the value at which

[8] In the case of diminishing impatience (Halevy, 2008), present bias will become weaker as a front-end delay becomes longer.

[9] In the experiment, the time preference elicitation task is called phase 1. When making decisions in phase 1, subjects know that there will be a phase 2 but do not know what they will be asked to do in phase 2.

they switch from the option with a fixed amount to the option with varying values. At the value stated and above, all the options with varying amounts are chosen. The 8 blocks differ in timings of sooner and later payments and in whether subjects are asked to switch from option A to option B, or vice versa. In particular, subjects are asked to decide between the following: (1) payment today and payment in 1 week, (2) payment today and payment in 1 month, (3) payment in 1 month and payment in 1 month and 1 week, and (4) payment in 1 month and payment in 2 months. For the 4 blocks, the sooner option pays \$8.00, the later option pays \$0.01 for question 1, and the amount increases in an increment of \$0.01, reaching \$20.00 at question 2,000. For the other 4 blocks, the later option pays \$20.00, and the amount for the sooner option varies between \$0.01 and \$20.00.[10] To make subjects' decisions incentive-compatible, only one question in one of the blocks is randomly selected, and the result is only known to subjects at the end of the experiment. This is to avoid the possibility that a realized outcome from the elicitation task affects behavior in repeated games.

2.3 Payment method

As explained above, subjects in our experiment (possibly) receive several payments over time. To precisely reflect subjects' time preferences on their intertemporal choices, it is extremely important to control for the following issues: (1) the credibility of payments in the future and (2) balancing the transaction costs between immediate and future payments. To minimize such concerns, all earnings (including the show-up fee) from the experiment are paid to subjects through Venmo.[11] Money deposited in a Venmo account can be transferred at no cost to a bank account that is registered to the account, and immediate payments are made right after sessions. It is important for our experimental design that our electronic payment method can guarantee that the payment right after the session would be perceived as "immediate" by our subjects. For instance, Balakrishnan et al. (2017) elicit subjects' time preferences in Kenya by paying them with the Kenyan mobile money system M-Pesa. They find that a substantial portion of subjects exhibit present bias only if they are paid immediately after the session.[12] At the end of the experiment, we collect multiple sources

[10]The instructions and screen shots are available in the Appendix.

[11]Venmo is a mobile payment service provided by PayPal in which account holders can transfer money to others via a mobile application. It handled \$12 billion in transactions in the first quarter of 2018. For more information, visit https://venmo.com/about/product/. At the time of the experiment (fall of 2016), Venmo was widely used among Brown University undergraduates. In the recruiting announcement, it is clearly stated that subjects with a Venmo account are eligible to participate in the experiment.

[12]See Andreoni and Sprenger (2012) and Balakrishnan et al. (2017) for discussions on payment methods in intertemporal decision-making experiments.

of Venmo account information—subject name, username, phone number, and email—so any typo in entering a subject's information cannot cause delayed payments. At the same time, subjects are also given the experimenter's contact information, and they can "request" a payment from the experimenter if any payment is not made on time.[13]

2.4 Procedure

We conducted 12 sessions (4 sessions for each treatment) between September and October 2016. A total of 226 Brown University undergraduate students participated in the experiments, with an average of 18.83 subjects per session, a maximum of 24 and a minimum of 14. The average earning for subjects was \$33.03, with a maximum \$57 and a minimum \$14. The average number of rounds per repeated game was 4.2, with a maximum of 17 and a minimum of 1. We used z-Tree (Fischbacher, 2007) to program our experiment.[14] Since Dal Bó and Fréchette (2018) show in their meta-analysis that the length of a previous match is positively related to cooperative behavior in the next match, we tried to control for such effects as follows. We created 4 different sequences of random numbers in advance of the experiment, and then, the nth sequence of random numbers was applied to the nth session of each treatment, so every nth session of each treatment had exactly the same number of rounds realized.

3 Research Questions

3.1 Discount factor, Present bias, and Cooperation

Our research questions are driven by the theoretical predictions and experimental evidence in infinitely repeated games. Previous theoretical works show that discount factors play an important role in determining whether cooperation can be supported as an equilibrium outcome. For instance, Fudenberg and Maskin (1986) show that individually rational payoffs can be supported in a subgame perfect equilibrium if the discount factor, δ, is sufficiently high. Abreu et al. (1990) show that the set of subgame perfect equilibrium payoffs expands in δ. With exponential discounting, the threshold over which cooperation can be supported as an equilibrium outcome is 0.33, implying that an actual discount factor needs to be at

[13] No such case happened in making the actual payments.

[14] I am grateful to Guillaume Fréchette for sharing the z-Tree code that was used in Yuksel and Fréchette (2016) and Vespa (2015).

least 0.44.[15] Based on previous experiments of measuring time preferences (e.g., Andersen et al. (2008) and Andreoni and Sprenger (2012a)), it is expected that our subjects have weekly and monthly discount factors above the threshold. Furthermore, as shown by Dal Bó and Fréchette (2018), many experiments report that higher continuation probabilities promote greater cooperation, and we anticipate that the same argument still holds true for our experimental design. This leads us to the following question.

Question 1. *Do subjects cooperate more in the weekly treatment than in the monthly treatment?*

While there are many theoretical and experimental endeavors to explore dynamic inconsistency in individual decision-making, relatively limited attention is paid to strategic interactions.[16] Chade et al. (2008) is the first study that introduces a model of quasi-hyperbolic discounting (β-δ) into repeated games. Although, in general, the set of equilibrium payoffs does not have monotonic relationships with δ and β, the authors have a nice characterization for the repeated PD game such that with δ (β) fixed, the set of equilibrium outcomes expands as β (δ) increases.[17] This result implies that present bias may hamper cooperation while discount factors are fixed. The monthly and delay-monthly treatments are deliberately designed to test for such a result. For both treatments, subjects' average discount factors are fixed as monthly given that subjects receive stage game payoffs at monthly intervals. A frond-end delay of one month on the initial stage game payoff, in principle, eliminates present bias in the delay-monthly treatment. Given that there exists ample experimental evidence that a non-negligible proportion of people have present bias, comparing cooperation in the monthly treatment with that in the delay-monthly treatment will shed light on the following question.

Question 2. *Do subjects cooperate more in the delay-monthly treatment than in the monthly treatment?*

3.2 Who is going to cooperate?

The two questions above are the effect of time preferences on cooperation at the treatment level. By using measured subjects' time preferences in the elicitation task, we can also look

[15] That is, an actual discount factor $\times$ 0.75 is at least as high as 0.33.

[16] For a discussion on dynamic inconsistency in individual decision-making, see Andreoni and Sprenger (2012a) and Halevy (2015).

[17] More precisely, this result applies to the class of games in which the minimax point of the stage game coincides with a Nash equilibrium.

at the correlations between time preferences and cooperation at the individual level. It is intuitive to expect that subjects with higher discount factors/less present bias are more likely to cooperate.However, strategic uncertainty may make it difficult to draw clear predictions. In other words, depending on the beliefs that subjects have about their opponents' time preferences and resulting behavior, it is possible that there may not be robust correlations between measured time preferences and cooperative behavior. Therefore, the following is an open question that can be empirically answered.

Question 3. *Are subjects' time preferences in the elicitation task correlated with cooperation?*

4 Results

4.1 Discount factors and cooperation

To study the effect of discount factors on cooperation, we compare subjects' behavior in the weekly treatment with that in the monthly treatment. Figure 1 shows the evolution of the rate of cooperation over matches in each treatment. Note that our primary focus is on looking at the first round cooperation since different matches may end up with a different number of rounds, and the rate of cooperation may depend on the number of rounds.

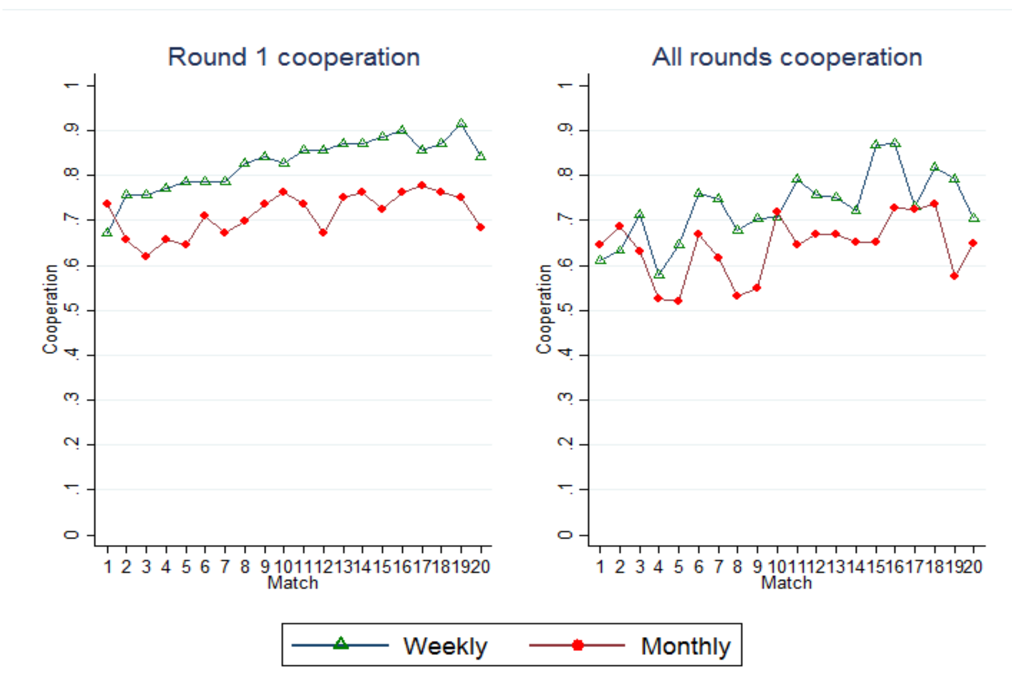

Figure 1: Average cooperation of the Weekly and Monthly treatments

The left panel of Figure 1 represents the rate of round 1 cooperation in each match. It seems that the effects of time widows on cooperation tend to be clear as subjects gain experience. In the first match, the rate of cooperation in the monthly treatment is higher than that in the weekly treatment. However, at the onset of the second match, the rate of cooperation in the weekly treatment becomes higher than that in the monthly treatment, and the difference in cooperation between the two treatments seem to become more persistent in the later matches. The right panel of Figure 1 represents the rate of cooperation for all rounds, and a similar pattern appears in a less clear manner.

Table 3: Percentage of Cooperation in Weekly and Monthly Treatments

First 10 matches					
Round 1			All rounds		
Weekly		Monthly	Weekly		Monthly
0.78	>**	0.69	0.67	>	0.60
Last 10 matches					
Round 1			All rounds		
Weekly		Monthly	Weekly		Monthly
0.87	>***	0.74	0.76	>*	0.66

Notes:
***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Table 3 supports the pattern above by statistically assessing the differences between the two treatments.[18] The rates of round 1 cooperation for the first 10 matches are significantly different. As subjects gain experience, not only the significance level but also the difference in the rate of round 1 cooperation become more clearly differentiated between the weekly and monthly treatments. As a result, round 1 cooperation in all matches in the weekly treatment remains significantly higher than that in the monthly treatment.[19] Resonating with the right panel of Figure 1, the difference in cooperation in all rounds is only marginally significant for the last 10 matches (p-value=0.099).

[18] Unless specified, statistical significance throughout the paper is assessed by probit regressions with a binary variable indicating one of the two relevant categories. Standard errors are clustered at the level of the session.

[19] Statistical differences do not hinge on the selection of specific matches in later matches. For instance, for the last five matches, the rate of cooperation in the weekly treatment is higher than the in the monthly treatment with the p-value of 0.013.

This result unambiguously suggests that the rate of cooperation is higher in the treatment in which the time window for payment is shorter, i.e., discount factors are higher. The differences in cooperative behavior across the treatments grow more significant and differentiated as subjects gain experience in learning how to play repeated games.

4.2 Present bias and cooperation

Next, we examine the effect of present bias on cooperation by comparing behavior in the delay-monthly and monthly treatments. Figure 2 shows the rate of cooperation over matches in both treatments. The left panel of Figure 2 considers round 1 cooperation; in the early matches, round 1 cooperation is slightly higher in the delay-monthly treatment than in the monthly treatment. As subjects get more experienced, they show clearly different cooperative behavior. After the decrease in cooperation for the first 4 matches, it is clear that the difference in cooperation between the two treatments becomes clear as the trend of cooperation in the delay-monthly treatment is increasing, while the rate of cooperation in the monthly treatment fluctuates around 70 percent. The right panel of Figure 2 shows the rate of cooperation for all rounds, with the general pattern being similar to the comparison of round 1 cooperation between those treatments.

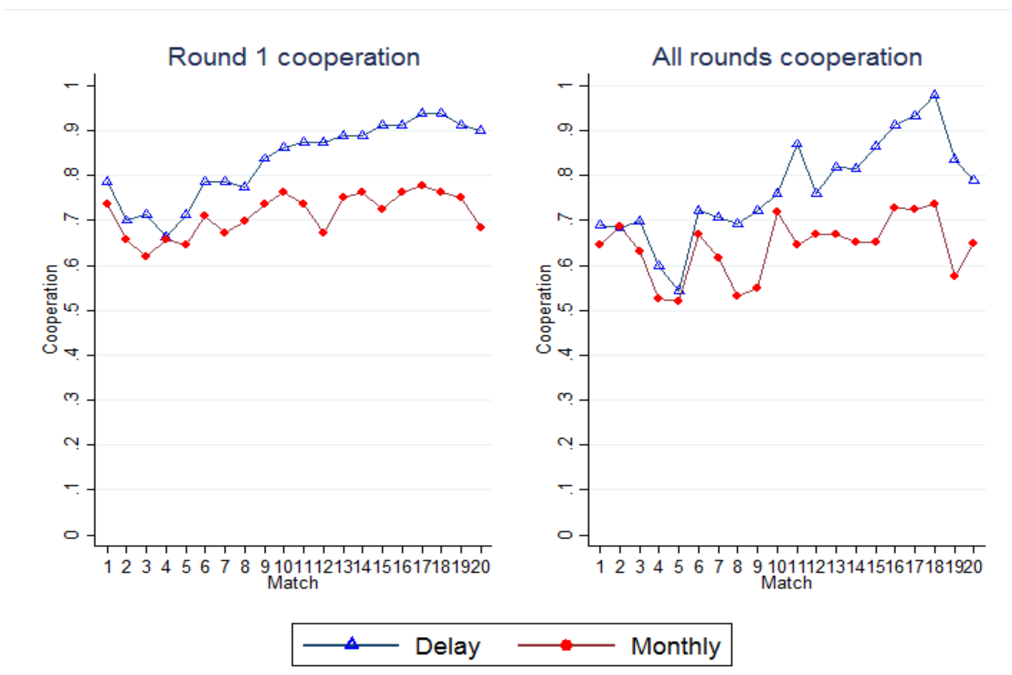

Figure 2: Average cooperation of the Delay and Monthly treatments

The statistical significance of these differences is assessed in Table 4. The effect of learning is more prominent for comparison between the monthly and delay-monthly treatment. The difference in round 1 cooperation in the first 10 matches is not significant. However, as subjects gain experience, round 1 cooperation in the last 10 matches is significantly higher in the delay-monthly treatment than in the monthly treatment. A similar dynamics also applies to the comparison of cooperation in all rounds. Taken together, we can conclude that present bias has a negative effect on cooperation, and this finding also lends support to the theoretical prediction.

Table 4: Percentage of Cooperation in Delay and Monthly Treatments

First 10 matches					
Round 1			All rounds		
Delay		Monthly	Delay		Monthly
0.76	>	0.69	0.68	>	0.60

Last 10 matches					
Round 1			All rounds		
Delay		Monthly	Delay		Monthly
0.90	>**	0.74	0.85	>**	0.66

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

4.3 Are individual time preferences related to cooperation?

In this section, we turn our attention to subjects' time preferences measured from the elicitation task and their relationship to behavior in a repeated game. Based on a model of quasi-hyperbolic discounting (Laibson, 1997), we calculate subjects' β (present/future bias) and δ (discount factor) with some simplified assumptions. First, risk neutrality is presumed as the monetary stakes are small (maximum \$20). Second, we also assume that subjects are narrowly bracketed in the sense that they do not take into account the arbitrage opportunities outside the lab. We calculate δ from the decisions in blocks without immediate payment, e.g., payment in 1 month vs. payment in 1 month and 1 week. Then, we compute $\beta \times \delta$ from the decisions between immediate and future payments, e.g., payment today and payment in 1 week. Using weekly and monthly time windows, we have four parameters, δ_i and β_i, $i \in \{w, m\}$, where w and m are weekly and monthly, respectively. In calculating the parameters, we exclude subjects who do not seem to understand the tasks. For instance, some subjects make extreme choices and their computed discount factors are either 800 or 0.01. The exclusion of such subjects leaves 204 subjects (90.3%).

Table 5: Average Elicited Time Preferences

	Treatments		
	Weekly	Monthly	Delay-Monthly
β_w	0.99 (0.24)	0.98 (0.18)	0.97 (0.20)
δ_w	0.88 (0.15)	0.93 (0.13)	0.90 (0.13)
β_m	1.05 (0.22)	0.96 (0.17)	1.00 (0.18)
δ_m	0.84 (0.17)	0.91 (0.13)	0.87 (0.15)
# Obs.	63	66	75

*Note: Numbers in the parentheses are standard deviations.

We first look at the distribution of β and δ across the treatments. Table 5 shows average weekly and monthly β and δ for each treatment. On average, subjects are exponential discounters and this finding is similar to that of Andreoni and Sprenger (2012a). At the individual level, however, a substantial proportion of subjects has present bias. As shown in Figure A.1 and A.2 in the Appendix, about 30% of subjects across the treatments turn out to have present bias. We test whether subjects are randomly assigned to the treatments regarding their underlying time preferences. For weekly and monthly β, no pairwise comparison between the treatments shows any significant difference.[20] Subjects in each treatment have similar weekly discount factors (KS test, p-value>0.173). However, for discount factors that are elicited from monthly time horizons, subjects in the monthly treatment are significantly more patient than subjects in the weekly and delay-monthly treatments (p-value=0.019 and 0.042, respectively). This result is important for interpreting our main results above. Note that the cooperation rate in the monthly treatment is lower than those in the weekly and delay-monthly treatments. The fact that subjects in the monthly treatment are more patient than subjects in other treatments provides us with more conservative condition (than random assignment) to have such results in repeated games. In other words, our main results do not hinge on the unbalanced distribution of under time preferences of subjects.[21]

[20]With Kolmogorov-Smirnov (KS) test, while other pairwise comparisons are insignificant (p-values>0.416), the difference between the weekly and monthly treatments is only marginally significant (p-value=0.108).

[21]We also check whether the raw choice data does not differ over the treatments. We compare switching

An interesting observation is that for each treatment, the elicited δ over weekly and monthly time horizons do not differ significantly (KS test, p-values>0.545). This finding is consistent with the subadditivity of elicited time preferences that subjects often become more patient when the time horizon gets longer (Read, 2001). Also, the elicited β over weekly and monthly time horizons are very similar (KS test, p-values>0.294). Therefore, when we investigate the relationship between elicited time preferences and cooperation below, we take the average of weekly and monthly estimates for each of β and δ.

Table 6: β, δ, and Round 1 Cooperation (Probit - Marginal effects)

	Weekly		Monthly		Delay-Monthly	
	(1) First 10 Matches	(2) Last 10 Matches	(3) First 10 Matches	(4) Last 10 Matches	(5) First 10 Matches	(6) Last 10 Matches
β	-0.199	-0.109	0.275	0.364	0.449*	0.172
	(0.243)	(0.227)	(0.403)	(0.328)	(0.260)	(0.117)
δ	0.099	0.334	-0.109	0.476	0.273	0.001
	(0.137)	(0.232)	(0.380)	(0.498)	(0.264)	(0.316)
Obs.	630	630	660	660	750	750

Notes: Dependent variable: cooperation=1, defection=0. Clustered standard errors in parentheses. Marginal effects are taken at the mean.
***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Table 6 presents the marginal effects from probit regressions. The correlations between measured time preferences and round 1 cooperation are examined for each treatment. Overall, we find no robust and significant correlations. One common feature for the treatments with no delay (weekly and monthly) is that as subjects gain experience, the magnitudes of correlation coefficients are more likely to align with cooperative behavior. For instance, except for β in the weekly treatment, β and δ are more positively correlated with cooperation in later matches of the experiment, although all coefficients are not significant. For β in the weekly treatment, its negative correlation gets smaller in the last 10 matches. In the delay-monthly treatment, while subjects with higher β and δ are more likely to cooperate, the magnitudes of coefficients decline as subjects gain experience.

The lack of significant correlations in the analysis above is possibly due to the fact that in contrast to individual decision making, behavior in strategic interactions can be governed by other determinants besides a player's own time preferences. For instance, even if a subject is very patient, a pessimistic belief about the opponent's cooperation may lead her to

points in all 8 blocks. Among all pairwise comparisons between the treatments, only 2 out of 24 comparisons shows some significant differences (KS test, p-values$=0.017$ and 0.068).

defect. Therefore, we try to control for subjects' belief about the opponent as follows. Given that discount factors need to be higher than 0.44 to support cooperation as an equilibrium outcome, most of our subjects facing the weekly or monthly time windows are more patient than the threshold. If subjects defect in all matches, it might be driven by their pessimistic beliefs about the opponent, rather than their impatience. That said, we focus on subjects whose (1) behaviors are conditional on positive beliefs and (2) beliefs are not altered drastically, so that variations in cooperation can be explained by differences in time preferences. One way to approximate such subjects is to look at who (1) cooperated and (2) observed the opponent's cooperation in the previous match. The former is to make sure that subjects had positive beliefs about the opponent, which made them cooperate, and the latter is to prevent subjects' beliefs from negatively affected by the opponent's defection.

Table 7: β, δ, and Round 1 Cooperation in last 10 Matches (Probit - Marginal effects)

	Mutually Cooperated in Round 1 of the previous Match			Mutually Cooperated in the previous Round		
	Weekly	Monthly	Delay-Monthly	Weekly	Monthly	Delay-Monthly
	(1)	(2)	(3)	(4)	(5)	(6)
β	0.004	0.049	0.046**	0.043	0.234**	0.136
	(0.010)	(0.061)	(0.020)	(0.040)	(0.100)	(0.095)
δ	0.082**	0.079	0.061***	0.068**	0.097	0.045
	(0.035)	(0.091)	(0.010)	(0.033)	(0.140)	(0.091)
Obs.	484	355	616	440	352	592

Notes: Dependent variable: cooperation=1, defection=0. Clustered standard errors in parentheses. Marginal effects are taken at the mean.
***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Table 7 looks at the relationship between elicited time preferences and round 1 cooperation again. Column (1)-(3) focus on subjects who mutually cooperated in round 1 of the previous Match. Not only we have all positive coefficients, but also β and δ become meaningful predictors for cooperation in round 1 of the next match. In particular, 1% increase in δ of the weekly treatment is related to 0.08% increase in cooperation among subjects who mutually cooperated in round 1 of the previous match. Both β and δ in the delay-monthly treatment are also scientifically correlated with round 1 cooperation. To check the robustness of the results, in column (4)-(6), we look subjects who mutually cooperated in the last round of the previous match. This is because subjects may not precisely recall behavior in round 1 of the previous match if the length of the previous match gets longer. We also find similar pictures that β and δ are positively correlated with round 1 cooperation in all columns. While no parameter in the delay-monthly treatment is no longer significant, instead, 1% increase in β in the monthly treatment scientifically leads to 0.23% increase in round 1 cooperation.

Taken together, our result, consistent with those of previous studies, implies that investigating correlations between individual characteristics and cooperation in strategic interactions is not trivial.[22] Once we try to control for beliefs about the opponent, elicited time preferences are positively and significantly related to cooperation in many circumstances.

4.4 Does "time" matter for the play of a repeated game?

The innovation of this study is to incorporate "time" into the random termination method for implementing repeated games in the lab. One question to explore is whether time makes subjects play repeated games differently. We provide anecdotal evidence that compares our data with that of the previous study. In particular, our main focus is on the selection of strategies.

We estimate strategies used in the last ten matches following the strategy frequency estimation method (SFEM) in Dal Bó and Fréchette (2011). The SFEM estimates a mixture model in which the frequency of each strategy from a pre-specified set of strategies is measured, assuming that each subject uses the same (mixed) strategies in every repeated game with a possibility of mistakes. In this estimation, we assume 6 strategies that are frequently used in the literature: always defect (AD), always cooperate (AC), grim (G), tit for tat (TFT), win-stay, lose-shift (WSLS), and a trigger strategy with two periods of punishment (T2). WSLS is a strategy that begins with cooperation and then depends on the combination of behavior chosen in the previous round. If both cooperate or defect, then cooperation will be selected. Otherwise, WSLS will defect. T2 begins with cooperation, and if the other defects, then T2 triggers two rounds of defection. After the punishment phase, T2 goes back to cooperation.[23]

Given that there is no previous experiment that used the exactly same payoff parameters with ours, we compare our study with the previous experiment based on the size of the basin of attraction of AD against G strategy. The basin of attraction is the maximum proportion of G that playing AD can be optimal, and Dal Bó and Fréchette (2018) show that the size of

[22]See Davis et al. (2016) for time preferences, Sabater-Grande and Georgantzis (2002), Proto et al. (2019), and Davis et al. (2016) for risk preferences, and Dreber et al. (2014) for social preferences in repeated games..

[23]AD, AC, G, and TFT are the strategies that are most frequently identified in Dal Bó and Fréchette (2011). Dal Bó and Fréchette (2019) show that these strategies are robustly identified if T2 is replaced with another strategy such as STFT, which is equivalent to TFT, except that it defects in round 1. See Dal Bó and Fréchette (2011) for the estimation procedure in detail.

the basin of attraction well captures cooperative behavior in repeated game experiments.[24] Without considering the "actual" discount factors, the size of the basin of attraction for our experiment is 0.1667. With the estimated discount factors in table 6, the basin of attractions by replacing the continuation probability with 0.75× discount factors are between 0.2326 and 0.2663. Among previous experiments reviewed by Dal Bó and Fréchette (2018), the two treatments from Dal Bó and Fréchette (2011) provide us with the closest basin of attractions. In Dal Bó and Fréchette (2011), the payoffs for CC, CD, DC, and DD are R, 12, 50, and 25, respectively, and the treatments differ in R. The basin of attraction for the R=48 treatment is 0.1625, and the comparison between our experiment and the R=48 treatment allows us to see whether introducing actual time preferences into random termination makes subjects less cooperative. The R=40 treatment has the basin of attraction of 0.2708, and we can check whether similar basin of attractions lead to coherent levels of cooperation across the experiments.

Table 8: Estimation of Strategies Used by Treatment

	Weekly	Monthly	Delay	DB & F (R=48)	DB & F (R=40)
Round 1 Cooperation	0.83	0.71	0.83	0.85	0.61
AD	0.087*	0.131*	0.027	0.000	0.109
	(0.050)	(0.074)	(0.026)	(0.000)	(0.096)
AC	0.097	0.124	0.107	0.079	0.296**
	(0.065)	(0.098)	(0.092)	(0.085)	(0.123)
GRIM	0.244***	0.268**	0.287	0.116	0.267
	(0.149)	(0.103)	(0.181)	(0.195)	(0.202)
TFT	0.572**	0.477***	0.515***	0.561***	0.327*
	(0.113)	(0.145)	(0.187)	(0.185)	(0.186)
WSLS	0.000	0.000	0.000	0.000	0.000
	(0.000)	(0.001)	(0.000)	(0.000)	(0.000)
T2	0.000	0.000	0.064	0.244	0.000
Gamma	0.368***	0.467***	0.329***	0.287***	0.435***
	(0.048)	(0.049)	(0.052)	(0.061)	(0.126)

Notes: Bootstrapped standard errors in parentheses.
***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Table 8 presents the estimates of the frequency of each strategy.[25] The first row is the

[24]Blonski et al. (2011) also propose an index that reveals how risk dominance determines cooperation in the play of repeated games.

[25]The frequency of T2 is computed by the fact that the frequencies for all strategies sum to one, and gamma captures the amount of noise, with infinite gamma implying that behavior would be purely random. Our estimation is for the last 10 matches and the estimation from Dal Bó and Fréchette (2011) is for matches

rate of round 1 cooperation for all matches. First, the cooperation rates in our experiment lie below that in the R=48 treatment. Having actual time preferences on top of random termination induces subjects less patient and lowers cooperation as it is expected. Second, the comparison between our experiment and the R=40 treatment is surprising in that although the underlying basin of attractions are very similar, the rate of cooperation is much higher in our experiment. This indicates that the extent to which actual time preferences affect play of a repeated game might be different from the way that variations in payoffs (R=48→40) reduces cooperation.

Looking at the estimated strategies gives us clearer intuition into such differences. First, there are not much variations over the three treatments of our experiment. A slight difference is that AD is the highest and TFT is the lowest in the monthly treatment where the rate of cooperation is the lowest. A substantial proportion of subjects use cooperative strategies. In particular, on average, more than 75 percent of the data in all treatments can be identified as either G or TFT. Second, in comparing the weekly and delay-monthly treatments with the R=48 treatment, the proportion of subjects who use TFT is similar, but almost no one uses AD and the frequency of G is much lower in the R=48 treatment. As a result, T2 becomes more popular in the R=48 treatment. This result implies that even with similar levels of cooperation, the subtle mechanisms to support cooperation depend on underlying time preferences.[26] Third, a big drop in cooperation between our experiment and the R=40 treatment is mainly due to the fact that less subjects use TFT in the R=40 treatment. Although subjects in the R=40 treatment are more likely to use AC than subjects in the monthly treatment, clear reduction in frequency of TFT in the R=40 treatment hampers cooperation. This result clearly shows that while the basin of attractions are similar, the choice of strategies is influenced by whether actual time preferences kick in subjects' decision making. In sum, the anecdotal evidence provided here sheds light on the finding that behavior and underlying mechanisms to support cooperation are shaped by the presence of time preferences beyond the scope of random termination.

after 110 interactions.

[26]Some recent theoretical works also show that time preferences matter for the shape of strategies in dynamic games. For instance, Obara and Park (2017) show that in repeated games with the punishment strategy that is more severe than the grim trigger strategy (e.g., repeated Cournot duopoly), present bias affects the pattern of the worst punishment.

5 Discussion

In the theoretical analysis of a dynamic game with players who have dynamic inconsistency (e.g., present bias), it is important to take into account conflicts between the current self and the future self—the incentives of the current self for the future outcome do not coincide with those of the future self. Our experimental design makes it difficult to address such dynamic inconsistency because all decisions are made at once. In other words, accommodating dynamic inconsistency in a dynamic game needs a longitudinal design in which the same subjects make decisions over time.

For the reason above, the experimental design of Kim (2017) is complementary to the design of the current study. In Amazon's Mechanical Turk, subjects play one repeated prisoner's dilemma game over weeks—one stage game each week.[27] Although this can be considered the most intuitive design to mimic the environments modeled in theory, some limitations exist. First, the attrition of subjects is extremely difficult to avoid.[28] Therefore, careful considerations that take into account the effect of attrition on behavior are needed. Second, as pointed out in Dal Bó and Fréchette (2018), subjects' experience with repeated games plays an important role in determining their behavior. Playing one repeated game is not enough for subjects to fully learn about the game. With a longitudinal design in which one stage game is played once a week or a month, running multiple repeated games will not be easy to implement. Taken together, the experimental design in this paper is appropriate to study the causal effects of time preferences on behavior in a dynamic game.

6 Conclusions

In this paper we implement a novel experimental design for repeated games in the laboratory. Subjects play all repeated games at once in the lab, but stage game payoffs are paid to them over a long period of time. Varying the time window for payment (weekly or monthly) allows us to investigate the effects of patience on cooperation. Subjects with the weekly payment cooperate more than subjects with the monthly payment, and this confirms that higher discount factors promote greater cooperation. We also introduce one month delay

[27]Subjects also receive stage game payoffs every week, and all payment are electronically made through Mechanical Turk's payment system.

[28]To prevent attrition, Kim (2017) uses a stick-and-carrot strategy. If subjects participate in all weeks of the experiment, they receive a bonus pay of $3 in addition to their earnings from a repeated game. Once subjects miss one stage game, they are not allowed to participate in the remaining weeks of the experiment. The average attrition rate between the first and the second week is about 15%, and it drops to lower than 5% after the second week.

to the initial stage game payoffs to study the effect of present bias on cooperation. By comparing the monthly payment treatments with and without a delay, we find that the rate of cooperation is higher when there is a delay. This implies that present bias reduces cooperation.

In the literature on time preferences, exclusive attention has been paid to individual intertemporal decision-making—measuring time preferences and their relationship with behavior such as financial decision-making and demand for a commitment device. We extend the scope of the literature in two ways. First, we apply the experimental method of delayed payments to a strategic interaction, which is a more complicated environment than individual decision-making. Our results clearly show that subjects are able to understand the consequences of time preferences on interactions between two players. In particular, given that it is an important question of whether people are sophisticated enough to understand their dynamic inconsistency, the comparison between the treatments with and without a delay provides strong evidence on the sophistication of our subjects. Second, while the previous studies look at correlations between measured time preferences and behavior, our design allows us to assess the causal effects of time preferences on cooperation. This implies that our results lend support to previous correlation results found in individual decision-making.

The experimental framework introduced in this paper can be applied to the recent developments in dynamic games that accommodate dynamic inconsistency. For instance, Kim et al. (2020) directly study the role of time preferences in the Rubinstein bargaining of alternating offers, based on the predictions of Schweighofer-Kodritsch (2018). On one hand, a bargainer with the weekly time window for the agreed payoff play with another bargainer with the monthly time window. On the other hand, only one bargainer experiences a delay for the initial stage game payment, while another bargainer has no delay. They show that discount factors and present bias are the essential determinants for bargaining outcomes even in an environment where there exist clear asymmetries among bargainers' time preferences. The same idea also has interesting applications for multilateral bargaining, dynamic games with Markov strategies, and Cournot duopoly games where the worst punishment strategy is harder than permanent recursion to static Nash equilibrium.

This paper shows that changing time windows for payment is effective in modifying subjects' underlying time preferences and inducing different cooperative outcomes. This approach has meaningful implications for individual decision making contexts as well. For instance, can the introduction of the front end delay for payment provide people with better incentives for work? How frequently do we need to pay workers for their best outcome? I hope this paper can help better understanding of dynamic incentives and the role of time preferences in individual and strategic decision making over time.

References

[1] Abreu, Dilip, David Pearce, and Ennio Stacchetti (1990). “Toward a theory of discounted repeated games with imperfect monitoring.” *Econometrica*: 1041-1063.

[2] Andreoni, James, and Charles Sprenger (2012a). “Estimating time preferences from convex budgets.” *American Economic Review*, 102(7): 3333-3356.

[3] Andreoni, James, and Charles Sprenger (2012b). “Risk preferences are not time preferences.” *American Economic Review*, 102(7): 3357-3376.

[4] Andersen, Steffen, Glenn W. Harrison, Morten I. Lau, and E. Elisabet Rutström (2008). “Eliciting risk and time preferences.” *Econometrica*, 76(3): 583-618.

[5] Augenblick, Ned, Muriel Niederle, and Charles Sprenger (2015). “Working over time: Dynamic inconsistency in real effort tasks.” *Quarterly Journal of Economics*, 130(3): 1067-1115.

[6] 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): 164-192.

[7] Chakraborty, Anujit, Yoram Halevy, and Kota Saito (Forthcoming). “The Relation between Behavior under Risk and over Time.” *American Economic Review: Insights*.

[8] Chade, Hector, Pavlo Prokopovych, and Lones Smith (2008). “Repeated games with present-biased preferences.” *Journal of Economic Theory*, 139(1): 157-175.

[9] Dal Bó, Pedro (2005). “Cooperation under the shadow of the future: experimental evidence from infinitely repeated games.” *American Economic Review*, 95(5): 1591-1604.

[10] Dal Bó, Pedro and Guillaume R. Fréchette (2011). “The evolution of cooperation in infinitely repeated games: Experimental evidence.” *American Economic Review*, 101(1): 411-429.

[11] Dal Bó, Pedro and Guillaume R. Fréchette (2018). “On the determinants of cooperation in infinitely repeated games: A survey.” *Journal of Economic Literature*, 51(1): 60-114.

[12] Dal Bó, Pedro and Guillaume R. Fréchette (2019). “Strategy choice in the infinitely repeated prisoners’ dilemma.” *American Economic Review*, 109(11): 3929-3952.

[13] Davis, Douglas, Asen Ivanov, and Oleg Korenok (2016). “Individual characteristics and behavior in repeated games: an experimental study.” *Experimental Economics*, 19(1): 67-99.

[14] Dreber, Anna, Drew Fudenberg, and David G. Rand (2014). “Who cooperates in repeated games: The role of altruism, inequity aversion, and demographics.” *Journal of Economic Behavior and Organization*, 98: 41-55.

[15] Fischbacher, Urs (2007). “z-Tree: Zurich toolbox for ready-made economic experiments.” *Experimental Economics*, 10(2): 171-178.

[16] Fréchette, Guillaume R., and Sevgi Yuksel (2017). ”Infinitely repeated games in the laboratory: Four perspectives on discounting and random termination.” *Experimental Economics*, 20(2): 279-308.

[17] Frederick, Shane, George Loewenstein, and Ted O’donoghue (2002). “Time discounting and time preference: A critical review.” *Journal of Economic Literature*, 40(2): 351-401.

[18] Fudenberg, Drew, and Eric Maskin (1986). “The folk theorem in repeated games with discounting or with incomplete information.” *Econometrica*: 533-554.

[19] Giné, Xavier, Jessica Goldberg, Dan Silverman, and Dean Yang (2018). “Revising commitments: Field evidence on the adjustment of prior choices.” *Economic Journal*, 128 (608): 159-188.

[20] Halevy, Yoram (2015). “Time consistency: Stationarity and time invariance.” *Econometrica*, 83(1): 335-352.

[21] Kim, Jeongbin (2017). “Discounting, dynamic consistency, and cooperation in an infinitely repeated game experiment.” *Working paper.*

[22] Kim, Jeongbin , Wooyoung Lim, and Sebastian Schweighofer-Kodritsch (2020). “Bargaining and time preferences: An experimental study.” *Working paper.*

[23] Laibson, David (1997). “Golden eggs and hyperbolic discounting.” *Quarterly Journal of Economics*: 443-477.

[24] Lehrer, Ehud, and Ady Pauzner (1999). “Repeated games with differential time preferences.” *Econometrica*, 67(2): 393-412.

[25] Meier, Stephan, and Charles Sprenger (2010). "Present-biased preferences and credit card borrowing." *American Economic Journal: Applied Economics*, 2(1): 193-210.

[26] Obara, Ichiro, and Jaeok Park (2017). "Repeated games with general discounting." *Journal of Economic Theory*, 172: 348-375.

[27] Phelps, Edmund S., and Robert A. Pollak (1968). "On second-best national saving and game-equilibrium growth." *Review of Economic Studies*, 35(2): 185-199.

[28] Proto, Eugenio, Aldo Rustichini, and Andis Sofianos (2019). "Intelligence, Personality, and Gains from Cooperation in Repeated Interactions." *Journal of Political Economy*, 127(3): 1351-1390.

[29] Read, Daniel (2001). "Is time-discounting hyperbolic or subadditive?" *Journal of risk and uncertainty* , 23(1): 5-32.

[30] Roth, Alvin E., and J. Keith Murnighan (1978). "Equilibrium behavior and repeated play of the prisoner's dilemma." *Journal of Mathematical Psychology*, 17(2): 189-198.

[31] Sabater-Grande, Gerardo, and Nikolaos Georgantzis (2002). "Accounting for risk aversion in repeated prisoners' dilemma games: An experimental test." *Journal of Economic Behavior and Organization*, 48(1): 37-50.

[32] Schweighofer-Kodritsch, Sebastian (2018). "Time preferences and bargaining." *Econometrica*, 86(1): 173-217.

[33] Sherstyuk, Katerina, Nori Tarui, and Tatsuyoshi Saijo (2013). "Payment schemes in infinite-horizon experimental games." *Experimental Economics*, 16(1): 125-153.

[34] Sutter, Matthias, Martin G. Kocher, Daniela Glätzle-Rützler, and Stefan T. Trautmann (2013). "Impatience and uncertainty: Experimental decisions predict adolescents' field behavior." *American Economic Review*, 103(1): 510-531.

[35] Vespa, Emanuel (2015). "An experimental investigation of strategies in dynamic games." *Working paper.*

Appendices

A Distribution of elicited time preferences

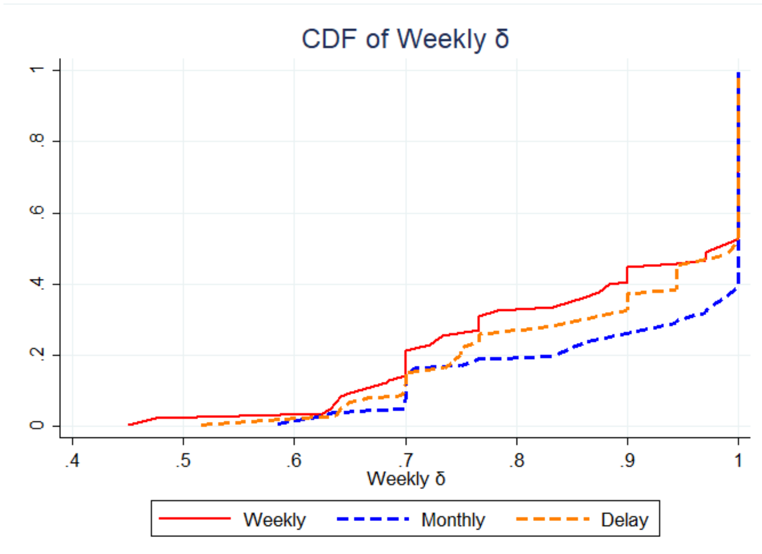

(a) Weekly δ

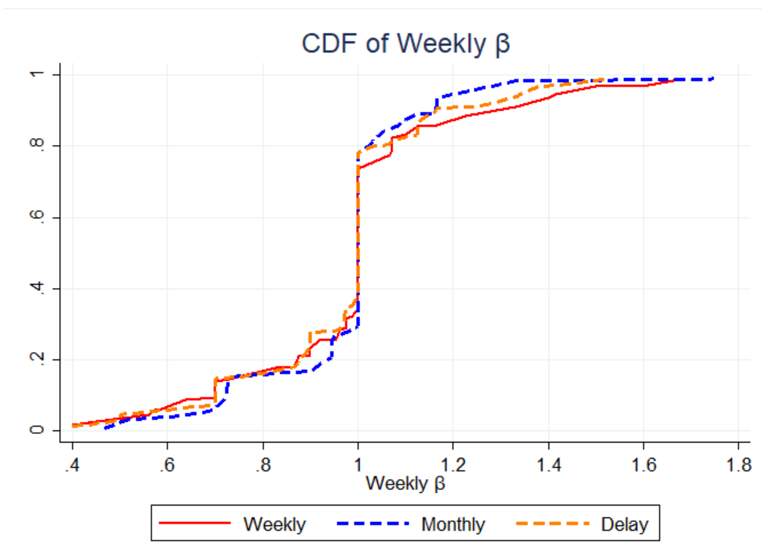

(b) Weekly β

Figure A.1: Distribution of Weekly Parameters

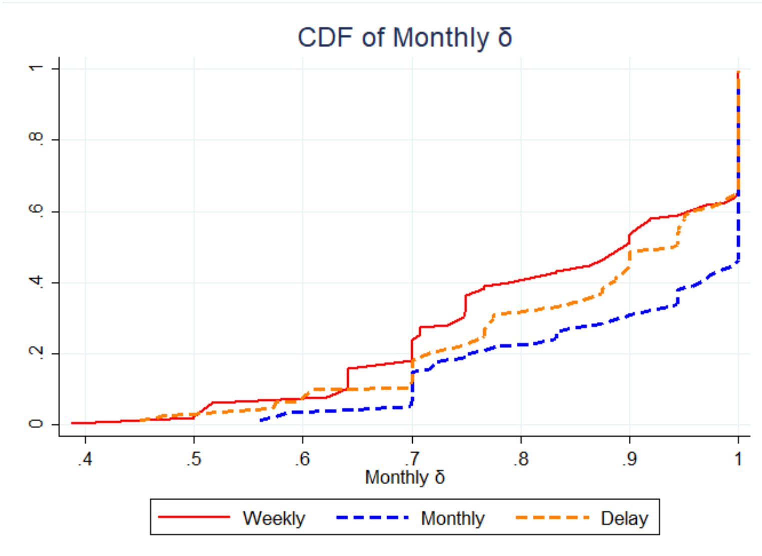

(a) Monthly δ

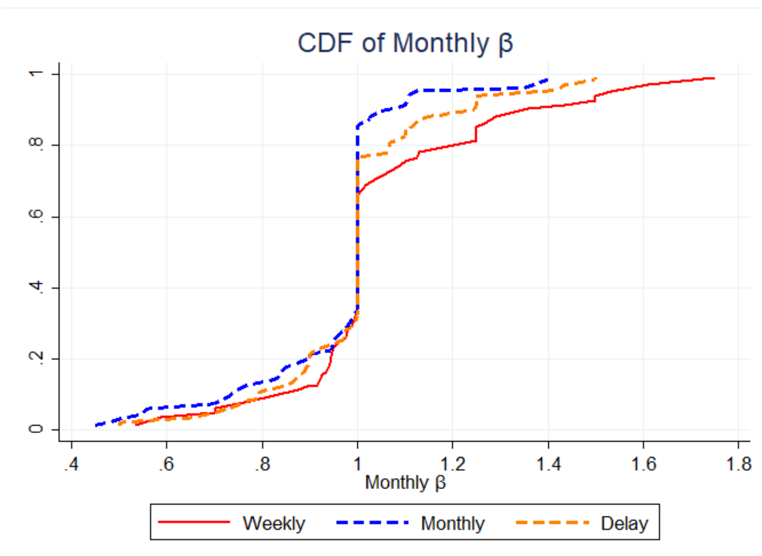

(b) Monthly β

Figure A.2: Distribution of Monthly Parameters

B Instructions for the experiment

B.1 Phase 1 instructions

Instructions (phase 1)

Welcome

You are about to participate in a session on decision-making, and you will be paid for your participation. What you earn depends partly on your decisions and partly on chance. The payment you earn will be paid to you through VENMO.

The entire session will take place through computer terminals. Please do not talk or try to communicate in any way with other participants during the session.

The entire session consists of two phases. The instructions for phase 1 are given below. After phase 1 ends, you will be given the instructions for phase 2.

We will start with a brief instruction period for phase 1. During this instruction period, you will be given a description of the main features of phase 1. If you have any questions during this period, raise your hand. Your question will then be answered publicly so everyone can hear.

General Instructions

1. In phase 1, you will be asked to make decisions for 8 blocks of questions. In each block, there are 2,000 questions. For each question, you can choose either: Option A, which pays you sooner, or Option B, which pays you later.

2. After you answer all questions, I will randomly pick one question and pay you the option you chose on that question. Each question is equally likely to be chosen for payment. Obviously, you have no incentive to lie on any question, because if that question gets chosen for payment, then you would end up with the option you like less.

3. For example, the questions in one block are as follows (note that each row corresponds to a question, and so you will have to choose an option in each row):

Questions	Payment Option A (Pays the Amount Below Today)	Payment Option B (Pays the Amount Below in 1 month)
1	$8.00	$0.01
2	$8.00	$0.02
3	$8.00	$0.03
⋮	⋮	⋮
1,999	$8.00	$19.99
2,000	$8.00	$20.00

I assume you will choose Option A for at least the first few questions, but at some point switch to choosing Option B. In order to save time, you can answer at which dollar value you'd switch. I can then 'fill out' your answers to all 2,000 questions based on your switch point (choosing Option A for all questions before your switch point, and Option B for all questions at or after your switch point). I will still draw one question randomly for payment. Again, if you lie about your preferred switch point, you might end up with an option that you like less.

4. The 8 blocks will differ in two ways: (1) the timings of sooner and later payments:
- Between payment **today** and payment in **1 week**.
- Between payment **today** and payment in **1 month**.
- Between payment in **1 month** and payment in **1 month and 1 week**.
- Between payment in **1 month** and payment in **2 months**.

and (2) whether you are asked to switch from A to B or B to A.

Payment

1. At the end of the experiment, one question in one of the blocks will be randomly selected for payment and will be displayed on your screen. Depending on your decision for that question, you will be paid on the designated date through VENMO. If in the question that is randomly selected, your decision was to receive a payment today, then you will be paid through VENMO within a few hours of the end of the experiment. If, on the other hand, your decision was to receive a payment in the future, you will be paid on the designated date through VENMO.

2. In addition, you will receive a $5 show-up fee through VENMO after the experiment.

- Are there any questions?

Before we start, let me remind you that:

- There are 8 blocks of questions in each of which you will be asked to state your switch point.
- Only one question in one of the blocks will be randomly selected for payment.
- Depending on your decision, you will be paid on the designated date through VENMO.
- A $5 show-up fee will be paid to you through VENMO after the experiment.

B.2 Phase 2 instructions

Instructions (phase 2) - Weekly treatment

We will start with a brief instruction period for phase 2. During this instruction period you will be given a description of the main features of phase 2. If you have any questions during this period, raise your hand. Your question will then be answered publicly so that everyone can hear.

General Instructions

1. In phase 2 you will be asked to make decisions in several rounds. Each sequence of rounds is referred to as a match. You will be randomly paired with another person for a match.

2. The length of a match is determined randomly. After each round, there is a 75% probability that the match will continue for at least another round. This is as if we were to randomly choose an integer between 1 and 100 and continue if the number chosen is less than or equal to 75 and end if the number chosen is larger than 75. So, for instance, if you are in round 2, the probability that there will be a third round is 75%, and if you are in round 9, the probability that there will be another round is also 75%.

3. Once a match ends, you will be randomly paired with another person for a new match. You will have 20 matches in phase 2.

4. In each round, you will be asked to choose between action 1 and 2. The payoffs are determined by your action and the action chosen by the person paired with you. The payoffs are described in the table below:

	The other's choice	
Your choice	1	2
1	\$4.00, \$4.00	\$1.00, \$5.00
2	\$5.00, \$1.00	\$2.00, \$2.00

- The first entry in each cell represents your payoff, while the second entry represents the payoff of the person you are paired with. That is, if:

You select 1 and the other selects 1, you each make \$4.00.
You select 1 and the other selects 2, you make \$1.00 while the other makes \$5.00.
You select 2 and the other selects 1, you make \$5.00 while the other makes \$1.00.
You select 2 and the other selects 2, you each make \$2.00.

- Once you and the person you are paired with have made your choices, those choices will be highlighted and your payoff for the round will appear.

Payment

1. At the end of the experiment, one of the matches will be randomly selected for payment.

2. For the selected match, you will receive payment for the first round today. After that, you will receive payment for the following rounds once a week. That is, you will receive payment for the second round in 1 week, payment for the third round in 2 weeks, and so on. The schedule of payment is summarized in the table below.

Payoffs (round)	Payment schedule (from today)
1st round payoff	Today
2nd round payoff	in 1 week
3rd round payoff	in 2 weeks
⋮	⋮

3. In the same way that payments are made for phase 1, you will be paid on the designated dates through VENMO.

- Are there any questions?

Before we start, let me remind you that:

- The length of a match is randomly determined. After each round, there is a 75% probability that the match will continue for at least another round. You will play with the same person for the entire match.

- After a match is finished, you will be randomly paired with another person for a new match. You will have 20 such matches.

- One match will be randomly selected for payment.
You will receive your payment for the first round today. After that, you will receive payment for the following rounds once a week. That is, you will receive payment for the second round in 1 week, payment for the third round in 2 weeks, and so on.

Instructions (phase 2) - Monthly treatment

We will start with a brief instruction period for phase 2. During this instruction period you will be given a description of the main features of phase 2. If you have any questions during this period, raise your hand. Your question will then be answered publicly so that everyone can hear.

General Instructions

1. In phase 2 you will be asked to make decisions in several rounds. Each sequence of rounds is referred to as a match. You will be randomly paired with another person for a match.

2. The length of a match is determined randomly. After each round, there is a 75% probability that the match will continue for at least another round. This is as if we were to randomly choose an integer between 1 and 100 and continue if the number chosen is less than or equal to 75 and end if the number chosen is larger than 75. So, for instance, if you are in round 2, the probability that there will be a third round is 75%, and if you are in round 9, the probability that there will be another round is also 75%.

3. Once a match ends, you will be randomly paired with another person for a new match. You will have 20 matches in phase 2.

4. In each round, you will be asked to choose between action 1 and 2. The payoffs are determined by your action and the action chosen by the person paired with you. The payoffs are described in the table below:

	The other's choice	
Your choice	1	2
1	\$4.00, \$4.00	\$1.00, \$5.00
2	\$5.00, \$1.00	\$2.00, \$2.00

- The first entry in each cell represents your payoff, while the second entry represents the payoff of the person you are paired with. That is, if:

You select 1 and the other selects 1, you each make \$4.00.
You select 1 and the other selects 2, you make \$1.00 while the other makes \$5.00.
You select 2 and the other selects 1, you make \$5.00 while the other makes \$1.00.
You select 2 and the other selects 2, you each make \$2.00.

- Once you and the person you are paired with have made your choices, those choices will be highlighted and your payoff for the round will appear.

Payment
1. At the end of the experiment, one of the matches will be randomly selected for payment.

2. For the selected match, you will receive payment for the first round today. After that, you will receive payment for the following rounds once a month. That is, you will receive payment for the second round in 1 month, payment for the third round in 2 months, and so on. The schedule of payment is summarized in the table below.

Payoffs (round)	Payment schedule (from today)
1st round payoff	Today
2nd round payoff	in 1 month
3rd round payoff	in 2 months
⋮	⋮

3. In the same way that payments are made for phase 1, you will be paid on the designated dates through VENMO.

- Are there any questions?

Before we start, let me remind you that:

- The length of a match is randomly determined. After each round, there is a 75% probability that the match will continue for at least another round. You will play with the same person for the entire match.

- After a match is finished, you will be randomly paired with another person for a new match. You will have 20 such matches.

- One match will be randomly selected for payment.
You will receive your payment for the first round today. After that, you will receive payment for the following rounds once a month. That is, you will receive payment for the second round in 1 month, payment for the third round in 2 months, and so on.

Instructions (phase 2) - Delay-Monthly treatment

We will start with a brief instruction period for phase 2. During this instruction period you will be given a description of the main features of phase 2. If you have any questions during this period, raise your hand. Your question will then be answered publicly so that everyone can hear.

General Instructions

1. In phase 2 you will be asked to make decisions in several rounds. Each sequence of rounds is referred to as a match. You will be randomly paired with another person for a match.

2. The length of a match is determined randomly. After each round, there is a 75% probability that the match will continue for at least another round. This is as if we were to randomly choose an integer between 1 and 100 and continue if the number chosen is less than or equal to 75 and end if the number chosen is larger than 75. So, for instance, if you are in round 2, the probability that there will be a third round is 75%, and if you are in round 9, the probability that there will be another round is also 75%.

3. Once a match ends, you will be randomly paired with another person for a new match. You will have 20 matches in phase 2.

4. In each round, you will be asked to choose between action 1 and 2. The payoffs are determined by your action and the action chosen by the person paired with you. The payoffs are described in the table below:

	The other's choice	
Your choice	1	2
1	\$4.00, \$4.00	\$1.00, \$5.00
2	\$5.00, \$1.00	\$2.00, \$2.00

- The first entry in each cell represents your payoff, while the second entry represents the payoff of the person you are paired with. That is, if:

You select 1 and the other selects 1, you each make \$4.00.
You select 1 and the other selects 2, you make \$1.00 while the other makes \$5.00.
You select 2 and the other selects 1, you make \$5.00 while the other makes \$1.00.
You select 2 and the other selects 2, you each make \$2.00.

- Once you and the person you are paired with have made your choices, those choices will be highlighted and your payoff for the round will appear.

Payment
1. At the end of the experiment, one of the matches will be randomly selected for payment.

2. For the selected match, you will receive payment for the first round in 1 month from today. After that, you will receive payment for the following rounds once a month. That is, you will receive payment for the second round in 2 months, payment for the third round in 3 months, and so on. The schedule of payment is summarized in the table below.

Payoffs (round)	Payment schedule (from today)
1st round payoff	in 1 month
2nd round payoff	in 2 months
3rd round payoff	in 3 months
⋮	⋮

3. In the same way that payments are made for phase 1, you will be paid on the designated dates through VENMO.

- Are there any questions?

Before we start, let me remind you that:

- The length of a match is randomly determined. After each round, there is a 75% probability that the match will continue for at least another round. You will play with the same person for the entire match.

- After a match is finished, you will be randomly paired with another person for a new match. You will have 20 such matches.

- One match will be randomly selected for payment.
You will receive your payment for the first round in 1 month from today. After that, you will receive payment for the following rounds once a month. That is, you will receive payment for the second round in 2 months, payment for the third round in 3 months, and so on.

C Screen shots for the experiment

C.1 Phase 1 screen shots

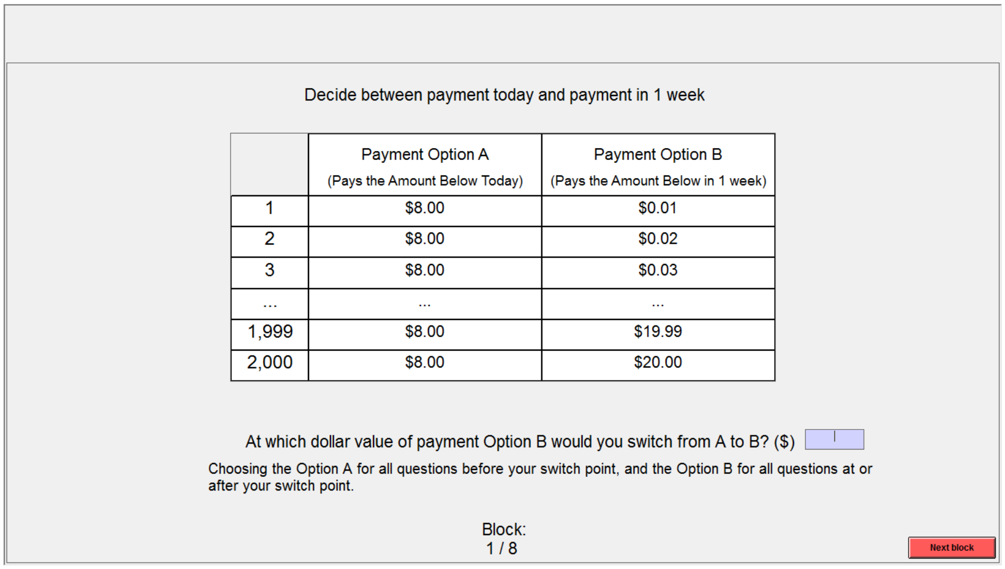

Figure C.1: The screen shot of phase 1 (block 1)

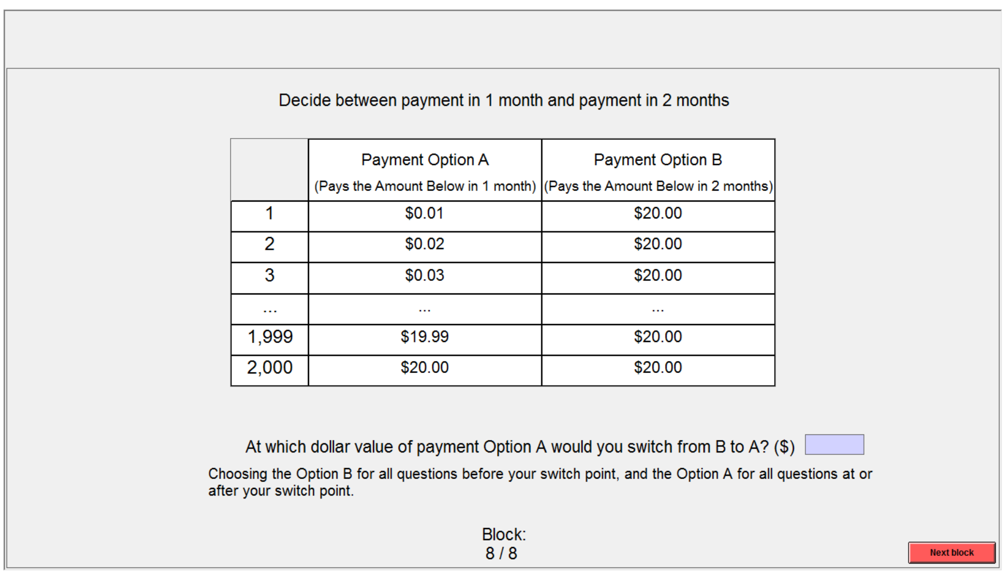

Figure C.2: The screen shot of phase 1 (block 8)

C.2 Phase 2 screen shots (Delay-Monthly treatment)

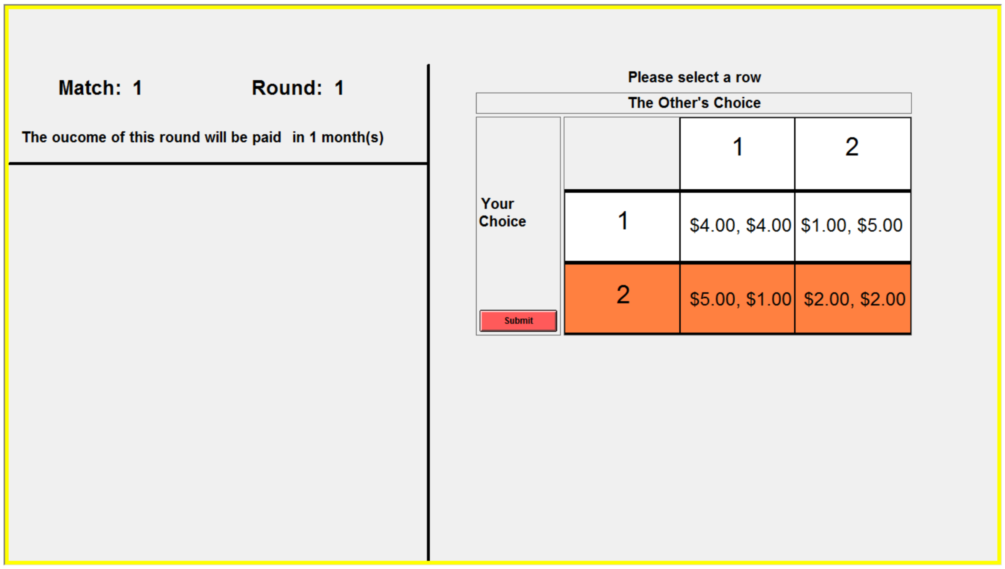

Figure C.3: The screen shot of the round 1 decision stage

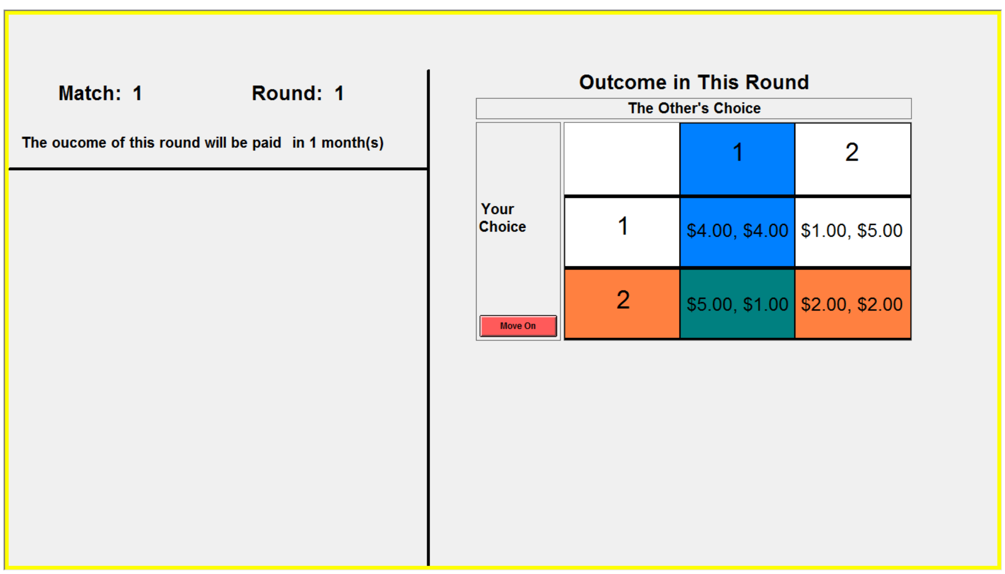

Figure C.4: The screen shot of the round 1 feedback stage

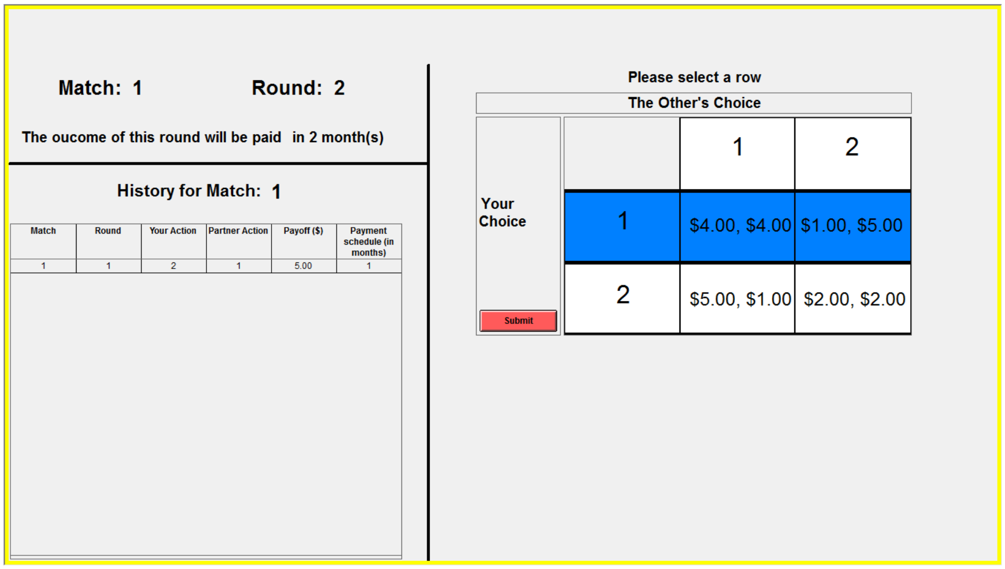

Figure C.5: The screen shot of the round 2 decision stage

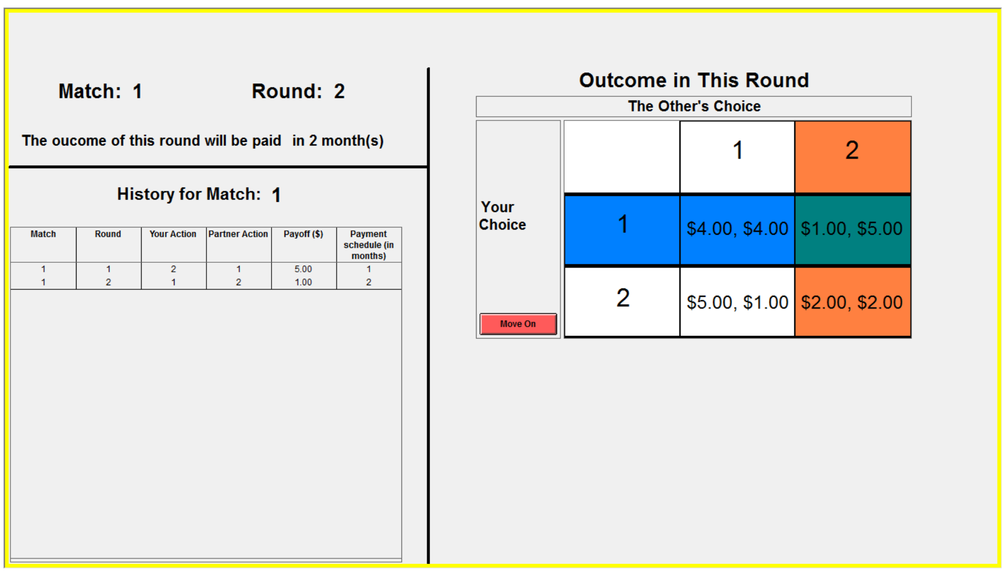

Figure C.6: The screen shot of the round 2 feedback stage