ESSAYS IN EXPERIMENTAL METHODOLOGY

by

Christopher James Waldron Candreva

A Dissertation Submitted to the Faculty of the

DEPARTMENT OF ECONOMICS

In Partial Fulfillment of the Requirements

For the Degree of

DOCTOR OF PHILOSOPHY

In the Graduate College

THE UNIVERSITY OF ARIZONA

2019

THE UNIVERSITY OF ARIZONA
GRADUATE COLLEGE

As members of the Dissertation Committee, we certify that we have read the dissertation prepared by Christopher James Waldron Candreva, titled *Essays in Experimental Methodology* and recommend that it be accepted as fulfilling the dissertation requirement for the Degree of Doctor of Philosophy.

_______________________________ Date: 6/14/19

[Committee Chair Name] Julian Romero

_______________________________ Date: 6/16/19

[Committee Member Name] Martin Dufwenberg

_______________________________ Date: 6/17/19

[Committee Member Name] Charles Noussair

Final approval and acceptance of this dissertation is contingent upon the candidate's submission of the final copies of the dissertation to the Graduate College.

I hereby certify that I have read this dissertation prepared under my direction and recommend that it be accepted as fulfilling the dissertation requirement.

_______________________________ Date: 6/14/19

[Committee Chair Name] Julian Romero
Dissertation Committee Chair
[Academic Department] Economics

Acknowledgements

I would like to thank the Economics Science Laboratory at the University of Arizona for providing me with both the funding to run my experiments, and a venue to run them in. I would like to thank my Advisor Julian Romero for assisting me with the code necessary to run my experiments, for the advice he gave me about how to implement my ideas, and for the necessary funding to complete these projecst. Additional Thanks to Charles Noussair and Martin Dufwenberg for their insights and feeback. Finally I like to thank all the faculty, students and presenters at the University of Arizona that allowed me to bounce ideas off of them, which in turn helped me to come to the formal conclusions that I present here in this dissertation. Your help has been invaluable.

Contents

List of Figures

List of Tables

Abstract

This thesis focuses on questions about experimental methodology, examining novel treatments and subjects' understanding of experimental environments. The first chapter examines differences between the traditional discounted model of repeated Prisoner's Dilemmas used in the theory, and the random termination model used in the laboratory. Under relatively general assumptions, the threshold δ^*-value, above which cooperation can be supported as a subgame perfect equilibrium, varies based on subject's specific utility function under random termination, but not under discounting. To test this, a new experimental design was created that captures features of the infinite horizon discounting model absent in random termination Using a between-subject, subjects played six repeated Prisoner's Dilemmas at $\delta = 0.98$, and either the new treatment or random termination. Cooperation rates were higher in the random termination treatment. Maximum likelihood estimation was used to determine what repeated games strategies subjects used during this experiment. Subjects in the discounted treatment tended to use more defective, yet forgiving strategies, while subjects in the random termination treatment used more cooperative yet less forgiving strategies. This finding suggests that more work needs to be done to understand the differences between random termination and discounting

The second chapter further examines the differences between these two treatments. One specific channel that could created a difference in behavior between random termination and discounting is subject's risk attitude. Using the theory model from chapter one, the threshold δ^*-value, above which cooperation can be supported as a subgame perfect equilibrium, varies based on subject's risk attitude in random termination, but is independent of subject's risk attitude under discounting. Using a within-subject design, subjects played 32 unique supergames. Each

supergame featured either this new treatment or random termination, one of eight distinct δ-values, and one of two different game matrices. This experiment was used to determine if risk attitude caused subjects' decision making to differ between the two settings. First period cooperation rates did not differ between the two treatments, even though differences were predicted by the risk attitudes elicited. Subjects' behavior did vary in more complex ways, though risk attitude did not explain these differences.

The final chapter examines the influence of pre-experiment tasks on subject understanding. This experiment used a 2x2 design varying the type of instructions and whether or not the pre-experiment quiz was incentivized. One set of instructions was based on instructions used in prior studies. The other was written using techniques from the Multi-Media Learning literature, which aims to find ways to maximize subject understanding from instruction. After the pre-experiment procedures, subjects participated in ten modified BDM selling markets with induced values. High ability subjects in this experiment showed higher understanding of the BDM mechanism after receiving Multi- Media Learning instructions compared to Standard instructions. This was true in both the first market, and across all markets. Incentivizing the quiz had a negligible impact on subject understanding.

Chapter 1

Long Horizon Repeated Games

1.1 Introduction

This paper looks to see if subjects treat long, randomly terminated repeated games the same way they treat discounted, "infinitely" repeated games. Random termination was been used as a proxy for infinite horizon discounting for decades. The reason it is considered a reasonable proxy for discounting is because for risk neutral, expected utility maximizers the two settings are ex ante equivalent. More generally, these two setting are ex ante equivalent so long as the decision maker has an affine utility function. However, years of experimental work has shown that subjects do not have affine preferences, with multiple channels potentially changing the shape of their utility functions. Furthermore, I show in section 1.2 that the critical discount factor, δ^*, above which the Grim Trigger strategy is a subgame perfect equilibrium of the repeated Prisoner's Dilemma, depends directly on the shape of the decision maker's utility function under random termination, but not under discounting. The exact value of δ^* under random termination differs from its value under discounting when utility is not affine. Thus it is possible that prior repeated games experiments are not having subjects participate in the environment they believe they are participating. Therefore, it is important to test if subjects treat randomly terminated repeated games the same way they treat discounted, “infinitely” repeated games.

To do this, I used a new, unique methodology for running repeated games experiments. In this methodology, the payoffs in the stage game matrix are discounted every period. That is the stage game in the second period is the the first period stage game with every payoff multiplied by δ, the stage game in the third period is the second period stage game with every payoff multiplied by δ, etc... Subjects then played this game for an indefinite number of periods. This design has two potential issues. The first is that it is not feasible to keep subjects in the lab for an infinite amount of time, no matter how much we may want to. The second is that, because payoffs in this game are continuously falling, at some point they will fail to satisfy dominance, as defined in Smith (1982). Since they exact point when payoffs fail to satisfy dominance will be different for each subject, an endogenous opt out was included to end the repeated interaction. Full details on this opt out and the rest of this treatment can be found in section 1.3.

To test if subjects treat these treatments differently, I ran an experiment using a between subject design. Subjects in each session played six repeated Prisoner’s Dilemma games with $\delta = 0.98$. This δ was chosen to ensure that there were no potential issued with cooperation being or not being a subgame perfect equilibrium, and to allow for a variety of experiences and behaviors within supergames. I find that subjects’ cooperation rates are statistically, significantly higher in the random termination treatment compared to the discounted treatment in both the first period of play, and across all periods of play (**Result 1.1**). Additionally, cooperation rates in the discounted treatment evolved differently over time, both within and across supergames, compared to what has been traditionally observed in the literature. In these supergames, first period cooperation rates remained stagnant across supergames, while all period cooperation rates rose across supergames, and were even higher than first period cooperation rates in the last couple of supergames (**Result 1.2** and **Result 1.3**). To further investigate these differences, the Strategy Frequency Estimation Method (SFEM) was used to determine if subjects' strategy choice varied by treatment type. Subjects in the discounted treatment tended to play more initially defective, yet forgiving strategies, while subjects in the random termination treatment tended to play more initially cooperative, yet less forgiving strategies (**Result 1.4**).

This is not the first study to question whether the random termination implementation is the optimal methodological choice for repeated games experiments. One important feature of the random termination treatment is that there is a direct mapping from the value of δ to the expected number of periods each supergame will last.[1] Fréchette and Yuksel (2017) looked to divorce the expected number of periods from the discount factor and document any changes in behavior that arise from doing so. To do this, they looked at three alternatives to the standard random termination treatment: discounting for four periods followed by a coordination game that represented the infinite number of period remaining, block random termination, and discounting for four periods followed by random termination at the same δ value. The authors also used a standard random termination treatment with the same stage game parameters as a baseline for behavior. They found that subjects were most likely to cooperate in the standard random termination treatment, compared to their alternative treatments. This is true for both action frequency and choice of repeated game strategy, where subjects in the discounted treatments played more defective

[1]The expected number of periods is always $\frac{1}{1-\delta}$ in the standard random termination implementation.

Figure 1.1: Generic Prisoner's Dilemma

	C	D
C	a, a	c, b
D	b, c	d, d

Notes: This is the game matrix of a standard Prisoner's Dilemma game if the following are true: $b > a > d > c$ and $2a > b + c$

strategies than those in the randomly terminated treatment.

The only other study that uses the same discounted methodolgy as this chapter is the one presented in chapter 2. That chapter looks to see if departures from risk neutrality caused subjects to behave differently in randomly terminated supergames compared to those that are "infinitely" discounted. Using a similar theory model to the one presented in section 1.2, that paper showed that as subjects became more risk averse, the threshold δ^* value, above which cooperation could be supported as a subgame perfect equilibrium in the infinitely repeated Prisoner's Dilemma, increased when the game was randomly terminated, but did not change with risk attitude when the game was discounted. The paper tested that theoretical result in the lab using a within subject design that had subjects play 32 distinct supergames varying between two ending rule, random termination and the discounted methodology, two stage games, the Prisoner's Dilemma and Battle of the Sexes, and eight distinct δ, ranging between 0.5 and 0.9. That paper found no difference in subject behavior in the first period of a new supergame, even after subject risk attitudes were accounted for. Simple, two period patterns of stage game outcomes were more likely to emerge in discounted supergames than they were in randomly terminated supergames, however.

1.2 Theoretical Background

Assume that two players are playing a repeated 2×2 matrix game with an indefinite time horizon. An example of one of these types of matrix games is provide in figure 1.1. I will assume that each player seeks to maximize some strictly monotonic utility function $U(\cdot)$. I will further assume that $U(\cdot)$ is applied to their entire earnings stream rather than their earnings in each period. This is because each player cares about their earnings from the entire repeated interaction and not necessarily their period by period earnings. This leads to the following being a player's ex ante utility for a supergame featuring discounted payoffs over an infinite time horizon:

$$U(\sum_{t=0}^{\infty} \delta^t x_t)$$

where δ is the discount rate and x_t is the payoff the player earns in period t. If the supergame instead features random termination, then the following is a player's ex ante utility:

$$(1-\delta)\sum_{t=0}^{\infty} \delta^t U(\sum_{k=0}^{t} x_k)$$

where δ is the continuation probability and x_t is the payoff the player earns in period t.

At the start of each supergame players will have to choose which action to play in period 1. If the stage game they are playing is a generic Prisoner's Dilemma, as in figure 1.1, cooperation can be supported as a subgame perfect equilibrium if the discount factor for that game is $\delta \geq \delta^*$, where δ^* solves the following equation in a supergame featuring discounting over an infinite time horizon:

$$U(\sum_{t=0}^{\infty} \delta^t a) = U(b + \sum_{t=1}^{\infty} \delta^t d) \tag{1.1}$$

Or instead the following equation if the supergame features random termination:

$$(1-\delta)\sum_{t=0}^{\infty}\delta^t U(\sum_{k=0}^{t} a) = (1-\delta)[U(b) + \sum_{t=1}^{\infty}\delta^t U(b+\sum_{k=1}^{t} d)] \tag{1.2}$$

In both these equations it is assumed that the other player is playing a grim trigger strategy.[2]

To find this threshold value in a discounted supergame, δ^{D*}, note that $U(\cdot)$ is strictly monotonic, and thus has an inverse. This means that equation 1.1 can be rewritten as the following:

$$\delta = \frac{b-a}{b-d} \tag{1.3}$$

Thus in discounted supergames, only the payoff parameters of the stage game and not the player's utility function, impacts the cutoff value of delta, δ^{D*}, that dictates whether or not cooperation can be supported as a subgame perfect equilibrium.

In supergames that feature random termination however, equation 1.2 can only be simplified to:

$$\sum_{t=0}^{\infty}\delta^t U(\sum_{k=0}^{t} a) = [U(b) + \sum_{t=1}^{\infty}\delta^t U(b+\sum_{k=1}^{t} d)]$$

which implies that δ^{RT*} directly depends on $U(\cdot)$ and not just the payoff parameters. Thus player's utility function directly impacts the cutoff value δ^{RT*}. An analytical solution to this problem requires further assumptions on the functional form of $U(\cdot)$. One functional form for which analysis has already been performed is $U(x) = x^{1-r}$. With this utility specification δ^{RT*} varies with r, while δ^{D*}, does not vary with r.

[2]This assumption is standard in the literature. Grim Trigger is the strategy that establishes the lower bound for what δ values support cooperation as a subgame perfect equilibrium for the repeated Prisoner's Dilemma, because it is the harshest possible punishment. See Dal Bóand Fréchette (2011) and Blonski, Ockenfels and Spagnolo (2011).

Figure 1.2: Changes in δ^* for $U(x) = x^{1-r}$

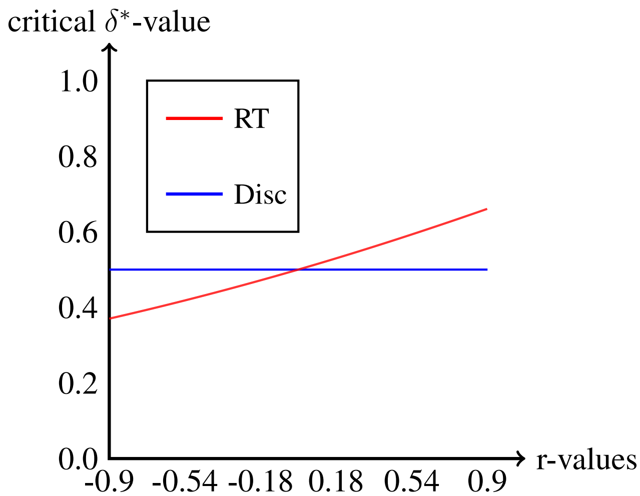

Notes: These δ^*-values are for the stage game shown in figure 1.3. These two critical δ^*-values, δ^{D*} and δ^{RT*}, are equal to one another when $r = 0$. This is the parameter specification that makes someone with this utility function risk neutral.

Figure 1.2 shows how these δ^*-values change with r for the the stage game shown in figure 1.3. This is the stage game used in the experiment. This figure follows the model prediction that δ^{RT*} depends on $U(\cdot)$, while δ^{D*} does not. Thus the model leads to the following theoretical result:

Theoretical Result 1.1. *The critical δ^*-value depends on the treatment type of the game:*

1. *In discounting, δ^{D*} depends only on the payoff parameters of the game, and NOT the subject's utility function.*

2. *In random termination, δ^{RT*} depends on the payoff parameters of the game AND the subject's specific utility fucntion.*

There are multiple different channels through which $U(\cdot)$ can change δ^{RT*}. Chapter 2 of this dissertation explores risk attitude as one potential channel. Other potential channels include higher order risk attitudes, specifically whether a subject is prudent or imprudent, ambiguity aversion, and models of noisy decision making, such as logistic choice utility. This paper does not look to see if a specific channel causes behavior to differ between random termination supergames and discounted supergames. Instead it looks to see if long run decision making differs across the two treatments, allowing for the possibility that any or all of these channels could be influencing these outcomes. This model does still provide an explanation of how differences between the two treatments can arise.

1.3 Experimental Design

1.3.1 Overview

In this experiment, subject's participated in six repeated Prisoner's Dilemma games using the stage game in figure 1.3. All supergames featured $\delta = 0.98$. $\delta = 0.98$ was chosen to ensure that the repeated games were sufficiently long, allowing for a wide range of possible difference in behavior between the two treatments to emerge. Neutral action labels were used during in the experiment in place of the C and D labels used in this paper. This stage game comes from Ioannou and Romero (2014). This stage game was chosen because $\delta^{D*} = 0.5$, while the Blonski, Ockenfels, and Spagnolo (2011) threshold, henceforth δ^{BOS}, is $\delta^{BOS} = \frac{2}{3}$. The δ^{BOS} threshold determines when it is risk dominate to cooperate in the first period of a supergame when the other player plays Always Defect 50% of the time and Grim Trigger the other 50% of the time. Both of these values are significantly below $\delta = 0.98$, so there are no structural barriers to cooperation in this game.

Six repeated games were chosen based on simulation results. Given a predetermined population of repeated game strategies used by simulation subjects with

Figure 1.3: PD Matrix

	C	D
C	3, 3	1, 4
D	4, 1	2, 2

Notes: This is the stage game subjects played with in the random termination treatment and in the first period of the discounted treatment. Neutral action labels were used during the experiment (U and L for the row player and L and R for the column player).

noisy decision making, it was found that it took at least four supergames worth of data to get an accurate estimate of the repeated games strategies for supergames of this δ-value. Two additional supergames were added at the start of the experiment to give subjects a chance to learn the environment. Thus a total of six supergames were run in each session. At the start of each supergame subjects were randomly paired with one another. In each period subjects were asked to make a choice about which action they preferred to play and which action they thought the other player was going to play. After both subjects made their choices, they were shown the outcome for that period and their total payoff for that repeated game was updated appropriately.

In each session, the supergames were either randomly terminated or featured discounted payoffs with the option to opt out. This study, in part, aimed to be the first to document subjects' choice of repeated game strategy under discounting. If subjects were to play under both treatments in the same session, then their experiences under discounting could potentially affect their play under random termination, or vice versa. Thus to eliminate this potential contamination, a between-subject design was used.

1.3.2 Random Termination

In the random termination treatment, game lengths were predrawn by the experimenter. This is a departure from the design used in chapter 2. In that study dice were rolled at the front of the room to determine whether random termination supergames would continue or end. They were used to help make the random component of random termination more salient in the eyes of the subjects compared to the discounted supergames they also participated in. The use of dice was reasonable in that study, since that experiment used relatively low δ-values.[3] This meant that the supergames were not particularly long and it did not slow down the experiment by too much to have all the subjects progress through those games at the same pace.[4]

In this experiment, with $\delta = 0.98$, each supergame lasts significantly longer, in expectation, than they did in the previous study.[5] This incentivizes having each group progress at their own pace, so that an individual group taking longer in a particular period does not slow down the entire experiment. One way to achieve independent pacing is to have the computer randomly determines whether or not a given group continues with the supergame after each period. In each session the order that subjects experienced these game lengths was randomized i to help control for any order effects.

1.3.3 Discounted

In the discounted treatment, subjects played a repeated game closer to how they are typically describe in theory. At the start of each period, after the first, all the

[3]The highest δ-value used was 0.9, which means those games had 10 periods in expectation. Subjects played two random termination supergames at that δ-value per session.

[4]The longest any random termination supergame lasted was 28 periods. Only three supergames lasted for more than 20 periods.

[5]Expected length is 50 periods. Additionally there is a 13.5% chance any individual supergame goes for 100 or more periods.

payoffs in the stage game were multiplied by $\delta = 0.98$,[6] and subjects played for an indefinite number of periods. The subjects' repeated game payoffs were the sum of all their individual stage game payoffs. To prevent these supergame from lasting forever, one departure was made from theory. Subjects were allowed to opt out of continuing to play the supergame following any period. This option was included for two reasons. First, without a way to end these supergames, subjects would be stuck in the lab for an infinite amount of time, which is not feasible. Second, eventually the payoffs from a discounted supergame get so small they fail to satisfy dominance, as defined in Smith (1982). Since the exact point in time this is true will vary from subject to subject, the opt out option is the cleanest way to give the subjects control over this, while still maintaining an indefinite time horizon. Another option for maintaining dominance would be to end the game after a fixed payoff level was reached. This however would artificially create a final period for the game. This would turn the game into a finitely repeated game, which fundamentally alters any predictions one could make.

Additionally, in order to maintain equivalence between the two treatments, subjects in the discounted treatment where given a continuation payoff to compensate them for the infinite number of periods they chose not to play. This continuation payoff was the infinitely discounted sum of the sucker payoff in the period the supergame ended for them.[7] Subjects were told what their continuation would be

[6]This is the same way Fréchette and Yuksel (2017) implemented discounting in their two discounted treatments.

[7]The sucker payoff was chosen for two reasons. First, I wanted to ensure that the continuation payoff did not influence the decisions subjects made during play. This is why the discounted value of subjects' average payoff was not used. Second, I wanted as strong a distinction as possible between the random termination supergames and the discounted supergames. This is why I did not just give them the infinitely discounted sum of a random payoff from the matrix in the period the supergame ended. The maximum payoff could not be selected, because then subjects would be incentivized to opt out immediately and receive their highest possible payoff. The mutual defection payoff was not chosen, since then subjects' have not incentive to play the repeated game once they reach the mutual defection outcome, which reduces the likelihood that subjects attempt to return to mutual

following every period to help them make their decision of when to opt out. Subjects were also told what the maximum amount they could possibly earn, i.e. the infinitely discounted sum of the temptation payoff, was following each period. For example, in period 10, the continuation payoff to a subject in this experiment would be 40.85 and the maximum remaining payoff would be 163.42. Subjects were informed about these two payoffs in the instructions and on each periods summary screen.

Lastly, these discounted supergames lasted until both subjects chose to opt out. This choice was made in order to help maintain certainty of experience. If the supergame ended after the first subject chose to opt out then half the subjects would not know when their supergame would end. Since uncertainty of the exact game length is a feature of random termination this treatment was designed to avoid, it was decided to end the supergame after both players chose to opt out. This choice does still introduce uncertainty of when the supergames ends for the subject that chose to opt out first in the periods after they chose to opt out. For this reason, only the periods prior to either subject opting out was considered in the analysis of the results of this experiment. This meant that in many instances, subjects continued manually making decisions even after they chose to opt out.

1.3.4 Hypotheses

The hypotheses for this experiment come from the results of chapter 2 and Fréchette and Yuksel (2017). The first hypothesis has to do with the difference in cooperation rates between the discounted treatment and the random termination treatment. chapter 2 found no difference in first or all period cooperation rates between the

cooperation. Thus the sucker payoff was chosen, even though both players receiving the sucker payoff falls outside the set of feasible payoffs.

two treatments. This is the first hypothesis for this paper:

Hypothesis 1.1. *There will be no difference in either first period or all period cooperation rates between the discounted treatment and the random termination treatment.*

The other hypothesis for this experiment has to do with strategy estimation from SFEM. This procedure was not used on the data from chapter 2 (working paper), because the structure of that data set does not fit the necessary conditions of SFEM.[8] Other analysis for long run behavior was looked at using that data set, however. That analysis showed a higher propensity for defective behavior in the discounted treatment compared to the random termination treatment. Additionally, Fréchette and Yuksel (2017) found higher subjects in their discounted treatments were more likely to play defective, repeated games strategies. These two findings together lead to the following hypothesis on which strategies subjects use:

Hypothesis 1.2. *Subjects in the discounted treatment will be more likely to use defective strategies (All D and various suspicious strategies) than subjects in the random termination treatments.*

1.4 Results

Eight sessions of this experiment, four for each treatment, were run at the Economics Science Laboratory (ESL) at the University of Arizona between December 2017 and February 2018. The random termination sessions lasted for one hour on average and subjects earned an average of $17.48, with minimum earnings of $11.00 and maximum earnings of $24.00. The discounted sessions lasted an hour

[8]SFEM requires multiple supergames worth of data for each subject per parameter specification. In chapter 2 each subject only played one supergame with each parameter specification.

Table 1.1: Cooperative Behavior

	Overall	Disc.		RT
First Period	62.8%	53.4%	$<^{***}$	73.7%
		$\vee$		$\vee^{***}$
All Periods	56.7%	51.6%	$<^{**}$	62.5%

Notes: **: Significant at 5% level. ***: Significant at the 1% level. Significance determined by Wilcoxon rank-sum test. The unit of observation is the cooperation rate of one subject across all supergames. Overall is both treatments combined, Disc. is the discounted treatment and RT is the random termination treatment. First Period is the first period of a supergame. All periods is every period subjects played in RT and every period prior to either subject in the pair opting out in Disc.

and 45 minutes on average and subjects earned an average of $16.29, with minimum earnings of $10.25 and maximum earnings of $23.00. Each session had between 8 and 18 subjects. Final subject counts by treatment type were 50 for random termination and 58 for discounted.

Table 1.1 shows the cooperation rates, both in the overall experiment and by treatment type in the first period of supergames and across all periods of the supergames.[9] Unlike in Chapter 2, cooperation rates in the first period in the random termination treatment were significantly higher than they were in the discounted treatment. All period cooperation rates were also significantly higher in the random termination treatment compared to the discounted treatment. This leads to the following result in contrast to **Hypothesis 1.1**:

Result 1.1. *Cooperation rates are significantly higher in the random termination treatment compared to the discounted treatment.*

Table 1.1 also shows further differences exists between the two treatments. Across all supergames first period cooperation rates in the randomly terminated treatment were significantly higher than the all period cooperation rates. In the dis-

[9] For the discounted treatment, only periods that were played prior to either subject opting out were considered.

Figure 1.4: Cooperation Rates Across Matches

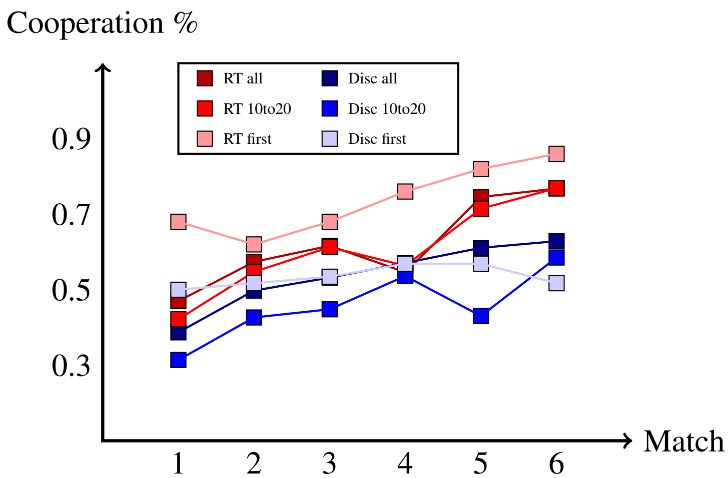

Notes: This graph shows how cooperation rates evolved across supergames, with each square representing one supergame. The unit of observation is the cooperation rate of each subject over the given number of periods. RT stands for the random termination treatment, Disc. for the discounted treatment. First was the first period of a supergame. All is every period in RT and every period prior to either subject in a pair opting out in Disc. 10to20 looks at periods 10 through 20 in RT and all periods between 10 and 20 that occur prior to either subject in a pair opting out in Disc.

counting treatments, on the other hand no significant differences exist between first period and all period cooperation rates.

To help determine why these differences exist, I looked to see how cooperation rates varied over time. Figure 1.4 shows cooperation rates by supergame in the first period, all periods and periods 10-20 in both treatment types. For the random termination treatment, cooperation rates evolve as is typically seen in the literature. Cooperation rates increased as subjects gained more experience, with the first period cooperation rates being greater than the all period cooperation rates. Furthermore the cooperation rates in periods 10 to 20 closely follow the all period cooperation rates in the random termination treatment, suggesting that differences from the first period arise early in the supergames.

Cooperation rates in the discounted treatment however, did not follow the typically observed patterns. In this treatment, first period cooperation rates did not

Table 1.2: Cooperation Rate Across Supergames

		First		Sixth
RT	First Period	68.0%	$<^{**}$	86.0%
		$\vee^{**}$		$\vee^{***}$
	All Periods	44.5%	$<^{***}$	76.0%
Disc	First Period	50.0%	$<$	51.%
		$\vee$		$\wedge$
	All Periods	35.7%	$<^{***}$	62.4%

Notes: **: significant at the 5% level. *** significant at the 1% level. Significance determined by Wilcoxon rank-sum test. The unit of observation is the average cooperation rate of one subject within the given timeframe. First is for cooperation rates in the first supergame subjects played. Sixth is for cooperation rates in the sixth supergame subjects played. RT stands for random termination. Disc stands for discounting. First period is the first period of a supergame. All periods is subjects average cooperation level across all periods in RT and across all periods before either player opted out in Disc.

significantly vary across supergames. All period cooperation rates, however do increase as subjects gained more experience, and by the last supergame all period cooperation rates were higher than first period cooperation rates. Additionally, the all period cooperation rates were higher than the cooperation rates for periods 10 to 20 in the discounted treatment. This suggests that subjects are converging to cooperation later in supergames rather than early in them, contrary to how cooperation rates behave in the random termination treatment.

To confirm that these observations were significant, I looked to see how cooperation rates in the first supergame varied from those in the sixth supergame. Table 1.2 shows first and all period cooperation rates in the first and sixth supergames in the random termination and discounted treatments. As was expected from figure 1.4, the cooperation rates in supergame six is significantly higher than those in supergame one for the random termination treatment. Additionally, cooperation rates in the first period we significantly higher than the all period cooperation rates in both these supergames. In the discounted treatment, on the other hand, only the all period cooperation rates between the first supergame and sixth supergame were significantly different. This provides further evidence for the observations from figure 1.4. This leads to the following result:

Result 1.2. *Cooperation rates raise as subjects gain more experience in the random termination treatment. In the discounted treatment, first period cooperation rates do not change as subjects gain more experience. Subjects do learn to cooperate more across all periods in the discounted treatment.*

In addition to investigating how cooperation evolved across supergames, I also looked at how cooperation rates evolved within supergames. To do this I divided each supergame into 5-period bins and looked at how cooperation rates changed from bin to bin. Figure 1.5a shows this evolution for the random termination treatment. For the discounted treatment the display was separated into two graphs. Figure 1.5b includes bins up through period 60, matching the random termination supergames, and figure 1.5c shows all bins after period 60.

For the random termination treatment the within supergame cooperation rates evolve as expected. The first bin is slightly higher than the later bins, and cooperation rates in later supergames are higher on average than those in earlier supergames. The first part of the discounted treatment looks somewhat similar to the random termination treatment. The main difference is the lower overall cooperation rates and a slight upward trend in the later supergames following an initial decrease. The latter part of the discounted treatment, however is much different than the earlier part. Here the cooperation rates diversify, going from being bunched between 40% and 80% to the extremes of either 100% or 0% cooperation, depending on the supergame. These extremes are primarily driven by the decreased sample size[10]

[10]Starting with bin 61-65, the number of subjects still participating in a supergame varied between

Figure 1.5: Cooperation Rates Within Supergames

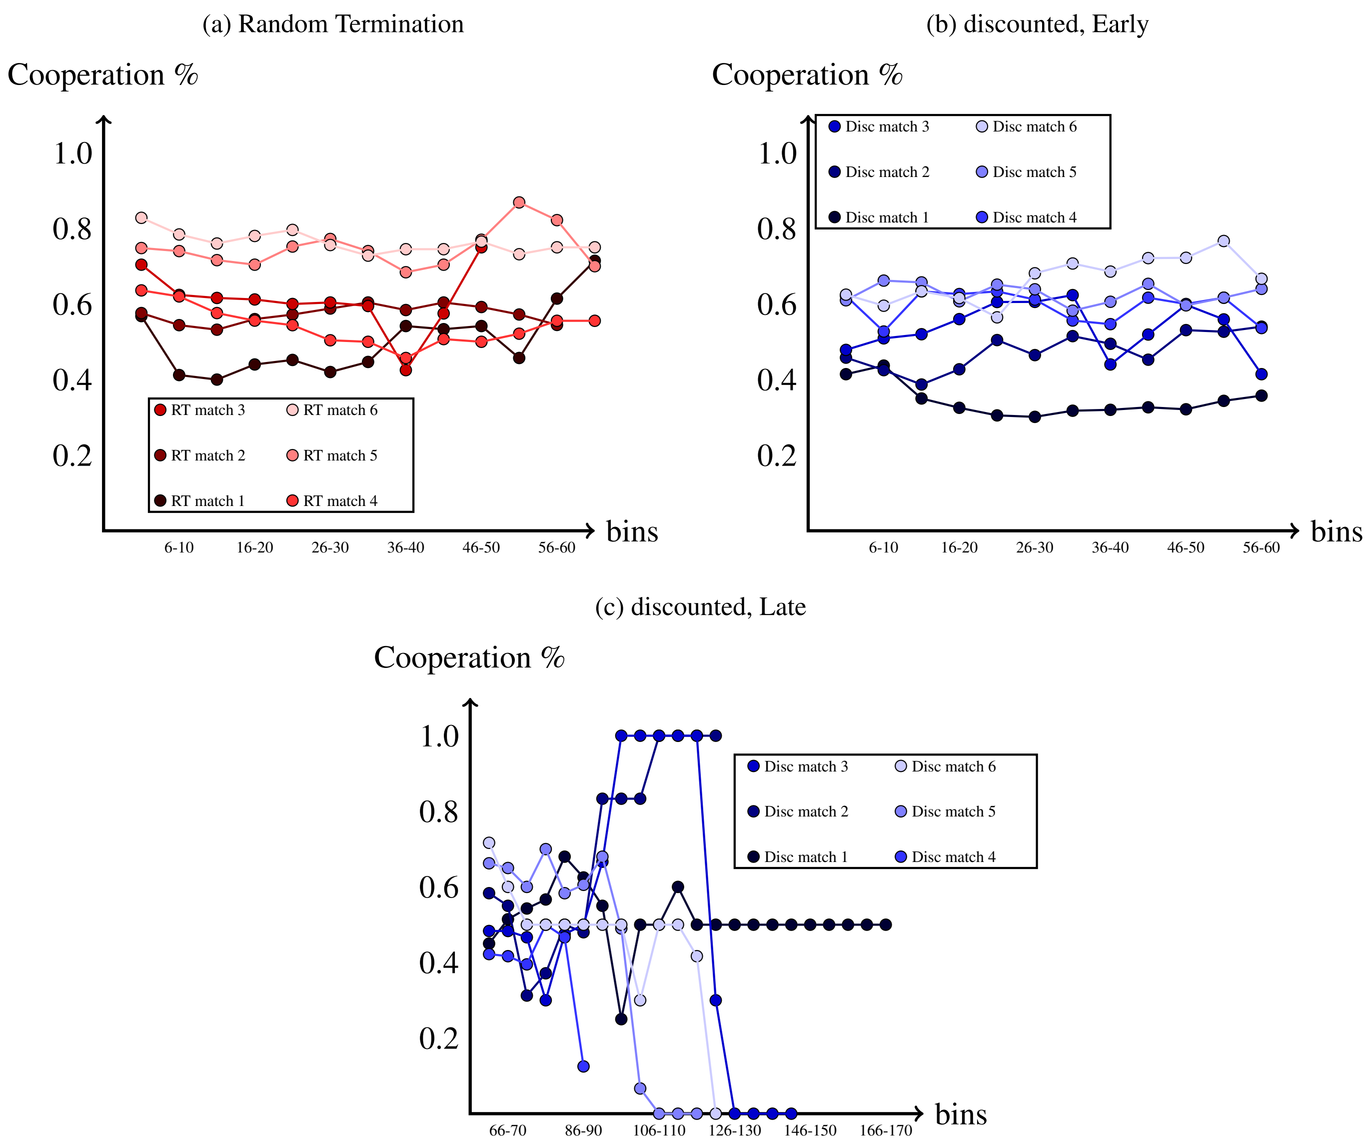

Notes: Each line represents the evolution of cooperation within a given supergame. Each dot represents the cooperation rate in that 5-period bin. The unit of observation is the cooperation rate of a subject within that 5-period bin. Lines continue until no subjects had observations remaining. RT stands for the random termination treatment, and Disc stands for the discounted treatment. RT lasted for each subject until that supergame ended for them. Disc lasted until either subject in a pair opted out.

Table 1.3: Cooperation Rates Within Supergames

(a) Random Termination

Supergame	First 5		Last 5
One	56.8%	>	47.2%
Two	57.6%	>	56.8%
Three	70.4%	>	65.2%
Four	63.6%	>	56.8%
Five	74.8%	<	77.6%
Six	82.8%	>	70.8%

(b) Discounting

Supergame	First 5		Last 5
One	33.3%	>*	27.4%
Two	40.0%	>	34.6%
Three	38.6%	>	34.5%
Four	48.3%	>	43.0%
Five	51.9%	>*	36.2%
Six	48.0%	<	48.5%

Notes: *: significant at the 10% level. Significance determine by Wilcoxon ran-sum test. The unit of observation is the average cooperation rate of one subject within the given timeframe. First 5 is the average cooperation rate in the first five periods of a supergame. Last 5 was the average cooperation rate in the last five periods of a supergame in random termination, or the last 5 periods before either player opted out in discounting. For the discounting treatment, only supergames were neither subject opted out before period 10 were included, to prevent a period from appearing in both First 5 and Last 5.

towards the end of the discounted treatment polarizing the cooperation rates.

To confirm the trends observed in figure 1.5 are true, I looked to see if the cooperation rates in the first five periods of each supergame significantly differed from those in the last five periods of supergame. The results of these tests are on table 1.3. In the random termination treatment, cooperation rates in the first five periods are statistically indistinguishable from those in the last five periods for all six supergames. In the discounted treatment, the cooperation rates in the first five periods of supergames one and five was significantly higher than it was in the last five periods, but only at the 10% level. This difference is insignificant in the other four supergames. Thus even though the actual level of cooperation differs between the two treatments, the way cooperative behavior changes within supergames does not vary significantly. Therefor:

Result 1.3. *Cooperation rates in both treatment types evolve similarly within su-*

12 and 24, with the final few bins having as few as two subjects, compared to bin 1-5 which had 56 subjects.

pergames.

As a final check on differences in behavior between the two treatments, I looked to see what repeated game strategies subjects played. To do this I followed the Strategy Frequency Estimation Method (SFEM) originally proposed by Dal Bó and Fréchette (2011). This methodology assumes that each subject chooses a fixed strategy at the beginning of the session, or for this project starting in supergame 3,[11] and that subjects have a chance of making a mistake when choosing their action each period. Thus every sequence of choices has a positive probability of occurring given a fixed strategy.

Specifically, it is assumed that in each round r of match m, subject i playing strategy s^k makes a mistake with probability $(1-\beta)$. That is they will play D when the strategy said they should play C or vice versa. Letting y_{imr} be an indicator function that is equal to 1 if subject i played the action that strategy s^k prescribed in period r of match m leads to the following expression for $p_i(s^k)$, the probability that subject i played strategy s^k:

$$p_i(s^k) = \prod_M \prod_R \beta^{y_{imr}}(1-\beta)^{1-y_{imr}}$$

where M is the set of matches subject i played, and R is the set of rounds in a given match.

Then given a set of strategies K and proportions ρ, a probability distribution on K, I can derive the a likelihood function for the entire sample:

$$\sum_S \sum_I \ln\left(\sum_{s^k \in K} \rho(s^k) p_i(s^k)\right)$$

[11]Using this technique on simulated data showed that at least four supergames were needed for accurate estimates. Thus the first two supergames were dropped to allow subjects a chance to learn.

Where S is the set of experimental sessions for a given treatment, and I is the set of subjects in a given session. Note that this set up assumes that all subjects are ex ante identical. Additionally since ρ is a distribution over K the final estimates that come from maximizing the above likelihood function gives us estimates of $\rho(s^k)$, which is the portion of the population that plays strategy s^k and β, the probability that subjects play the action their strategy tells them to. The set of strategies K that are considered in this paper are on table 1.4. This is the standard set of 20 strategies typically used in the literature.[12] After these estimates were found, 1000 bootstrap samples were constructed for each treatment by randomly sampling the appropriate number of subjects with replacement. Standard errors for each estimate were found using the estimates from the 1000 bootstrap samples.

Table 1.5 shows the SFEM estimates for each strategy who's population portion was statistically significantly greater than 0. Full estimates are provided in appendix A.1. The data used to generate these estimates are the first 20 periods of each of the last four supergames subjects played. Only the first 20 periods were used to ensure that longer supergames do not have a larger impact on the estimated strategies simply because they are longer. As one can see, the subjects in each treatment use strategies in different proportions that provides some support for **Hypothesis 1.2**. The subjects in the discounted treatment use more defective and forgiving strategies, as seen by the higher levels of All D and the various TFT strategies, while the subjects in the random termination treatment tend to use less forgiving cooperative strategies, as seen by the higher levels of Grim strategies.

One possible explanation for this difference is how subjects perceive the supergame. In both the random termination treatment and the discounted treatment,

[12] Other studies that use this set include Fudenberg, Rand and Dreber (2012), Fréchette and Yuksel (2017), and Romero and Rosokha (2018).

Table 1.4: Set of Strategies

Strategy	Abbreviation	Description
Always Cooperate	All C	Always play C
Always Defect	All D	Always Play D
Tit-for-Tat	TFT	Play C unless partner played D last round
Tit-for-2-Tats	TF2T	Play C unless partner played D in both of the last 2 rounds
Tit-for-3-Tats	TF3T	Play C unless partner played D in all of the last 3 rounds
2-Tits-for-1-Tat	T2FT	Play C unless partner played D in either of the last 2 rounds (2 rounds of punishment if partner plays D)
2-Tits-for-2-Tats	T2F2T	Play C unless partner played 2 consecutive Ds in the last 3 rounds (2 rounds of punishment if partner plays D twice in a row)
Suspicious Tit-for-Tat	STFT	Play D in the first round, then play TFT
Suspicious Tit-for-2-Tats	STF2T	Play D in the first round, then play TF2T
Suspicious Tit-for-3-Tats	STF3T	Play D in the first round, then play TF3T
Grim Trigger	Grim	Play C until either player plays D, then play D forever
Lenient Grim 2	Grim2	Play C until 2 consecutive rounds occur in which either player played D, then play D forever
Lenient Grim 3	Grim3	Play C until 3 consecutive rounds occur in which either player played D, then play D forever
Suspicious Grim 2	SGrim2	Play D in the first period, then play Grim2
Suspicious Grim 3	SGrim3	Play D in the first period, then play Grim3
Win-Stay-Lose-Shift	WSLS	Play C if both players chose the same action last round, otherwise play D
Win-Stay-Lose-Shift with 2 rounds of punishment	WSLS2	Play C if both players played C in the last 2 rounds, both players played D in the last 2 rounds, or both players played D 2 rounds ago and C last round. Otherwise play D
Punish Twice	T2	Play C until either player plays D, then play D twice and return to C (regardless of all actions played during punishment rounds)
False cooperator	False C	Play C in the first round, then D forever
Alternator	ALT	Start with D, then alternate between C and D

Notes: This table provides the full list of strategies used in the Strategy Frequency Estimation Method (SFEM) procedure. This first column provide the name of the strategy, the second it's abbreviation that will be referenced in the text and the final column provides a description of the strategy and how it determines its action choice.

Table 1.5: SFEM Results

Strategies	Disc	RT
All D	0.155^{***}	0.060^{*}
	(0.049)	(0.034)
TFT	0.194^{***}	0.123^{*}
	(0.060)	(0.071)
TF2T	0.108^{**}	0.051
	(0.050)	(0.052)
STFT	0.249^{***}	0.082^{*}
	(0.061)	(0.048)
Grim	0.036	0.103^{*}
	(0.024)	(0.059)
Grim2	$5.73E^{-15}$	0.198^{**}
	($3.38E^{-3}$)	(0.087)
Grim3	0.089^{**}	$2.73E^{-13}$
	(0.044)	(0.021)
β	0.863	0.914
	($4.48E^{-4}$)	($4.55E^{-4}$)

Notes: ***: Significant at the 1% level. **: Significant at the 5% level. *: Significant at the 10% level. Disc is the estimation using data from the discounted treatment. RT is the estimation using the data from the random termination treatment. The first 20 periods (or every period if the supergame lasted less than 20 periods) of each of the last four supergames were used to generate these estimates. Only strategies that were statistically significantly greater than zero for at least one treatment were included in this table. Standard errors generated using 1000 bootstrap samples.

earlier periods are worth more to subjects than later periods. In the random termination treatment, this is because the probability of encountering earlier periods in higher than the probability of encountering later periods. In the discounted treatment this is more explicit since subjects per period payoffs decrease as they play more periods. This explicit difference makes it easier for subjects to see the gains from defection in the discounted treatment compared to the random termination treatment; since the per period gain from defection in the random termination treatment is static. Because these gains are more visible, they promote a higher level of defection early, resulting in higher levels of STFT and All D in the discounted treatment.

Continuing along this train of thought, as subjects in the discounted treatment get further into a given supergame, the absolute gain from defecting decreases, even though the relative difference remains the same. If subjects make their decisions based on the absolute difference instead of the relative difference, then they would be more likely to try and cooperate later in the discounted treatment compared to the random termination treatment. This could be a reason behind why subjects play more forgiving TFT-type strategies in the discounted treatment, but less forgiving Grim-type strategies in the random termination treatment, where the absolute difference between defection and cooperation is fixed throughout the supergame.

Another potential influence on which strategies subjects chose is how the supergames end. In the discounted treatment, subjects can know with certainty whether or not the supergame will continue, since they know it will not end if they have yet to opt out. In the random termination treatment subjects have no such certainty, since following any period there is a nonzero probability of the supergame ending. This certainty of continuance alongside the decrease in absolute difference between cooperation and defection could cause subjects to try and return to cooperation at a higher rate in the discounted treatment compared to the random termination treatment. This incentive could lead to higher levels of TFT-type strategies in the discounted treatment and more Grim strategies in the random termination treatment in the absence of these incentives.

This explanation seems even more plausible when one considers the β values estimated using this data. In the discounted treatment, a estimated β-value of 0.863 implies that subjects make a mistake almost 14% of the time. This β-value is significantly lower than the 0.914 value estimated for the random termination treatment and is also significantly lower than has been previously estimated in the literature.[13] One potential reason why this estimate is as low as it is is that subjects are not actually playing TFT-type strategies, but something similar where they make the conscious switch to cooperation later in a supergame even though a TFT-type strategy would not tell them to do so. They do this because the absolute gains from defection are no longer enticing, and they prefer to try and establish a cooperative relationship with the other subject they were matched with. Future work is needed to determine if these types of strategies lead to a better fit. For now, we can say the following:

Result 1.4. *Subjects in the discounted treatment are more likely to play intially defective and/or forgiving strategies (TFT,STFT,All D,...). Conversely subjects in the random termination treatment are more likely to play intially cooperative, yet less forgiving strategies (Grim, Grim2,...).*

[13] Fréchette and Yuksel (2017) find β-values between 0.898 and 0.936. Dal Bó and Fréchette (2016) find β-values between 0.92 and 0.97.

1.5 Conclusion

This paper looked to see if subject behavior differed between randomly terminated repeated games, and those that that were discounted over an "infinite" time horizon. To do this, I used a unique experimental design that allowed subjects to play a discounted repeated game, with the option to opt out if they believed payoffs had fallen too low to continue playing. A between subject design was used, with subjects playing six repeated games with $\delta = 0.98$ and randomly rematched partners at the start of each new interaction. I find that subjects cooperate at a higher rate in random termination supergames compared to discounted supergames. This is true both in the first period of these supergames and across all periods of these supergames.

Additionally, while cooperation rates in the random termination treatment had their traditional upward trajectory as subjects learned the environment, the cooperation rates in the discounted treatment evolved differently. First period cooperation rates remained static across all six supergames subjects played, while all period cooperation rates increased as subjects gained experience. In fact all period cooperation rates in the discounted treatment were actually higher than first period cooperation rates in supergames five and six. Closer examinations of how cooperation evolved within given supergames showed that over the first half of discounted supergames, cooperation evolved similarly to how it evolved in the random termination supergames, but over the back half wide swings in cooperation rates occurred as fewer and fewer pairs of subjects continued to play.

In addition to examining differences in cooperation rates between the two treatments, SFEM was used to see if subjects' played different repeated games strategies in the two different treatments. Subjects in the discounted treatment played initially defective strategies at a higher rate than those in the randomly terminated treatment.

Subjects in this treatment also tended to play more forgiving Tit-for-Tat style strategies, while subjects in the random termination treatment were more likely to play more punishing, Grim trigger style strategies.

These differences may be due to the fact that subjects base their decisions on absolute differences in payoffs between cooperation and defection instead of relative differences. This could make the early periods of a random termination treatment look less important, even though relatively speaking they are just as import as the early periods of a discounted supergame. Thus subjects are more likely to try and cooperate early in the random termination treatment while they try and get the high temptation payoff as soon as possible in the discounted treatment. Later in the supergame, when the absolute differences between cooperating and defecting fall, subjects instead try to maximize their flow payoffs by cooperating. This switch to cooperation from defection for reasons outside of the behavior of the subject one is matched with maybe why a lower value of β was estimated for the discounted treatment compared to the random termination treatment in this experiment, or those that have been previously estimated in the literature. Further research is need to see if these types of strategies that switch to cooperation later in the supergame better fit the data generated in the discounted treatment.

For now, all we can say is that it appears that subjects behave differently in these two treatments when playing at relatively high δ-values. Further research looking at smaller δ-values using a between subject design could be useful to help determine the results in this paper hold when the set of possible equilibria is restricted. It would also be useful to determine exactly which channel influences subjects decisions across these two environments. Helping to pin down the extract reason why subject behavior in discounted supergames differs from their behavior in randomly terminated ones.

Chapter 2

Risk and Repeated Games

2.1 Introduction

This paper seeks to add to the experimental literature on repeated games. It does so by addressing a potential issue in these experiments: that subjects who are not risk neutral may approach random termination differently than they would approach an infinite time horizon, discounted environment. To address this, subjects participated in a series of risk elicitation tasks to get a well rounded estimate of their risk attitude. They then played a series of repeated games featuring either random termination of the game, or a new repeated game treatment that capture elements of infinite horizon discounting not present in random termination. While some differences in behavior were observed between the two treatments, those differences can not be attributed to subjects' risk attitudes.

Repeated games have long been studied in the economics literature. Fudenberg and Maskin (1986) formalized the Folk Theorem; they proved that if games are infinitely repeated and players are sufficiently patient, i.e. the discount factor is high enough, then any strategy that gives players at least their minmax payoff can be supported as a subgame perfect equilibrium. Since then, many other papers have continued to look at Folk Theorem results in other settings, trying to show the full possibility of what players could do in equilibrium. A more recent focus of the literature, starting with Blonski, Ockenfels, and Spagnolo (2011), has been on equilibrium selection by real subjects in these games. The goal of these papers is to produce sharper predictions of how these games are played.

In these experiments, random termination of the supergame has been used as a proxy for discounting with an infinite time horizon.[1] The justification for using random termination is that for a risk neutral agent, the two settings are isomorphic to one another, ex ante, if the discount rate is equal to the continuation probability. However, there is also a well documented experimental literature that shows that the majority of people are at least somewhat risk averse, and not risk neutral. This could potentially mean that subject behavior in randomly terminated supergames could be different than it would be if subjects played infinitely discounted supergames. This is the question this paper seeks to address. Do departures from risk neutrality cause subjects to behave differently in randomly terminated supergames compared to infinitely discounted supergames?

To address this question I first had subjects participate in a series of risk elicitation tasks. Multiple risk elicitation tasks were used because of the findings in Crosetto and Filippin (2016). They showed that different risk elicitation tasks produced different estimates of subjects' risk attitudes. Multiple tasks allows for a more well rounded estimate of subjects' risk attitude. A full description of these tasks are available in section 2.3.1.

After the risk elicitation tasks, subject participated in a series of repeated games

[1]The first repeated game run using random termination of the supergame was done by Roth and Murnighan (1978)

played with different, randomly selected partners. Subjects played 32 unique supergames varying the stage game,[2] discount factor,[3] and ending rule. The two ending rules used were standard random termination and a unique experimental design that captures some features of discounting not present in random termination. With this ending rule, subjects play a repeated game against another subject with all payoffs in the stage game matrix being discounted by the same fixed percentage in every period.[4] While this design does fail to capture some elements of infinite discounting, namely stationarity, that are captured by random termination, it does have other features of infinite discounting, namely certainty of experience, that random termination fails to capture. A full description of this treatment can be found in section 2.3.2.

Fréchette and Yuksel (2017) is the closest paper I could find in the literature to what I study here. They also look into the differences between random termination and discounting. Their paper focuses on divorcing the expected number of period subjects play from the discount factor used in the experiment. They do this with three novel treatments: discounting followed by a coordination game, discounting followed by random termination, and block random termination. They also have subjects play under the standard random termination treatment as a benchmark. They find that cooperation rates in the Prisoner's Dilemma were highest in the standard random termination treatment.

Notably, their experiment does differ from mine in a couple of crucial ways. First, I believe my design more directly implements a discounted treatment, as it ends endogenously based on subjects' choices and features no uncertainty outside

[2] A Prisoner's Dilemma, and a Battle of the Sexes game.

[3] All values shown on Table 2.2

[4] This is how Fréchette and Yuksel (2017) implement discounting in their two discounted treatments.

the strategic uncertainty present in any repeated 2 x 2 matrix game. Second, I vary the discount factor as opposed to using only one discount factor for the entire experiment. I do this because according to my theoretical model outlined in section 2, small changes to subjects' risk attitudes could result in different choices at a given discount factor. I use the wide range of discount factors to try and see these differences in decision making among subjects.

One other study also uses the same discounted treatment in this paper. Chapter 1 looks at the differences in how subjects play long randomly terminated repeated games compared to long discounted repeated games. Using a between subject design, subjects in that study participated in six repeated, Prisoner's Dilemma games using one of the two treatments to determine how the supergame ended, with $\delta = 0.98$. Using the Strategy Frequency Estimation Method (SFEM), that study found that subjects in the discounted treatment tended to play more defective strategies, but cooperators tended to play more forgiving strategies such as Tit-For-Tat. Subjects in the random termination treatment tended to be more cooperative, but the cooperators tended to play less forgiving strategies, such as Grim Trigger.

In this study, I find no significant differences, between the two ending rules, in which action subjects choose in the first period of each supergame. This holds true even when controlling for subjects risk attitudes. However, I do find differences in how subjects play these supergames going long. Subjects are more likely to converge to simple patterns in play in discounted supergames than they are in randomly terminated supergames. Risk attitude does not significantly impact this difference.

Figure 2.1: Generic Prisoner's Dilemma

	C	D
C	a, a	c, b
D	b, c	d, d

Notes: This is the game matrix of a standard Prisoner's Dilemma game if the following are true: $b > a > d > c$ and $2a > b + c$

2.2 Theoretical Background

Assume two players are playing a repeated 2×2 matrix game with an indefinite time horizon. An example of one of these types of matrix games is provide in figure 2.1. I will assume that each player seeks to maximize some strictly monotonic utility function $U(\cdot)$. I will further assume that $U(\cdot)$ is applied to their entire earnings stream rather than their earnings in each period. This is because each player cares about how much they earn from the entire supergame, and not necessarily their period by period earnings. This leads to the following being a player's ex ante utility for a supergame featuring discounted payoffs over an infinite time horizon

$$U(\sum_{t=0}^{\infty} \delta^t x_t)$$

where δ is the discount rate for that supergame and x_t is the payoff the subject earns in period t.

If the supergame instead features random termination, then each player's ex ante utility is

$$(1 - \delta) \sum_{t=0}^{\infty} \delta^t U(\sum_{k=0}^{t} x_k)$$

where δ is the continuation probability and x_t is the payoff the subject earns in period t.

At the start of each supergame players will have to choose which action to play in period 1. If they are playing a generic Prisoner's Dilemma game as in figure 2.1, they will prefer to play C in the first period if the discount rate for that game is $\delta \geq \delta^*$, where δ^* solves the following equation in a supergame featuring discounting over an infinite time horizon

$$U(\sum_{t=0}^{\infty} \delta^t a) = U(b + \sum_{t=1}^{\infty} \delta^t d) \tag{2.1}$$

Or instead the following equation if the supergame features random termination

$$(1-\delta) \sum_{t=0}^{\infty} \delta^t U(\sum_{k=0}^{t} a) = (1-\delta)[U(b) + \sum_{t=1}^{\infty} \delta^t U(b + \sum_{k=1}^{t} d)] \tag{2.2}$$

In both these equations it is assumed that the other player is playing a grim trigger strategy.[5]

To find this δ^* in a discounted supergame, note that $U(\cdot)$ was assumed to be strictly monotonic, and thus it has an inverse. This allows me to rewrite equation 2.1 as the following equation

$$\delta = \frac{b-a}{b-d} \tag{2.3}$$

Thus in discounted supergames, a players utility function does not impact the cutoff value of delta, δ^*, that dictates what their optimal first period action is.

In supergames featuring random termination of the supergame however, equation 2.2 can only be simplified

$$\sum_{t=0}^{\infty} \delta^t U(\sum_{k=0}^{t} a) = [U(b) + \sum_{t=1}^{\infty} \delta^t U(b + \sum_{k=1}^{t} d)]$$

[5]This assumption is standard in the literature. Grim Trigger is the strategy that establishes the lower bound for what δ values support cooperation as a subgame perfect equilibrium for the repeated Prisoner's Dilemma because it is the harshest possible punishment. See Dal Bó and Fréchette (2011) and Blonski, Ockenfels and Spagnolo (2011).

which directly depends on $U(\cdot)$ and not just the payoff parameters and δ. Thus subjects utility function directly impacts the cutoff value δ^*. One possible channel for this impact is a player's risk attitude.

Specifically, as players get more risk averse, the critical δ^* value needed to support cooperation as a subgame perfect equilibrium increases above the value of δ^* for the equivalent discounted supergame. This is because risk averse players have concave utility functions and thus decreasing marginal utility. This means that even though $a > d$, the difference between $U(\sum_{k=0}^{t+1} a) - U(\sum_{k=0}^{t} a)$ and $U(b + \sum_{k=1}^{t+1} d) - U(b + \sum_{k=1}^{t} d)$ gets smaller as t gets larger. Thus the differences between cooperation and defection in the early periods have a greater impact on the decision making process. Thus at the discounted value of δ^*, risk averse players prefer to defect. Therefore in order to make a risk averse player indifferent between cooperating and defecting, more "probability weight" needs to be placed on later periods, and less need to be put on earlier periods. This is done by increasing δ.

Similarly, as players get more risk loving, the critical δ^* for a random termination supergame decreases below the value for the discounted supergame. This is due to the fact that risk loving players have convex utility functions and thus increasing marginal utility. This means that the additional a payoffs in later periods from cooperation have a greater impact on utility than the d payoffs in later periods following a defection. Thus risk loving players strictly prefer to cooperate in the first period at the discounted δ^* value. Thus to make these players indifferent less "probability weight" needs to be placed on these later periods and more needs to be put on the earlier periods. This is done by decreasing δ. These two arguments lead to the following result.

Result 2.1. *As players get more risk averse (loving), the critical discount factor* *needed to support cooperation in a random termination supergame, δ^{RT}, becomes larger (smaller) than the critical discount factor needed to support cooperation in a discounted supergame, δ^{DISC}.*

Finally for risk neutral players, the cutoff value δ^* is the same in both settings.[6]

[7]

2.3 Experimental Design

This experiment was presented to subjects in two parts. The first part was a series of risk elicitation tasks, followed by a short quiz on those tasks. After all subjects completed the first part of the experiment, they were given instructions for the second part of the experiment which involved a series of repeated games that were either randomly terminated or featured discounted payoffs over an "infinite" time horizon. After the subjects finished playing all the supergames, they were paid in cash, in private, and were dismissed from the laboratory.

[6]If the player is risk neutral then their utility function is just an affine transformation of their payoff. Therefor equation 2.2 can be rewritten as $(1-\delta)\sum_{t=0}^{\infty}\delta^t\sum_{k=0}^{t}a = (1-\delta)[b+\sum_{t=0}^{\infty}\delta^t(b+\sum_{k=1}^{t}d)]$ which is equivalent to equation 2.3

[7]This result is in contrast to what Zeng, Li, Chen and Nan (2016) find. They show in their model that cooperation goes up in the repeated game as players get more risk averse. However, there are two stark differences between the Zeng et al. model and the model I present here. First is that in their model interaction between the two players happens over a finite number of periods. Secondly they assume subjects have a utility function of the form $u_i = \frac{\bar{p}_i}{1+\alpha_i \cdot r_i}$ where $\bar{p}_i$ is the subjects average per period earnings, α_i is their measure of risk attitude with $\alpha_i = 0$ denoting a subject being completely risk loving and $\alpha_i > 0$ denoting some amount of risk aversion, and r_i is the standard deviation of $\bar{p}_i$. They then run players with these preferences through an evolutionary learning model to see what type of strategies subjects with different risk attitudes settle on. An important thing to note about this model is that it takes an ex post perspective when evaluating subjects preferences, were as the model I present takes an ex ante perspective towards what to do in the first period of the supergame. This difference is enough to change what type of predictions arise from these two different models. Van Assen and Snijders (2010) also show this experimentally. Sabater-Grande and Georgantzis (2002) on the other hand show in an experiment that risk averse subjects are more likely to play the stage game NE in a repeated Prisoner's Dilemma.

2.3.1 Risk Elicitation Tasks

At the start of each session subjects were presented with three risk elicitation tasks to preform. These tasks included a multiple-price list, a la Holt and Laury (2002), The Eckel and Grossman (2008) lottery choice task, and the Bomb Risk Elicitation Task (BRET) by Crosetto and Filippin (2013). The exact specifications for these three tasks were adapted from Crosetto and Filippin (2016).[8] The multiple-price list and lottery choice list can be found on table 2.1. These three tasks were presented to subjects in a random order under the names RED, BLUE, and GREEN,[9] with all instructions appearing on subjects' computer screens along side the tasks themselves. After subjects went through all three tasks they answered a series of quiz questions relating to their payoffs from these tasks. At the end of the experiment subjects that correctly answered all the quiz questions where payed for the outcome of one randomly selected task. Subjects that missed one or more quiz questions were not paid for the risk elicitation tasks.

The reason three risk elicitation tasks were chosen instead of just using one risk elicitation task has to do with the findings in Crosetto and Filippin (2016). The authors show that the different risk elicitation tasks produce different estimates of subjects' risk attitudes simply due to the structural differences between the tasks. For this reason I include multiple tasks in order to get a more well rounded estimate of subjects' risk attitude instead of relying on one biased estimate.

In the Holt-Laury multiple price list, subjects are asked to chose which lottery they prefer for each of the 10 rows on the table. The specific table used in this experiment is shown on table 2.1a. Option A is always a relatively safe lottery,

[8]These specifications were used to ensure that a risk neutral subject would earn approximately $5 in expectation from the risk elicitation tasks.

[9]RED: Holt-Laury multiple price list. BLUE: BRET. GREEN: Eckel-Grossman lottery choice task. This was done to control for any sort of order effect influencing the results of these tasks.

Table 2.1: Risk Elicitation Tasks

(a) Multiple-Price List

	Option A				Option B			
1	$\frac{1}{10}$	\$4	$\frac{9}{10}$	\$3.20	$\frac{1}{10}$	\$7.70	$\frac{9}{10}$	\$0.20
2	$\frac{2}{10}$	\$4	$\frac{8}{10}$	\$3.20	$\frac{2}{10}$	\$7.70	$\frac{8}{10}$	\$0.20
3	$\frac{3}{10}$	\$4	$\frac{7}{10}$	\$3.20	$\frac{3}{10}$	\$7.70	$\frac{7}{10}$	\$0.20
4	$\frac{4}{10}$	\$4	$\frac{6}{10}$	\$3.20	$\frac{4}{10}$	\$7.70	$\frac{6}{10}$	\$0.20
5	$\frac{5}{10}$	\$4	$\frac{5}{10}$	\$3.20	$\frac{5}{10}$	\$7.70	$\frac{5}{10}$	\$0.20
6	$\frac{6}{10}$	\$4	$\frac{4}{10}$	\$3.20	$\frac{6}{10}$	\$7.70	$\frac{4}{10}$	\$0.20
7	$\frac{7}{10}$	\$4	$\frac{3}{10}$	\$3.20	$\frac{7}{10}$	\$7.70	$\frac{3}{10}$	\$0.20
8	$\frac{8}{10}$	\$4	$\frac{2}{10}$	\$3.20	$\frac{8}{10}$	\$7.70	$\frac{2}{10}$	\$0.20
9	$\frac{9}{10}$	\$4	$\frac{1}{10}$	\$3.20	$\frac{9}{10}$	\$7.70	$\frac{1}{10}$	\$0.20
10	$\frac{10}{10}$	\$4	$\frac{0}{10}$	\$3.20	$\frac{10}{10}$	\$7.70	$\frac{0}{10}$	\$0.20

Notes: The above is a Holt-Laury style Multiple-Price list risk elicitation table. For each row subjects choose which gamble they prefer, Option A or Option B. Subjects that are paid for this task receive the realization of the gamble they selected from one randomly chosen row. A risk neutral subject would choose Option A for rows 1-4 and Option B for rows 5-10.

(b) Lottery Choice Task

	Choice	Probability(%)	Outcome
1	A	50	\$4
	B	50	\$4
2	A	50	\$6
	B	50	\$3
3	A	50	\$8
	B	50	\$2
4	A	50	\$10
	B	50	\$1
5	A	50	\$12
	B	50	\$0

Notes: The above is an Eckel-Grossman style lottery choice task. Subjects in this task choose which of the five gambles they most prefer. Subjects paid for this task receive the realization of their chosen gamble. A risk neutral subject would choose gamble 5.

while option B is always a relatively risky lottery. The probability of the high payoff increases as one moves down the table, with the 10th row being a test of rationality. There subjects are simply asked if they prefer \$4 for sure or \$7.70 for sure. For the first four rows of the table, the expected value of option A is greater than the expected value of option B. This switches for the last six rows of the table. Thus a risk neutral subject would choose option A for the first four rows and option B for the last six rows. Subjects that were paid for this task were paid the realization of the lottery they chose on one randomly selected row.

For the Eckel-Grossman lottery choice task, subjects chose their preferred lottery out of a list of 5 lotteries. The lotteries used in this experiment are listed on table 2.1b. As one moves down the table the expected payoff of the lottery goes up, but so does the variance. Thus the most risk averse subjects will select lottery 1, while the least risk averse subjects choose lottery 5. Subjects who are risk neutral or risk loving also choose lottery 5.

In the BRET subjects are asked to collect any number of boxes they wish to out of 100. Boxes are collected in numerical order, so they are asked to chose a number between 0 and 100. For every box they chose to collect, they earn \$0.20. However one of the 100 boxes contains a mine, and each box is equally likely to contain the mine. If the subject collects the box that contains the mine, they instead earn \$0. In essence subjects are choosing their preferred lottery out of 101 total lotteries where each lottery is of the form

$$\begin{cases} \$0 & \frac{k}{100} \\ \$0.2k & \frac{100-k}{100} \end{cases} \tag{2.4}$$

where k is the number of boxes they chose to collect. The lottery with the highest expected value is $k = 50$, so this is the choice a risk neutral subject would make.

A risk averse subject would collect less than 50 boxes, while a risk loving subject would collect more than 50 boxes.

After a subject completed all three tasks they answered three quiz questions. For each quiz question subjects were told what the outcome of any random component of one of the three tasks was and then they were asked what their payoff was for that task. For example, in the lottery choice task, the quiz question would tell them whether choice A or B was randomly chosen, and remind them of what lottery they selected, and then they were asked to identify their payoff based on that information. Subjects answered these questions with a drop down menu. For Eckel-Grossman and Holt-Laury, all the payoffs from the tables were provided as an option for them to choose. For the BRET, subjects could choose between the number of boxes collected times $0.20, the number of the box that held the mine times $0.20, or $0.

I had subjects go through these quiz questions to see whether or not they actually understood the tasks . If subjects could not tell me how much they earned from these tasks after the realization of any randomness, it is likely they did not understand the tasks at all. If subjects did not understand the tasks then it makes the results of the risk elicitation tasks less reliable and increases the likelihood that the subjects will not be able to understand the repeated game section of the experiment either. In order to ensure that subjects took these quiz questions seriously, they were told that they would only be paid for the first three tasks if they completed all three quiz questions correctly prior to participating in the tasks. If they did so then they were paid for the outcome of one task chosen at random. They were told whether or not they got the quiz questions correct at the end of the experiment, after they completed all the repeated games.

Figure 2.2: Prisoner's Dilemma

	C	D
C	$10, 10$	$1, 17$
D	$17, 1$	$7, 7$

Notes: This is the Prisoner's Dilemma matrix used in the experiment. It was chosen so that $\delta^* = 0.7$. Neutral action labels were used in place of C and D for the experiment (U and D for the row player and L and R for the column player).

Figure 2.3: Battle of the Sexes

	A	B
A	$4, 4$	$8, 16$
B	$16, 8$	$4, 4$

Notes: This is the Battle of the Sexes matrix used in the experiment. The payoff ratio comes from Ioannou and Romero (2014) with the payoffs scaled up to be similar to those from the Prisoner's Dilemma Neutral action labels were used in place of A and B for the experiment (U and D for the row player and L and R for the column player).

2.3.2 Repeated Games

In the second part of the experiment, subjects participated in a series of repeated games. Each of these games featured one of the two stage games shown in figures 2 and 3.[10] This payoff specification for the Prisoner's Dilemma was chosen so that the critical δ^* in the discounted supergames, and in the random termination supergames when the subject is risk neutral, was equal to 0.7. $\delta^* = 0.7$ was chosen because I wanted the risk neutral threshold to be in the middle of the distribution of discount factors used in this experiment. This was done so I could see any potential differences in choices made by both risk averse subjects and risk seeking subjects.[11] The full list of δ values can be found on table 2.2. The ratio of the payoffs for the Battle of the Sexes game was taken from Ioannou and Romero (2014), with the payoffs scaled up so that they were comparable to the prisoner's dilemma payoffs. At the start of each supergame subjects were randomly matched with another sub-

[10]During the experiment the row player actions were labeled U and D, while the column player actions were labeled L and R during all supergames, regardless of payoff matrix. Additionally subjects always played as the row player to help prevent any confusion about which action was theirs and which was the other subject's.

[11]This threshold is not uncommon in the literature. Dal Bó and Fréchette (2011) use a similar threshold in one of their games.

Table 2.2: δ Values

Discounting	0.50	0.60	0.65	0.70	0.75	0.80	0.85	0.90
Random Termination	50	60	65	70	75	80	85	90

Notes: This is the full list of δ values used in this experiment. In the discounting supergames this number was referred to as Delta. In the random termination supergames it was referred to as the Threshold Number. The wide variety of δ values was chosen in order to help find variations in cooperation rates due to variations in subjects' risk attitudes.

ject to play these games. During each period of play subjects were asked to chose which action they prefer to play, as well as which action they believed the other player would play in that period.[12] After both subjects made their choices, they were then shown the outcome of that period, and their payoffs for that supergame were updated appropriately.

Each supergame featured either random termination of the supergame, or discounting with an "infinite" time horizon. In the random termination treatments all matches between subjects progressed together period by period. Once every subject made their decision, the experimenter rolled two ten sided dice, one representing the tens place and the other the ones place of a two digit number. If the number rolled was less than or equal to 100 times the δ value[13] used in that supergame, then that supergame continued for at least one more period. If instead the number rolled was strictly greater than 100 times δ, then that supergame ended. Subjects were reminded of this rule on their computer screen after each period. The choice to roll dice as opposed to have the computer randomly generate whether or not the supergame continued, or predrawing the number of periods was made in order to make the uncertainty about the length of the random termination supergames more salient. I wanted to ensure that subjects knew that these supergames ended ran-

[12]Beliefs were not incentivized in this experiment. They were apart of the original experimental code and were not removed. They were later decided to be used in analysis. Those results should be taken with a grain of salt.

[13]In the instructions this was referred to as the Threshold Number

domly, while the discounted supergames did not. Therefore I wanted to make the randomness as visible as possible. Thus dice were used to determine when the random termination supergames ended.

In the supergames with discounting over an "infinite" time horizon subjects played a supergame as it is traditionally described in theory. In each period of the supergame the payoffs in the game matrix were reduced by a fixed percentage δ and an infinite number of periods were played, where subjects payoffs were the sum of their earnings in the infinite number of periods. However, one departure from theory was made. Since I could not actually keep subjects in the lab to play an infinite number of periods, I gave the subjects the option to opt out of continuing to play the supergame after each period. Additionally, the payoffs from a discounted game get so small they fail to satisfy dominance, as defined in Smith (1982). Since the exact point in time this is true will vary from subjects to subject, the opt out option is the cleanest way to give the subjects control over this, while still maintaining an indefinite time horizon. In exchange for opting out, subjects were given a continuation payoff to compensate them for all the periods that they could have played but did not.

This continuation payoff was equal to the infinitely discounted sum of the minimum payoff[14] the subject could receive from the payoff matrix in the period the supergame ended. Subjects were made aware of the existence of the continuation payoff in the instructions, and were informed about exactly what the continuation

[14]The minimum payoff was chosen for two reasons. First, I wanted to ensure that the continuation payoff did not influence the decisions subjects made during play. This is why the discounted value of subjects' average payoff was not used. Second, I wanted as strong a distinction as possible between the random termination supergames and the discounted supergames. This is why I did not just give them the infinitely discounted sum of a random payoff from the matrix in the period the supergame ended. The maximum payoff could not be selected, because then subjects would be incentivized to opt out immediately and receive their highest possible payoff. Thus the minimum payoff was chosen.

payoff was in each period before they chose whether or not to opt out. Additionally, subjects were told what the maximum amount they could possibly earn was if the continued to play. For example, if subjects were playing with the Prisoner's Dilemma stage game and $\delta = 0.8$, then their they would be told at the end of period 5 that their maximum remaining payoff would be 34.82 and that their continuation payoff in the match if the match ended in that period would be 2.05. I provided them with this information to further assist them in making their decision about when to opt out.

Lastly, these discounted repeated games lasted until both subjects chose to opt out. This choice was made in order to help maintain certainty of experience. If the supergame ended after the first subject chose to opt out then half the subjects would not know when their supergame would end. Since uncertainty of the exact game length is a feature of random termination that this treatment was designed to avoid, it was decided to end the supergame after both players chose to opt out.[15] This meant that in many instances, subjects continued manually making decisions even after they chose to opt out.

In addition to the two stage games and the two ending rules, subjects also played these supergames with a wide variety of δ values. All eight values used in this experiment are listed on table 2.2.The table presents these values in the same way they were presented to the subjects in both the discounted supergames and the randomly terminated supergames. The reason I used eight distinct δ values was because minor variations in subjects' risk attitude could potentially lead to swings in whether or not they would consider cooperating in the prisoner's dilemma.

For example, if we assume subject's utility functions takes the form $U(x) = x^r$

[15]This choice does still introduce uncertainty of when the supergames ends for the subject that chose to opt out first in the periods after they chose to opt out. For this reason, only the periods prior to either subject opting out was considered in the analysis of the results of this experiment.

where r, the coefficient of risk attitude, is some real valued number,[16] then a subject who is indifferent between cooperation and defection against a grim trigger strategy in a random termination supergame when $\delta = 0.80$ would have $r = 0.319$. This is equivalent to saying that they would choose lottery number two in Eckel-Grossman or collect 24 boxes in the BRET. If instead they are indifferent when $\delta = 0.75$, then they would have $r = 0.656$ which implies choosing lottery number four in Eckel-Grossman or collecting 40 boxes in the BRET. These two subjects are quite different even though their cutoffs are somewhat close. Thus this variety in δ values is used to help show changes in behavior more fluidly than a courser set of δ values would allow.

The two payoff matrices, two ending rules, and eight δ values result in 32 distinct treatments. In each session all subjects participate in each treatment once. These treatments were presented in a random order in an attempt to control for any order effects.

2.3.3 Hypotheses

The model described in section 2.2 provides the basis for the main hypothesis of this paper. This hypothesis has to do with cooperation in the Prisoner's Dilemma.

Hypothesis 2.1. *Amongst Risk Averse Subjects, for each δ, first period cooperation rates will be higher in discounted supergames than they will be in random termination supergames.*

From result 1 we know that risk averse agents require higher δ values in order to want to cooperate in a randomly terminated repeated Prisoner's Dilemma. Thus, more cooperation is expected in the discounted supergames.

[16]For this utility specification a subject is risk averse when $r < 1$, risk neutral when $r = 1$, and risk loving when $r > 1$.

For the Battle of the Sexes, no sharp predictions exist. This is because any patterns of play that could emerge involve playing stage game Nash Equilibria. Thus these patterns can be supported for any value of δ. This means that subjects' risk attitudes and the supergames ending rule should have no impact on their decision making. However, there could be differences between how subjects view discounted supergames and random termination supergames other than their risk attitude. Therefore, the Battle of the Sexes supergames will be used as a bench mark to see if any unaccounted for differences exist.

2.4 Results

A total of seven session of this experiment were run at the Economics Science Laboratory at the University of Arizona. Each session had between 8 and 14 subjects. On average a session lasted approximately two hours and subjects were payed an average \$19.12, with a maximum payment of \$46.40 and a minimum payment of \$6.40. Subjects were paid for what they earned in one randomly selected supergame out of they 32 played, with an exchange rate of 3 experimental units to \$1, any earnings they received for the risk elicitation tasks, as well as a \$5 show up fee. 68% of subjects successfully completed the quiz and were paid for the outcome of one randomly selected risk elicitation task. Due to computer errors, no usable data was generated by either of the first two sessions. Thus all data in this paper comes from 50 subjects who participated in the last five sessions.

Table 2.3 shows the summary statistics from the three risk elicitation tasks the subjects participated in. For the Eckel-Grossman and Holt-Laury tasks, these findings are similar to findings elsewhere in the literature. For the BRET however, subjects in this study chose on average between six and eight fewer boxes than pre-

Table 2.3: Summary Statistics: Risk Elicitation Tasks

	Eckel-Grossman (EG)	Holt-Laury (HL)	BRET
Mean	3	5.64	34.8
SD	1.40	1.61	18.06
Max	5	10	87
Min	1	2	1
Median	3	5.5	32.5

Notes: For EG the statistic is which lottery they chose. For HL it is the number of times they chose the relatively safe lottery. For the BRET it is the number of boxes they chose to collect. These estimates are similar to those typically reported in the literature, with the exception of the BRET. For that task subjects in this experiment report being more risk averse.

Table 2.4: Risk Elicitation Correlations

r Value Estimates		
	BRET	HL
EG	-0.100	0.376**
HL	0.047	

Notes: **: significant at the 5% level after applying the Bonferroni Multiple Hypotheses Correction. Holt-Laury was measured in the number of safe choices. More safe choices is more risk averse, so negative correlation with HL implies risk aversion goes up with both simultaneously for the Raw Measures.

viously reported.[17] Table 2.4 provides evidence for why multiple risk elicitation tasks are necessary to get a well round view of subjects risk attitudes. None of the three tasks are highly correlated with one another. To address this, these measures were transformed into the measure of risk aversion, r, for the utility function $U(x) = x^r$.[18] These measure were then averaged to get a pooled estimate of subject's risk attitude.[19] Summary statistics for the r measures for each task and the pooled estimate can be found on table 2.5.

[17]In Crosetto and Filippin (2013) subjects collected on average 46 boxes, while in Crosetto and Filippin (2016) the average was around 40 boxes collected.

[18]Thresholds from Crosetto and Filippin (2016) were used to determine the r values for the subjects in this study.

[19]In addition to doing a straight average, a weighted average that minimized the squared distance between the average and the individual components was also examined. The weighted average that minimized this squared distance was simply using the BRET Measure. Using that measure resulted in categorizing 43 subjects as risk averse, as opposed to the straight average where 47 were categorized as risk averse. This difference is small enough that the straight average was used.

Table 2.5: r Values for $U(x) = x^r$

	Eckel-Grossman (EG)	Holt-Laury (HL)	BRET	Pooled Average
Mean	0.148	0.578	0.702	0.476
SD	0.661	0.401	0.910	0.413
Median	0.474	0.590	0.447	0.459

Notes: For all four columns, the estimated r values is the statistic reported. The Pooled Average is a straight, unweighted average of the first three measures.

An individual with this utility functions is considered to be risk averse if $r < 1$, risk neutral if $r = 1$, and risk seeking if $r > 1$. In this subject pool, both the mean and median subject are risk averse. Furthermore, 47 of the 50 subjects that participated in this experiment are risk averse based on their pooled estimate of r.[20] Thus for the remainder of the section, all data on subjects choices will be presented as a whole, instead of broken down by risk attitude. Appendix A.2 has more information about the risk a characteristics of the subjects in this study for the interested reader.

2.4.1 Actions

Table 2.6 shows first period behavior for both stage games, while table 2.7 shows subject behavior across all periods.[21]

For the Prisoner's Dilemma each cell shows the relevant cooperation rate, while for the Battle of the Sexes it shows how often subjects played the Strategy in Their Preferred Equilibrium Outcome (STPEO). The rates are shown for all δs taken together and broken down by each δ individually. Additionally, figures 2.4 and 2.5 show CDFs of subjects' first period behavior for both the Prisoner's Dilemma and

[20]A wider range of r values was considered when looking at whether or not someone was risk neutral when determining this 47 number. A subject was considered risk neutral if their pooled estimate for r was between 0.95 and 1.05. None of the subjects had a risk neutral estimate with this increased range.

[21]In the discounting supergames, only periods prior to the subject opting out were consider for the All Period cooperation and STPEO rates.

Table 2.6: Summary of First Period Behavior

(a) Cooperation in the Prisoner's Dilemma

	Overall	Disc	RT
Overall	21.4%	23.0%	19.8%
$\delta = 0.5$	17.0%	18.0%	16%
$\delta = 0.6$	19.0%	20.0%	18.0%
$\delta = 0.65$	22.0%	22.0%	22.0%
$\delta = 0.7$	28.0%	26.0%	30.0%
$\delta = 0.75$	19.0%	20.0%	18.0%
$\delta = 0.8$	19.0%	22.0%	16.0%
$\delta = 0.85$	23.0%	22.0%	24.0%
$\delta = 0.9$	24.0%	34.0%	14.0%

(b) STPEO in the Battle of the Sexes

	Overall	Disc	RT
Overall	78.8%	79.5%	78.0%
$\delta = 0.5$	78.0%	74.0%	82.0%
$\delta = 0.6$	81.0%	80.0%	82.0%
$\delta = 0.65$	81.0%	88.0%	74.0%
$\delta = 0.7$	80.0%	82.0%	78.0%
$\delta = 0.75$	77.0%	80.0%	74.0%
$\delta = 0.8$	77.0%	76.0%	78.0%
$\delta = 0.85$	76.0%	76.0%	76.0%
$\delta = 0.9$	80.0%	80.0%	80.0%

Notes: Significance determined by a Wilcoxon Rank-Sum test with the Bonferroni Multiple Hypotheses Correction. Overall in the top row is the relative rate across both treatments. Overall is the first column is the relative rate across all δs. Disc is for the discounting supergames. RT is for the random termination supergames. The unit of observation is a subjects average first period cooperation rate for the given parameters in the Prisoner's Dilemma or average first period STPEO rate in the Battle of the Sexes.

the Battle of the Sexes respectively. Comparing these results to the prediction from section 2.3.3, one can see that hypothesis 2.3.3 is soundly rejected. In only one instance, $\delta = 0.9$, is the cooperation rate in discounted Prisoner's Dilemma supergames higher than they are in random termination supergames in the first period.[22] Similarly for the all period cooperation rates, the only significant difference observed was for $\delta = 0.7$. This observation had random termination greater than discounting, the opposite of what was observed in the first period cooperation rates.

There is also only one δ value for which the rate that subjects play the STPEO is higher in random termination Battle of the Sexes supergames than it is in the discounted supergames, and that was $\delta = 0.5$ across all periods. For all other δs across all periods, and all δs in the first period, the rate of cooperation or STPEO play is statistically indistinguishable between random termination supergames and discounted supergames.

[22] Wilcox-rank sum tests were used to check for significant differences.

Table 2.7: All Periods Behavior

(a) Prisoner's Dilemma

	Overall	Disc	RT
Overall	18.3%	19.0%	17.1%
$\delta = 0.5$	16.4%	17.2%	15.6%
$\delta = 0.6$	14.3%	12.7%	18.1%
$\delta = 0.65$	16.8%	17.3%	23.1%
$\delta = 0.7$	23.6%	18.7%	27.4%
$\delta = 0.75$	16.6%	17.7%	15.7%
$\delta = 0.8$	19.1%	18.5%	18.6%
$\delta = 0.85$	17.7%	18.9%	20.8%
$\delta = 0.9$	17.6%	21.3%	14.5%

(b) Battle of the Sexes

	Overall	Disc		RT
Overall	72.1%	72.5%		72.8%
$\delta = 0.5$	76.0%	71.4%	$<^{**}$	83.0%
$\delta = 0.6$	72.8%	77.6%		74.8%
$\delta = 0.65$	73.7%	75.3%		73.0%
$\delta = 0.7$	78.1%	79.7%		77.0%
$\delta = 0.75$	69.5%	69.3%		68.8%
$\delta = 0.8$	74.6%	74.8%		77.3%
$\delta = 0.85$	72.3%	71.5%		72.5%
$\delta = 0.9$	71.7%	70.9%		74.1%

*Notes:*** significant at the 5% level. Significance determined by a Wilcoxon Rank-Sum test with the Bonferroni Multiple Hypotheses Correction. Overall in the top row is the relative rate across both treatments. Overall is the first column is the relative rate across all δs. Disc is for the discounting supergames, where all periods prior to the subject opting out were considered. RT is for the random termination supergames, where all periods of the supergame were considered. The unit of observation is a subjects average cooperation rate across all periods for the given parameters in the Prisoner's Dilemma or average STPEO rate across all periods in the Battle of the Sexes.

Figure 2.4: First Period Cooperation PD Overall

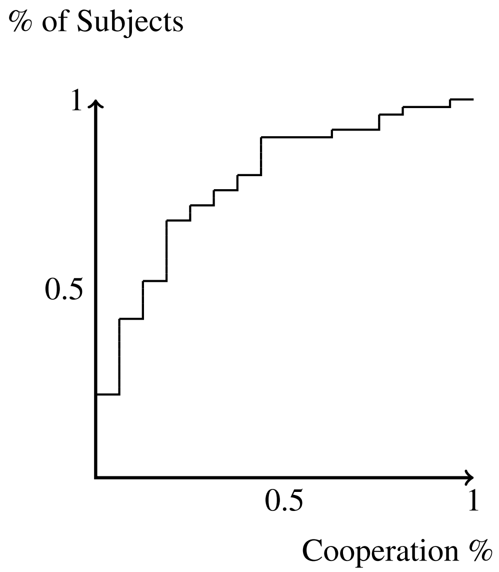

Figure 2.5: First Period STPEO Play BoS Overall

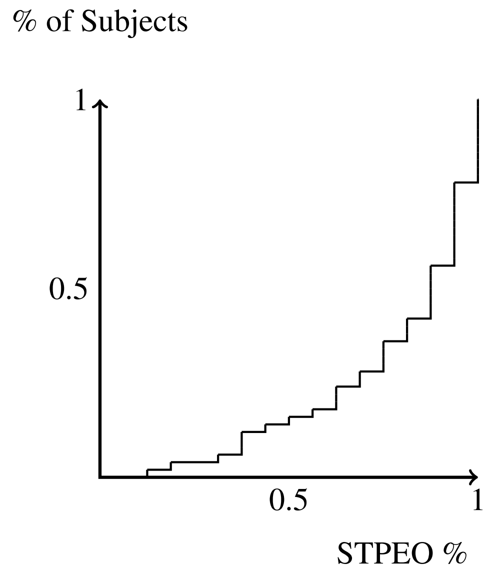

Notes: Figure 2.4 shows the CDF of first period cooperation rates in the Prisoner's Dilemma. The unit of observation is the cooperation rate of a subject in the first period across all Prisoner's Dilemma supergames. Figure 2.5 shows the CDF of first period STPEOs rates in the Battle of the Sexes. The unit of observation is the STPEO rate of a subject in the firs period across all Battle of the Sexes supergames.

In addition to examining differences in first period cooperation and STPEO play rates within δs, I also wanted to see how these rates varied across δs.[23] To do this I ran logit regressions with a dummy for action choice[24] as the endogenous variable on a dummy for treatment type[25] and a dummy variable for each δ-value, with $\delta = 0.50$ as the baseline. I ran these regression both with just these exogenous variables and also controlling for subjects risk attitudes and whether or not they correctly answered all three quiz questions. Risk attitudes were controlled for by using the BRET r value estimates. The BRET was used because the r value estimates from the BRET were most similar to what has been typically observed in the literature, and thus would be the most direct comparison to a typical subject population. Additionally this measure was interacted with the Treatment Type dummy. This was done because changes in Treatment Type is what is supposed to drive the impact of risk attitude. The regressions were run separately for the Prisoner's Dilemma supergames and the Battle of the Sexes supergames. Results from the regressions can be found on table 2.8.

Unsurprisingly, the treatment type that subjects played with had no impact on cooperation rates or how often subjects played the STPEO in the first period of a supergame. Changes in δ however did impact cooperation rates in the first period of the Prisoner's Dilemma supergames. Subjects cooperation rates were roughly eight percentage points higher for supergames with $\delta = 0.90$. No other δ value was significantly different from $\delta = 0.50$. Changes in δ did not have any statistically significant impact on how often subjects played the STPEO in the first period of a

[23]Regressions were only run for first period action choice. This is because after the first period, the actions of the other subject a subject is matched with has a potential influence on action choice, not just the parameters of the supergame.

[24]For the Prisoner's Dilemma, this dummy was equal to 1 if the subject chose to cooperate and 0 otherwise. For the Battle of the sexes, this dummy was equal to 1 if the subject chose the STPEO and 0 otherwise.

[25]Discounted supergames were 1, random termination supergames were 0.

Table 2.8: Logit Regression Marginal Effects

	PD	PD Risk	BoS	BoS Risk
Treatment Type	0.033	0.057	0.015	-0.007
	(0.051)	(0.061)	(0.022)	(0.043)
Delta=0.9	0.079**	0.079**	0.020	0.020
	(0.037)	(0.035)	(0.035)	(0.035)
Paid		0.064		-0.069
		(0.061)		(0.064)
Risk Est. * Treatment		-0.038		0.033
		(0.024)		(0.035)

Notes: ** significant at the 5% level. Standard errors were two-way clustered by subjectID and sessionID. PD is the Prisoner's Dilemma, where the dependent variable was a dummy variable for first period action choice equal to 1 for cooperate and 0 for defect. BoS is the Battle of the Sexes, where the dependent variable was a dummy variable for first period action choice equal to 1 for STPEO and 0 for the other action. Treatment Type was a dummy variable equal to 1 for discounting and 0 for random termination. Paid was a dummy variable equal to 1 if the subject successfully passed the pre-experiment quiz, and 0 if they did not. Risk Est. is the estimated r value coming from the BRET task.

Battle of the Sexes supergame. Finally controlling for subjects risk attitudes did not have any noticeable impact on how the ending rule or δ effected these rates.

2.4.2 Patterns of Play

In addition to looking at how subjects approached the first period of play, I also looked to see if any sort of patterns emerged in these supergames. This is used as a further check for differences between the random termination supergames and the discounted supergames. Because the δ values are relatively small, the lengths of many of the supergames played in this experiment were exceedingly short. For this reason patterns are defined in the following way:

1. In the first period of the sequence, the first outcome of the stage game in the pattern was played.

2. In the second period of the sequence, the second outcome of the stage game

Table 2.9: Existence by Stage Game

	Overall	Disc		RT
Prisoner's Dilemma	70.5%	80.7%	$>^{***}$	57.2%
Battle of the Sexes	79.1%	93.4%	$>^{***}$	61.5%

Notes: *** significant at 1% level. Significance determined by a Wilcoxon Rank-Sum test with the Bonferroni Multiple Hypotheses Correction. Disc is for the discounting treatment. RT is for the random termination treatment. Overall is the rate for both treatments combined. Each entry is the percentage of supergames in that category where at least one pattern was observed. Such supergame is said to have converged to a pattern.

in the pattern was played.

3. Additionally in the second period, the subjects' beliefs about the other subject's action were the same as the action the other subject actually took.

This specification means that I am looked for two period patterns. The specific patterns I looked for depended on the stage game:

- For the Prisoner's Dilemma the two patterns I looked for were both subjects played cooperate in both periods or both subjects played defect in both periods.

- For the Battle of the Sexes I looked for four patterns. These were settling on one Nash Equilibrium of the stage game for both periods, alternating between the two Nash Equilibria of the stage game, or both subjects playing the STPEO action in both periods.

Overall, the patterns emerged in 59.6% of all supergames.[26] This statistic includes random termination supergames that ended after the first period, and discounted supergames where one of the subjects opted out after the first period. By definition such a supergame could never converge to a pattern, since at least two periods are

[26] In 71.5% of all periods subjects reported beliefs matched the action the other subject they were matched with chose.

Table 2.10: Prisoner's Dilemma Patterns

	Overall	Disc	RT
Cooperation	8.2%	6.8%	10.7%
Defection	93.8%	95.9%	89.9%
Multiple	2.0%	2.7%	0.6%

Notes: Significance determined by a Wilcoxon Rank-Sum test with the Bonferroni Multiple Hypotheses Correction. Disc is the discounting treatment. RT is the random termination treatment. Overall is both treatments. Cooperation means a cooperation pattern emerged. Defection means a defection pattern emerged. Multiple means both patterns emerged. Each percentage is the percentage of supergames that converged that contain at least one instance of the appropriate pattern.

needed to define a pattern. After removing those supergames from the data, patterns emerged 74.8% of the time.[27] Table 2.9 further breaks down the existence by stage game and ending rule. Patterns were more likely to emerge in a Battle of the Sexes supergame as compared to a Prisoner's Dilemma supergame. Patterns were also more likely to emerge in discounted supergames as compared to random termination supergames, regardless of game matrix.

Tables 2.10 and 2.11 show how often each pattern occurred in the supergames where patterns emerged. For the Prisoner's Dilemma, defection patterns were significantly more likely to arise in discounted supergames than they were in random termination supergames. Supergames where both patterns emerged were also more likely to occur in discounted supergames compared to random termination supergames. However, none of the differences between the two treatments are significant.

For the Battle of the Sexes, the differences were even greater. Alterations occurred in a greater percentage of random termination supergames compared to discounted supergames. Discounted supergames, by comparison have a higher percentage of supergames where patterns where only one player's preferred equilib-

[27]The percentage of the time subjects reported correct beliefs remained the same.

Table 2.11: Battle of the Sexes Patterns

	Overall	Disc		RT
Alternation	32.1%	29.1%		37.5%
Column's Preferred	18.9%	22.4%	$>^{*}$	12.5%
Row's Preferred	15.7%	18.7%		10.2%
Fight	60.8%	65.0%	$>^{*}$	52.8%
Multiple	25.9%	32.8%	$>^{***}$	13.1%

Notes: *** significant at 1% level. * significant at 10% level. Significance determined by a Wilcoxon Rank-Sum test with the Bonferroni Multiple Hypotheses Correction. Disc is the discounting treatment. RT is the random termination treatment. Overall is both treatments. Alternation is the pattern where subjects play one equilibria in the first period then the other equilibrium in the next period emerged. Column's preferred is the pattern where they play the column player's preferred equilibrium emerged. Row's preferred is the pattern where they play the row player's preferred equilibrium emerged. Multiple means more than one pattern emerged. Fight is when each play plays the STPEO action each period. Each percentage is the percentage of supergames that converged that contain at least one instance of the appropriate pattern.

rium was played. This difference was significant for the Column player's preferred equilibrium. Fights, each subject playing the STPEO in both periods, were also significantly more common in discounted supergames. This is likely a driving force in why Multiple patterns existing within supergames was more common in discounted supergames. Subjects would start by fighting and then move to one of their preferred equilibria.

To determine what influenced emergence of patterns, I ran logit regressions with a dummy for emergence as the endogenous variable.[28] The exogenous variables were a dummy for treatment type,[29] a dummy for each δ value with $\delta = 0.50$ as the baseline, and the number of periods the match lasted. The reason I included the last exogenous variable was because the discounted supergames tended to last much longer than the random termination supergames, as can be seen on table 2.12. Because longer supergames are more likely to contain a pattern, I did not want this effect to be attributed ending rule when in reality it was just the fact that the

[28] 1 if a pattern emerged in that supergame, 0 otherwise.
[29] Same as before.

Table 2.12: Summary Statistics: Supergame Length

	Overall	Disc	RT
Mean	11.01	17.33	4.69
SD	11.89	13.18	5.42
Max	78	78	28
Min	1	2	1
Median	8	14	3

Notes: Overall provides the statistics for both treatments combined. Discounting shows the summary statistics on supergame length for the discounting treatment and Random Termination does the same for the random termination treatment. In all measure, discounting supergames are longer than random termination supergames.

supergames were longer. As before I ran two specifications for each stage game, one controlling for risk attitude and whether or not the subject was paid for the first part of the experiment, and another without these controls. The results of these two regression are on table 2.13.

There are notable differences in the significant parameters between the two stage games. For the Battle of the Sexes, both the Treatment Type, and the Supergame Length are significant. This makes sense, because longer supergames lead to more opportunities for patterns to exist, and discounted games also naturally lead to longer supergames. Things are a somewhat different for the Prisoner's Dilemma supergames. In the model without the risk controls, both Treatment Type and the dummy for $\delta = 0.75$ are significant. Since none of the other δ dummies are significant, this result is likely just an outlier. In the model with risk controls however, the coefficient for Treatment Type is not significant, but the interaction between Treatment Type and the estimate of subjects risk attitude is. This coefficient implies that as subjects get more risk seeking, they are more likely to play a pattern in discounted supergames. This provides some evidence for the theory model presented in section 2.2, that changes in risk attitude cause differences in behavior between discounted and randomly terminated supergames. However, since the Treatment

Table 2.13: Existence Regressions Marginal Effects

	PD	PD Risk	BoS	BoS Risk
Treatment Type	0.185***	0.116	0.137**	0.111**
	(0.057)	(0.072)	(0.069)	(0.047)
Delta=0.75	-0.104***	-0.114**	-0.005	-0.0003
	(0.035)	(0.046)	(0.097)	(0.092)
Supergame Length	0.007	0.007	0.017***	0.016***
	(0.006)	(0.005)	(0.001)	(0.003)
Paid		-0.098		0.020
		(0.088)		(0.031)
Risk Est. * Treatment		0.106*		0.036
		(0.054)		(0.074)

Notes: *** significant at the 1% level. ** significant at the 5% level. * significant at the 10% level. Standard errors were two-way clustered by subjectID and sessionID. PD is for the Prisoner's Dilemma. BoS is for the Battle of the Sexes. The dependent variable for all regressions was a dummy variable equal to 1 if the supergame converged to a pattern and 0 otherwise. Treatment Type was a dummy variable equal to 1 if it was a discounting supergame and 0 otherwise. Paid was a dummy variable equal to 1 if the subject successfully passed the pre-experiment quiz, and 0 if they did not. Risk Est. is the estimated r value coming from the BRET task.

Type dummy is still significant for the Battle of the Sexes regressions, risk attitude is not the only channel that is causing differences between the two treatments.

2.5 Conclusion

The goal of this study was to determine if departures from risk neutrality caused subjects to make different decisions in repeated games that are randomly terminated compared to those that are infinitely discounted. To do this I first had subject participate in a series of risk elicitation tasks to get a well round picture of their risk attitude. I then had them play a series of repeated games where the discount factor, stage game, and ending rule was varied to see their choices in a variety of situations. The two ending rules used was standard random termination and a new treatment designed to capture features present in theoretical discounting models that are not present in random termination. In this new treatment, the payoffs in the stage game are discounted every period and subjects play for an indefinite number of periods. Subjects are given the option to opt out of continuing to play these supergames. This opt out both allows these supergames to end in a finite amount of time, and ensure that these supergames satisfy dominance (Smith 1982).

I find no difference in first period action choice across the two ending rules in both the Prisoner's Dilemma and the Battle of the Sexes. This is true without controlling for subjects' risk attitudes and with controlling for them. I do find a difference in how often simple patterns emerge across the two treatments. These simple patterns formed more often in supergames featuring the new discounted treatment. This difference holds even when controlling for the length of the supergame. Differences in risk attitude lead to a difference in how often patterns existed in Prisoner's Dilemma supergames, but do not explain the treatment differences in the Battle of the Sexes. These results lead me to believe that subjects' risk attitude is not the sole influence of behavioral differences between randomly terminated repeated games and "infinitely" discounted ones.

Returning to the model presented in section 2.2, it is important to note that δ^* is the same between random termination and discounting when $U(\cdot)$ is an affine function, and not specifically when subjects are risk neutral. Thus, other potential transformations of $U(\cdot)$ could be what is causing the differences between these two treatments, and not subject's risk attitudes. These channels could potentially be higher order risk attitudes, such as prudence, models of noisy decision making, or some other channel. Future research is need to determine if the theoretical model in section 2.2 is an inaccurate description of how subjects think about repeated games, or if a channel other than risk attitude is driving the potential differences found there.

Chapter 3

Multi-Media Learning and Experimental Economics Instructions

3.1 Introduction

Instructions in experimental economics are like scientific instruments in other disciplines. While typically not the focus of study themselves, they are necessary tools to conduct experiments. Some previous work has explored how experimental instructions should be written, but little work has focused on making sure the quality of instructions are high enough to ensure subjects understand the decision environments they are placed in. This study borrows techniques from the Multi-Media Learning literature in psychology, which is well summarized by Mayer and Fiorella (2014), to devise a new set of experimental instructions for the Becker-DeGroot-Marschak, henceforth BDM, mechanism. The goals of Multi-Media Learning is to take advantage of multiple learning channels, typically audio and visual, to increase the amount of information subjects can learn at once. This two channel approach is coupled with techniques to reduce extraneous cognitive processing. High ability subjects, as measured by Cognitive Reflection Test (CRT) scores, that received these instructions had a higher understanding of the BDM mechanism than those that received standard experimental instructions.

Two sets of experimental instructions were adapted from the instructions used in Cason and Plott (2014). One of these instructions were as close to theirs as possible, with departures due to differences in experimental design. These instructions were used as Standard, baseline, experimental economics instructions. The other set of instructions used several of the techniques in Mayer and Fiorella (2014) to help reduce cognitive load. The specifics of how these instructions were put together can be found in section 3.2.1. The instructions from Cason and Plott (2014) were used because they were the simplest recent instructions used for the BDM mechanism.

The BDM mechanism was chosen as this experiment's primary decision environment for several reasons. One, the primary purpose of the BDM mechanism is to measure subjects Willingness to Pay or Willingness to Ask for goods. The optimal action is to report these values truthfully. Therefore if subjects have induced valuations, as they do in this experiment, they then have a unique, optimal action that is known by the experimenter. This makes it a perfect task to measure subject understanding. Two, the BDM mechanism is notoriously difficult for subjects to understand.[1] This makes it even easier to see differences in understanding levels across different treatments, since understanding levels are quite low to begin with.[2] Three, the BDM is a useful tool for researchers trying to determine how sub-

[1]See Healy (2018) and Burfurd and Wilkening (2018) for two recent examples of trying to improve teaching the BDM mechanism to subjects.

[2]In Cason and Plott (2014) subjects chose the optimal action between 10% and 44% of the time depending on timing and treatment.

jects value different things. If the mechanism is well understood, it provides better estimates for these valuations compared to similar tools, such as second price auctions. Thus, discovering any techniques to improve subject understanding of this mechanism is worth pursuing. The specifics of the BDM mechanism used in this experiment are in section 3.2.2.

In addition to two sets of instructions, subjects in this experiment received one of two pre-experiment quizzes. In half the treatments subjects received an Incentivized pre-experiment quiz. For these quizzes, subjects were paid $0.50 for each question they answered correctly. In the other treatments, subjects received an Unincentivized quiz. The Unincentivized quiz included the same questions as the Incentivized quiz, but subjects were not paid for their correct answers. Incentivized quizzes make incorrect answers more salient to subjects, which may in turn improve understanding in the main experiment. Including an incentivized quiz also creates situations where subjects receive multiple reinforcement methods ahead of the main experiment, which Freeman et al. (2018) shows improves understanding.

Most previous work on experimental instructions have focused on the content of the instructions. Many studies have looked into the affect of using instructions that are contextualized in some way, as opposed to instructions that describe an abstract decisions environment.[3] Other papers have looked more directly into how instructions are written. During the literature review for this project, three such papers were found. First, Bigoni and Dragone (2012) looked at the impact different ways of presenting instructions had on understanding in public goods games. They found that shorter instructions reduced subject understanding, based on the number of pre-experiment quiz questions answered incorrectly, unless they were coupled

[3] see Alekseev, Charness, and Gneezy (2017) for an excellent survey of the literature on contextual instructions.

with forced input and thus, active participation by the subjects.

Second, Ramalingam, Morales, and Walker (2018) also looked into the impact of instruction length on subject understanding in public goods games. Borrowing from Chandler and Sweller (1991), they argue that their longer instructions, with more worked out examples and an emphasis on the positive externalities present in public goods games, helps to reduce the germane (effective) cognitive load on subjects. Using decision making time as a measure of understanding, they find that subjects that received the longer instructions had greater understanding of public goods games.

Third, Freeman et al. (2018) use a large survey[4] to show the wide variation in how instructions and other pre-experiment tasks are used in experimental economics. They then used a variety of pre-experiment tasks to determine what most improved subject understanding in a unique, experimental environment of their own design. They found that combining multiple reinforcement methods improved understanding, with an incentivized pre-experiment quiz, going through the instructions twice, and providing paper instructions in addition to on screen instructions doing the best job of improving understanding of their task.

In this study, high ability[5] subjects who received Multi-Media Learning instructions had a better understanding of the BDM mechanism, as determined by the rate at which they chose the optimal action, than those that received Standard instructions. This was true in both the first BDM market they participated in and across all markets they participated in. Multi-Media Learning instructions also reduced how often subjects made a focal mistake following a mistake in the previous

[4]260 papers published in either top five general interest journals or *Experimental Economics* between January 2011 and December 2016

[5]Here a high ability subject is one that scored either a 2 or a 3 on the CRT. More details on these subjects can be found in section 3.3.4.

Table 3.1: Treatments

		Instruction Type	
		Multi-Media Learning	Standard
Quiz	Incentivized	**MMI**	**SI**
Structure	Unincentivized	**MMU**	**SU**

market. Incentivizing the pre-experiment quiz had a negligible effect on subjects understanding of the BDM mechanism.

3.2 Experimental Design

The design of this experiment is broken down into two parts. The pre and post experiment tasks subjects participated in, which looked to improve subject understanding of the BDM mechanism, or control for individual characteristics that could impact understanding. These tasks are described in 3.2.1. The other part is the description of the BDM mechanism subjects participated in. This mechanism is described in 3.2.2. Finally the Hypotheses for this study, ahead of the experiment, are found in 3.2.3.

3.2.1 Pre and Post Experiment Tasks

This experiment follows a 2x2 design, varying the different types of pre-experiment tasks subjects participated in. In all treatments, subjects first received instructions in the form of a short video, lasting between six and seven minutes depending on the treatment, and then they took a short comprehension quiz before beginning the primary experiment. There were two different types of instructions subjects could receive, as well as two different types of quizzes. The design can be visualized on Table 3.1.

The first dimension was the type of quiz subjects took following the instructions. In all treatments, subjects answered the same six quiz questions. The order of these questions was randomly determined for each session, but was the same order for all subjects in that session. After every subject in a session finished the quiz, subjects were given a chance to review their answers to the quiz. Questions were reviewed in the same order they were taken. When reviewing a question, subjects were told whether or not they answered the question correctly. If they answered it incorrectly, they were told what the correct answer was and why. After each subject in a session finished reviewing the quiz questions, they were told the total number of questions they answered correctly, and then they moved onto the primary experiment.

The six quiz questions were broken down into three types of questions. The first type was questions about the parameters of the game. One of these questions asked subjects what their possible valuations were. The other asked them what the possible range of posted prices was. The second type of question asked them about the rules of the BDM mechanism. Subjects were given a possible offer price and then were asked the range of posted prices that would either pay them their valuation, or pay them the posted price. The third type of question asked them to calculate their payoff. Subjects were given a valuation, an offer price, and a posted price and were asked to report their earnings in that market.

The difference between treatments was whether or not subjects were paid for their quiz answers. In the Incentivized treatments, subjects were paid $0.50 for each quiz question they answered correctly. This was in addition to the earnings they received from the show-up fee and the rest of the experiment. Subjects were told about this incentive during the instructions. In the Unincentivized treatments, subjects were not paid for their quiz answers. No mention of potential payment was included in the instructions these subjects received.

Because subjects were given feedback about their quiz answers before moving onto the main experiment, this design potentially creates wealth effects. Subjects that took an Incentivized quiz would know that they already made some money in addition to their show-up fee, while those that took an Unincentivized quiz would only have the show-up fee as earnings. To control for these potential wealth effects, subjects in an Unincentivized treatment were given a surprise payment following the quiz, but before they began the main experiment. This payment was equal to the average quiz payoff in the previous Incentivized treatments.[6]

The reason to include a pre-experiment quiz is to reinforce the information subjects learned in the instructions. Answering questions about the experiment forces subjects to think more about the decision environment before experiencing it. This hopefully improves subject performance in the experiment. Incentivizing the quiz is then a further step in getting subjects to pay attention to the instructions. If they know they can earn additional money from it, they may follow the instructions more closely and have an even greater understanding of the decision environment.

To this point, there has been little work in looking into the benefits of incentivizing quizzes. In Freeman et al.'s (2018) literature survey, they found that only three of the 260 experiments in their data set used an Incentivized quiz. In their own study, they found that incentivizing quizzes did improve subject understanding of the decision environment. This is because incentivizing the quiz made incorrect answers more salient during review of the quiz questions. This in turn improved future decision making. The inclusion of an Incentivized quiz here is a further check on the benefits of incentivizing pre-experiment quizzes.

The second dimension that treatments varied by was the type of instructions subjects received. As was previously mentioned, all instructions for this experiment

[6]This amount came out to be $1.50 for all sessions

were video instructions.[7] Video instructions were used to ensure some uniformity in how subjects received the instructions. If instructions were read out loud by the experimenter, then there could differences in how subjects perceived the instructions across sessions. This is because the experimenter may read the instructions differently in later sessions, either consciously or unconsciously, as they read them more often. To avoid this, video instructions read by the computer were used to ensure consistency across sessions.

The content of the instructions was the same across all treatments, but the way that content was presented to subjects differed by treatment type. In the Standard treatments, subjects received instructions that were based on those in Cason and Plott (2014), adapted to the differences between their design and the one used here. The Standard treatments were the baseline for which the other treatment was compared to. For this reason a set of instructions previously used in the experimental literature was ideal. Additionally, the instructions used in Cason and Plott (2014) were very simple, allowing for easy adaption to the video instruction format.

The other type of instructions used in this experiment were written using influences from the Multi-Media Learning literature in psychology. The idea behind Multi-Media Learning is to take advantage of multiple learning channels, here audio and visual, while also using different techniques to reduce the cognitive load subjects are placed under. This allows for subjects to process more information at once, while also assisting them by only having them think about what is most important. See Mayer and Fiorella (2014) for an overview of the findings in the Multi-Media Learning literature and how to apply them.[8]

[7] Examples of these video instructions can be found on the author's personal website, cjwcandreva.com. The Multi-Media Learning Incentivized and the Standard Unicentivized instructions can be found there.

[8] Schweitzer (2017) is a working paper that provides a good guide for applying these principles to experimental economics instructions

Mayer and Fiorella (2014) include five load-reducing methods to help improve learner understanding. Three of these load-reducing methods were used in writing the Multi-Media Learning instructions for this experiment. These methods were the principles of Signaling and Redundancy, to address extraneous materials, and Temporal Contiguity, to help reduce representational holding. The two that were not used were Coherence, which addresses extraneous materials, and Spatial Contiguity, which address confusing layouts. Coherence is the elimination of related, but irrelevant information. This is something that is typically done in experimental economics instructions already. Spacial Contiguity has to do with placing words near the graphics they describe. Since many words were removed from the Multi-Media Learning instructions, through the use of the Redundancy technique, this technique was not used here.

The idea behind Signaling is to provide visual cues to help subjects know what part of the screen they should be focusing on. This was done by using the computer mouse to point to the relevant part of the interface, as well as the timing of when different elements were added to the screen. New elements were added to the screen when the narration began describing them.

The Redundancy principle has to do avoiding presenting identical streams of information. For the Standard instructions, captions for the narration was provided on the bottom of the screen. They were not included in the Multi-Media Learning instructions. Additionally, pictures were used in place of words wherever possible to minimize the amount of identical information being provided.

Temporal Contiguity works to try and minimize representational holding. Representational holding is how many different pieces of information subjects have to hold in their mind at once. This is done by lining up the visuals subjects see with what they are hearing, and not requiring them to try and recall what they were shown earlier. Here this was done by explaining the BDM mechanism to the subjects by using the interface for the Multi-Media Learning instructions. For the Standard instructions, the BDM mechanism and interface were explained separately. This reduced the cognitive load on the subjects that received Multi-Media Learning instructions, since they did not need to think about how the BDM mechanism worked while learning the interface.

After subjects completed all the pre-experiment tasks, they moved onto the main experiment. This experiment was a series of ten BDM markets. The details of these markets can be found in section 3.2.2. After subjects completed the main experiment, they then participated in one or two additional tasks. The first of these tasks, which all subjects participated in, was a quick demographic survey.[9] Subjects were asked seven questions about themselves and their prior experiences that may have impacted their decision in this experiment.[10] The answers were used as controls for the regressions found in section 3.3.

The second of these task was the Cognitive Reflection Test (CRT). The Economics Science Laboratory (ESL) at the University of Arizona was in the process of adding CRT scores to their subject database at the time of this experiment. Some subjects that participated in this experiment had already participated in the CRT ahead of this experiment. Those subjects CRT scores were not recollected, and their scores from the database were used. For those subjects that had not previously done the CRT at the ESL, they were given the option to participate in the CRT af-

[9]For one session of the Standard Incentivized treatment, a coding error prevented the collection of this data

[10]The specific questions subjects were asked were (Possible Answers in parentheses): Gender (Male, Female, Non-binary, Other), Year in School (1st, 2nd, 3rd, 4th, 5th or more), GPA (3.5-4.0, 3.0-3.5, 2.5-3.0, 2.0-2.5, <2.0), Major (Business, Economics, Engineering or Math, Other), If they have taken Intermediate Microeconomics (Yes, No), If they have taken Experimental Economics (Yes, No), and Whether or not they have done a BDM experiment before (Yes, No).

ter the main experiment. Subjects that participated in the CRT[11] had three minutes to answer all three questions. Subjects were paid $1 for each of these questions they answered correctly. These earnings were in addition to their earnings from the primary experiment.

The CRT scores were collected as a measure of innate subject ability. If higher quality subjects happened to be in one particular treatment, or another, that could change the potential results of this experiment. Controlling for subjects' innate ability allows for a better understanding of the impact of the pre-experiment tasks used here on subjects' understanding of the BDM mechanism. In addition to using CRT, subjects Extended Cognitive Reflection Test 1 (ECRT1) and ECRT2 scores were also used to control for subjects' innate ability.[12] Put forth by Noussair, Tucker, and Xu (2016), the ECRT1 and ECRT2 scores take an additional step in controlling for subject ability. In ECRT1, in addition to getting +1 to their score for a correct answer, -1 is applied to a subjects score if they give the intuitive answer, i.e. successfully use spontaneous and intuitive reason, but not reflective reasoning. All other incorrect answers receive a score of 0. This help to highlight subjects that incorrectly use that type of reasoning. ECRT2 instead scores these intuitive answers as 0 and all other incorrect answers as -1. This helps to highlight subjects that fail to engage in either type of reasoning.

3.2.2 BDM Mechanism

After subjects finished all the pre-experiment tasks, they moved onto the main experiment. This experiment was 10 rounds of a BDM market with induced valua-

[11] Three subjects declined the opportunity to answer the CRT questions. These subjects were dropped from the data set in any regression that involved their CRT scores.

[12] Eight subjects' CRT forms were misfiled after their CRT scores were recorded and put in the ESL subject database. Thus ECRT1 and ECRT2 scores are not available for these subjects.

tions. Throughout the experiment, and the rest of the paper, each round was referred to as a market. The BDM was chosen for multiple reasons. It is a commonly used tool in experimental economics. It is a mechanism that is well understood by researchers, yet subjects tend to have some difficulty understanding. If subjects have induced valuations, it has a unique, optimal action known to the experimenter. These factors made it the perfect environment to test if either instructions inspired by Multi-Media Learning, or Incentivized pre-experiment quizzes help to improve subject understanding.

At the start of the experiment, each subject was given a piece of colored card-stock. In the instructions, subjects were told that this was their card-stock, and that they would have an opportunity to sell it over the course of the experiment. They could sell it in these 10 different BDM markets, one of which would count for payment. At the start of each market, subjects were given a valuation for their card-stock. This valuation could be $2, $3, $4, or $5. If subjects did not sell their card-stock in the BDM market, this is the amount the experimenter gave them for their card-stock at the end of the experiment, if that market was chosen for payment.

After subjects saw their valuation for their card-stock, they made a choice of what offer price they would like to sell their card-stock at. Once subjects submitted their offer price, a posted price was randomly drawn from a uniform distribution between $0 and $12. Subjects were aware of this when they chose their offer price. If the posted price was greater than or equal to half the subjects chosen offer price, then the subject sold their card-stock in that market and received the posted price as payment. If instead the posted price was less than half the subjects chosen offer price, the subject did not sell their card-stock in that market and received their valuation for their card-stock as payment for that market.

Using half the subjects chosen offer price to determine whether or not they sell their card-stock in the BDM market instead of using the offer price itself is a departure from the standard BDM design. The reason this choice was made was because of the structure of the mechanism itself. Because the possible valuations subjects could receive were \$2, \$3, \$4, or \$5, they likely valued the card-stock at this amount. This is because a piece of card-stock is only worth a few cents on its own. This would make the unique, payoff maximizing choice under a standard BDM mechanism to choose an offer price equal to the subjects valuation. Subjects choosing this offer price would show full understanding of the mechanism.

However, the valuation is the only number that subjects saw on the interface. Thus the subjects valuation could become focal, causing subjects with no understanding of the BDM mechanism to choose their valuation as their offer price. Thus it may be difficult to determine which subjects have a perfect understanding of the BDM mechanism, and those that have close to no understanding of the BDM mechanism. To avoid this issue, half the subject's chosen offer price was used as the threshold to determine whether or not the subject sold their card-stock in the BDM market, and not their actual offer price.

Additionally, multiple possible valuations were used as a further check on subject understanding. If subjects received the same valuation every market, they may figure out the payoff maximizing choice for that valuation, even if they do not truly understand the BDM mechanism. If subjects do understand how the BDM mechanism works, then changing their valuation in future markets will not impact their ability to choose the optimal offer price. They will still pick twice their valuation to ensure they receive at least their valuation as payment for that market. Subjects that just figured out the payoff maximizing choice for a particular valuation will not necessarily be able to do this. This design choice allows for further separation between subjects that fully understand this version of the BDM mechanism, and those with only a partial understanding. Thus, to help control for this, four unique valuations, \$2, \$3, \$4, and \$5 were used in this experiment.

3.2.3 Hypotheses

The aim of this project is to discover whether Multi-Media Learning instructions and/or Incentivized pre-experiment quizzes help to improve subjects understanding of decision environments commonly used in experimental economics. In order to determine if either of these are true, one needs a way of measuring the level of understanding that subjects have. This project will use three distinct measures to track subject understanding. These measures are:

1. The number of quiz questions subjects answered correctly.
   - Quiz scores were also used as a regressor for the other understanding measures. The questions were broken down by type when used in this way.
2. The percentage of the subjects that chose the unique, optimal offer price.
3. The percentage of subjects that made the focal mistake. That is, they chose an offer price equal to their valuation instead of one equal to twice their valuation.

Quiz scores provide a first check on whether or not subjects understand how the BDM works before they have any experience with it. Choosing the unique optimal offer price is an obvious measure of understanding. It is one that works even better when one considers that fact the understanding of the BDM can be thought of as a eureka moment. Subjects are likely to make the optimal choice or not, given a sufficient level of understanding, and they typically do not gradually get closer to the optimal offer price. Focal mistakes is a measure of misunderstanding. It helps to track if any of the pre-experiment tasks done in this experiment lead to lower subject understanding.

One other important note, before getting into the hypotheses, is that there is fundamental difference between the first market subjects participate in, and all other markets they participate in. That is because for the first market, the only information subjects have about the BDM mechanism comes from the pre-experiment instructions and quiz. After the first market, subjects also have their experiences with the mechanism to influence their choices. Thus, separating out the first market allows for a purer test of the impacts of these pre-experiment tasks.

However, even if subjects have a higher understanding of the BDM mechanism from either Multi-Media Learning instructions and/or an Incentivized quiz, they may not necessarily have a perfect understanding going into the first market. Thus, acquiring some experience maybe more beneficial for these subjects compared to subjects that received Standard instructions and/or an Unincentivized quiz. For this reason, each of the hypotheses are divided into two parts. One for the first market by itself, and one for all markets.

Hypothesis 3.1. *Subjects that received Multi-Media Learning instructions will show higher levels of understanding of the BDM mechanism:*

A. In the first market

B. Across all markets

Compared to subjects that received Standard instructions.

Hypothesis 3.2. *Subjects that received an Incentivized pre-experiment quiz will show higher levels of understanding of the BDM mechanism:*

A. In the first market

B. Across all markets

Compared to subjects that received an Unincentivized pre-experiment quiz.

Hypothesis 3.3. *Subjects that received Multi-Media Learning instructions AND an Incentivized pre-experiment quiz will show higher levels of understanding of the BDM mechanism:*

A. In the first market

B. Across all markets

Compared to subjects that received only one of Multi-Media Learning instructions or an Incentivized pre-experiment quiz.

3.3 Results

17 sessions of this experiment were run at the Economics Science Laboratory (ESL) at the University of Arizona between February and March 2019. Four sessions were run for the the Multi-Media Learning Unincentivized, Standard Incentivized, and Standard Unincentivized treatments. Due to a record keeping error, a fifth session of the Multi-Media Learning Incentivized treatment was needed in order to have a sample size comparable to the other three treatments. All sessions took approximately 30-45 minutes. Table 3.2 summarizes each treatment, including subject earnings.

As was mentioned in Section 3.2, three different measures were used to determine how well subjects understood the BDM environment:

1. The number of pre-experiment quiz questions subjects answered correctly.

2. The percentage of the time subjects chose the unique, optimal offer price.

Table 3.2: Treament Summaries.

	MMI	**MMU**	**SI**	**SU**
Sessions	5	4	4	4
Subjects	40	43	41	39
Avg. Earnings	$13.03	$13.74	$13.60	$13.35

Notes: Earnings include the $5 show-up fee. Subject payments were rounded to the nearest quarter.

3. The percentage of the time subjects chose the Focal Mistake. The Focal Mistake is choosing their valuation as their offer price instead of twice their valuation.

These measures are examined in the context of the first market subjects participated in, as well as over the course of the whole experiment. Additionally, throughout this paper, markets where subjects received a valuation of two are dropped from the data set. This was done to control for the fact that many subjects care less when they receive the lowest possible valuation,[13] and therefor are less likely to pay attention to their choices in those markets. Regressions and Summary Statistics with the full data set can be found in appendix A.3.

3.3.1 First Market

Table 3.3 has summary statistics for the understanding measures in the first market subjects participated in. The first market is presented separately because it is the only one in which subjects have no experience with the BDM mechanism. Thus any understanding of the BDM mechanism subjects have comes from the instructions and pre-experiment quiz.

Instead of by treatment, Table 3.3 breaks down the data by dimension. This al-

[13] See Cox, Smith, and Walker (1985)

Table 3.3: Summary Statistics, Understanding Measures, First Market.

	All	Standard		Multi-Media Learning	Incentivized		Unincentivized
Quiz Correct	2.63	2.59	$<$	2.66	2.78	$>^*$	2.48
Observations	163	80		83	81		82
Opt. Choice Rate	20.16%	14.75%	$<$	25.40%	10.71%	$<^*$	27.53%
Focal Mis.	3.23%	3.28%	$<$	3.17%	5.45%	$>$	1.45%
Observations	115	59		56	50		65

Notes: **All**: All Data. **Standard**: Data from the treatments with standard instructions. **Multi-Media**: Data from the treatments with Multi-Media Learning instructions. **Incentivized**: Data from the treatments with an Incentivized quiz. **Unincentivized**: Data from the treatments with an Unincentivized quiz. Significance determined using a Wilcoxon Rank Sum test with the Bonferroni Multiple Hypotheses Correction. *: Statistically significant at the 10% level.

lows for direct comparisons across Instruction Type and Quiz Structure. Starting on the Quiz Structure side, there are a couple of significant differences. Subjects that received and Incentivized quiz tended to score higher on the pre-experiment quiz, providing support for hypothesis 3.2. However, subjects that received an Unincentivized quiz were significantly more likely to choose the optimal offer price in the first market, providing support against hypothesis 3.2. Subjects in the Incentivized treatments were also more likely to make the focal mistake, though this difference is not significant. On the Instruction Type side, no differences are found between the number of quiz questions subjects answered correctly, how often they chose the optimal offer price, or how often they made the focal mistake.

To help determine if these differences, or lack there of, are based on the treatments themselves or subject characteristics, regressions were run with the understanding measures as the endogenous variables. The regressions in the body of this paper include ECRT2 scores in place of CRT scores and when number of correct quiz answers is not the endogenous variable, the results of the quiz broken down by question type. These questions variables have the value 0.5 when the subject got one of that question type correct, and 1 when they got both of that question type correct. This specification is used in the body of the paper because it maximized pseudo/adjusted R^2 across the various regressions using a variety of ranking methods.[14] The results of these regressions using the data generated in the first markets subjects participated in can be found on Table 3.4 and Table 3.5. For the Focal Mistakes measure, no regressions were run do to sample size concerns.[15] Thus the results on Focal Mistakes comes from the raw measures.

Result 3.1. *Neither Instruction Type nor Quiz Structure significantly impacts the rate at which subjects make the Focal Mistake in the first market they participate in.*

Starting with Correct Quiz Answers on Table 3.4, one can see in Model 1, neither Multi-Media Learning instructions, nor Incentivizing the quiz has any impact on how many quiz questions subjects answered correctly. Once the ECRT2 scores are accounted for in Model 2 and 3, the interaction between ECRT2 and receiving Multi-Media Learning instructions is significant. This leads to the following result:

Result 3.2. *Subjects with higher ECRT2 scores do better on the quiz if they received Multi-Media Learning Instructions, providing support for hypothesis 3.1 A.*

For the Optimal Choice Rate, only two different models are on table 3.5. This is because including full controls lead to more regressors than optimal choices,[16] so

[14]The four ranking systems used were Choice Rank Voting across all regressions performed, Highest psuedo/adjusted R^2 Average within endogenous variables/data set, Highest pseudo/adjusted R^2 for the specification with the most regressors within endogenous variables/data set, Highest pseudo/adjusted R^2 for the specification with the least regressors within endogenous variables/data set. ECRT2 with this recording of quiz scores performed best for two of these rankings, and second best in a third.

[15]4 subjects made Focal Mistakes in the first market they participated in. This is reduced to 3 in the full controls case.

[16]In the data set of first market regressions, no value 2 observations, and full controls there were 24 optimal choices made, and 33 regressors. Thus no clean conclusions could be drawn for that specification.

Table 3.4: Quiz Correct Regressions

Model	**Model 1**	**Model 2**	**Model 3**
Intercept	2.54	2.51	2.47
	(0.058)	(0.075)	(0.742)
Multi-Media (MM)	-0.120	-0.243	-0.231
	(0.144)	(0.141)	(0.205)
Incentivized (I)	0.096	0.200	0.142
	(0.190)	(0.214)	(0.224)
MM * I	0.411	0.567	0.590
	(0.413)	(0.473)	(0.420)
ECRT2		0.038	-0.008
		(0.094)	(0.104)
MM * ECRT2		0.222*	0.221*
		(0.113)	(0.104)
I * ECRT2		-0.142	-0.153
		(0.114)	(0.097)
Full Controls	**No**	**No**	**Yes**
Observations	163	152	135

Notes: *: Significant at the 10% level. Three subjects declined to participate in the CRT task resulting in a decreased sample size using that control. Eights subjects CRT forms were collected by other researchers and were misplaced, so their ECRT1 and ECRT2 scores could not be determined. Additionally, due to a coding error, control questions were not answered by one session of the SI treatment, resulting in the smaller sample size with full controls. OLS regressions were used for all three models. Standard errors clustered at the session level.

only the regressions with and without ECRT2 scores are included. For both models the receiving a higher valuation and doing better on the Question Type 2 quiz questions lead to a higher rate of making the optimal choice. In the regressions without the ECRT2 scores, Multi-Media Learning instructions lead to a higher rate of choosing the optimal offer price, while receiving an Incentivized quiz led to a lower rate of choosing the optimal offer price. These marginal effects are not significant after including ECRT2 scores. The interaction between receiving Multi-Media Learning instructions and the Question Type scores are still significant however, implying that the Multi-Media Learning instructions serve as a viable substitute for doing

Table 3.5: Optimal Choice Regressions, First Market

Model	**Model 1**	**Model 2**
Multi-Media (MM)	0.304*	0.308
	(0.157)	(0.191)
Incentivized (I)	-0.454***	-0.254
	(0.172)	(0.198)
MM * I	0.341	0.287
	(0.239)	(0.271)
Valuation	0.143***	0.152***
	(0.040)	(0.049)
Ques. Type 1 (Q1)	0.219**	0.177
	(0.097)	(0.119)
Ques. Type 2 (Q2)	0.249***	0.310***
	(0.083)	(0.084)
MM * Q1	-0.351***	-0.339**
	(0.114)	(0.153)
MM * Q2	-0.302**	-0.381***
	(0.120)	(0.124)
ECRT2		0.070***
		(0.019)
Observations	124	115

Notes: ***: Significant at the 1% level. **: Significant at the 5% level. *: Significant at the 10% level. Three subjects declined to participate in the CRT task resulting in a decreased sample size using that control. Eights subjects CRT forms were collected by other researchers and were misplaced, so their ECRT1 and ECRT2 scores could not be determined. Probit regressions were used for both models. Marginal Effects are reported on the table. Standard errors clustered at the session level.

well on the quiz. Thus, we have the following result:

Result 3.3. *Multi-Media Learning instructions have a similar impact on optimal choice rate as higher quiz scores do, providing support for hypothesis 3.1 A.*

3.3.2 All Markets

In addition to looking at subjects choices in the first market, their behavior across all markets was also examined. In markets after the first, it is possible that not

Table 3.6: Summary Statistics, Understanding Measures, All Markets.

	All	**Standard**		**Multi-Media Learning**	**Incentivized**		**Unincentivized**
Opt. Choice Rate	16.90%	15.30%	<	18.46%	15.22%	<	18.56%
Focal Mis.	5.17%	6.09%	>	4.25%	5.89%	>	4.43%
Observations	1219	608		611	610		609

Notes: **All**: All Data. **Standard**: Data from the treatments with standard instructions. **Multi-Media**: Data from the treatments with Multi-Media Learning instructions. **Incentivized**: Data from the treatments with an Incentivized quiz. **Unincentivized**: Data from the treatments with an Unincentivized quiz. Significance determined using a Wilcoxon Rank Sum test with the Bonferroni Multiple Hypotheses Correction.

only pre-experiment instructions and quizzes influenced subjects choices, but also their experience in the previous markets. However as was stated in Hypotheses 3.1 *B*., 3.2 *B*, and 3.3 *B*, if Multi-Media Learning instructions and/or Incentivized pre-experiment quizzes actually improve subject understanding of the environment, then they should still cause subjects to perform better, even if they are learning from previous markets. To look at whether or not this is true, the same understanding measures from before were used, aside from quiz scores.[17] Summary statistics for these measures can be found on Table 3.6.

Again, starting with the Quiz Structure side, there is little difference in the understanding measures for subjects that received an Incentivized quiz, compared to those who received an Unincentivized quiz. Even the difference between the optimal choice rates, which was significant for the first market, is insignificant here. Thus, there is little evidence in the raw data that an Incentivized quiz improved subject understanding across the whole experiment. On the Instruction Type side, there is still no significant difference between optimal choice rates, or the rate of focal mistakes.

No significant differences were found between the first market measures and

[17] Since the quiz was only taken once, at the beginning of the experiment, subjects quiz scores cannot change as they get more experience.

the all market measures. This would imply that subjects tended to not learn from previous markets. An idea that is further supported by the drop in Optimal Choice Rate for subjects that received an Unincentivized quiz. As a further check that the different treatment effects where the drivers of these results, regressions were run with understanding measures as the endogenous variables. Additionally, the market number was included as an exogenous variable as a further check to see if any learning occurred across markets. The results of these regressions can be found on Table 3.7.

For the Optimal Choice Rate, neither Incentivized nor Multi-Media Learning is significant in Model 1, without controls. This more or less matches up with the raw statistics on Table 3.6. For Model 3, with full controls, receiving Multi-Media Learning instructions led to subjects choosing the optimal offer price at a significantly higher rate. Subjects were also more likely to choose the optimal offer price if they did better on all three quiz questions in Model 3. This is somewhat tempered by the negative interaction term between Multi-Media Learning and the quiz scores, but this is likely explained by Multi-Media Learning instructions being a reasonable substitute for the understanding gained by doing better on the quiz. Thus Model 3 gives the following result:

Result 3.4. *Subjects that received Multi-Media Learning instructions before the experiment are more likely to choose the optimal offer price across all markets. This provides support for Hypothesis 3.1 B.*

For the Focal Mistake regressions the marginal effect for Multi-Media Learning is positive and significant. However the magnitude of the marginal effect is quite small for Model 3, with these subjects making the focal mistake roughly have a percentage point more often. Furthermore, if the subjects that received Multi-Media

Table 3.7: All Market Regressions

	Optimal Choice Dummy			Focal Mistake		
Model	**Model 1**	**Model 2**	**Model 3**	**Model 1**	**Model 2**	**Model 3**
Multi-Media Learning (MM)	0.114	0.191	0.202*	0.043	0.074**	0.006*
	(0.124)	(0.119)	(0.119)	(0.033)	(0.036)	(0.003)
Incentivized (I)	-0.019	-0.090	-0.084	-0.054	-0.025	0.002
	(0.088)	(0.079)	(0.102)	(0.043)	(0.035)	(0.003)
Valuation	0.041***	0.043***	0.035**	0.013***	0.012**	0.0014*
	(0.016)	(0.016)	(0.017)	(0.005)	(0.005)	(0.0008)
Ques. Type 1 (Q1)	0.203***	0.202***	0.153*	0.025	0.017	0.001
	(0.065)	(0.065)	(0.089)	(0.024)	(0.021)	(0.005)
Ques. Type 2 (Q2)	0.070	0.115	0.135***	-0.134*	-0.099*	-0.059***
	(0.087)	(0.074)	(0.036)	(0.077)	(0.060)	(0.022)
Ques. Type 3 (Q3)	0.121	0.083	0.144**	0.020*	0.062***	0.007*
	(0.100)	(0.092)	(0.069)	(0.012)	(0.024)	(0.004)
ECRT2		0.044*	0.029		-0.006	-0.001
		(0.023)	(0.019)		(0.006)	(0.001)
MM * Q1	-0.179**	-0.211**	-0.201*	-0.043	-0.055**	-0.004
	(0.075)	(0.088)	(0.113)	(0.028)	(0.026)	(0.004)
MM * Q2	0.071	-0.027	-0.116*	-0.002	-0.015	0.054***
	(0.095)	(0.096)	(0.068)	(0.060)	(0.054)	(0.020)
MM * Q3	0.114	0.131	0.021	-0089***	-0.099***	-0.012**
	(0.134)	(0.111)	(0.119)	(0.034)	(0.036)	(0.005)
Full Controls	**No**	**No**	**Yes**	**No**	**No**	**Yes**
Observations	1219	1139	1014	1219	1139	1014

Notes: ***: Significant at the 1% level. **: Significant at the 5% level. *: Significant at the 10% level. Three subjects declined to participate in the CRT task resulting in a decreased sample size using that control. Eight subjects CRT forms were misfiled after the CRT scores were recorded, preventing the formulation of their ECRT2 scores. Additionally, due to a coding error, control questions were not answered by one session of the SI treatment, resulting in the smaller sample size with full controls. Probit regressions were used for all models. Marginal Effects are reported on the table. Standard errors were clustered by subject and session.

instructions do better on the Type 2 and 3 quiz questions, that works to lower the rate they make the focal mistake. This leads to the following result.

Result 3.5. *Among subjects that received Multi-Media Learning instructions, correctly answering the Type 2 and 3 quiz questions leads to a lower rate of making the Focal Mistake.*

3.3.3 Learning from Markets

In addition to looking at how subjects learned across all markets, how they behaved following specific situations was also studied. To do this a subset of the data featuring just the markets immediately following a subject making a mistake was created. Here a mistake is defined to be choosing an offer price other than the optimal offer price. There are two types of mistakes that subjects could have made, either a known mistake, or an unknown mistake. A Known mistake is one where one of two things happened:

1. Half the subject's chosen offer price was greater than the randomly drawn posted price, which in turn was greater than the subject's valuation in that market, i.e Half Offer Price $>$ Posted Price $>$ Valuation.

2. Half the subject's chosen offer price was less than the randomly drawn posted price, which in turn was less than the subject's valuation in that market, i.e. Half Offer Price $<$ Posted Price $<$ Valuation.

These types of mistakes are referred to as known mistakes because a subject that commits one earns less money because of it. Therefore they are directly effected by their mistake, which may in turn alter future choices. An unknown mistake, on the other hand, does not impact a subjects earnings. Their realized payoff is the

Table 3.8: Understanding Measures Following Known Mistakes

	All	Standard		Multi-Media Learning	Incentivized		Unincentivized
Opt. Choice Rate	10.26%	12.31%	>	8.79%	11.54%	>	8.97%
Focal Mis.	5.77%	7.69%	>	1.10%	3.85%	≈	3.85%
Observations	156	65		91	78		78

Notes: **All**: All Data. **Standard**: Data from the treatments with standard instructions. **Multi-Media**: Data from the treatments with Multi-Media Learning instructions. **Incentivized**: Data from the treatments with an Incentivized quiz. **Unincentivized**: Data from the treatments with an Unincentivized quiz. Significance determined using a Wilcoxon Rank Sum test with the Bonferroni Multiple Hypotheses Correction.

same as the one they would have received if they had not made a mistake in that market. Thus these subjects are less likely to learn from their mistake and alter their future choices. Understanding measure for subjects that made a known mistake in the previous market are on Table 3.8.

The understanding measures for subjects in markets directly following a known mistake are noticeably different from the general measures. The optimal choice rates are lower across the board, though none of the differences are significant.[18] The focal mistake rates are also smaller, suggesting that there is some learning happening after a known mistake, though again this difference is not significant.[19] Though this seems to suggest no impact of a known mistake, looking at what happens after a unknown mistake provides a full picture of the impact of making a mistake. The summary statistics for markets following unknown mistakes can be found on table 3.9.

Subjects in these markets have a significantly lower rate of make the optimal choice than subjects in the general data set do.[20] These subjects also make the focal

[18] P-values in table order after applying the Bonferroni Multiple Hypotheses Correction: 0.338, 1, 0.224, 1, 0.360.

[19] P-values in table order after applying the Bonferroni Multiple Hypotheses Correction: 1, 1, 1, 1, 1.

[20] P-values in table order after applying the Bonferroni Multiple Hypotheses Correction: <0.001, 0.006, <0.001, 0.023, <0.001.

Table 3.9: Understanding Measures Following Unknown Mistakes

	All	**Standard**		**Multi-Media Learning**	**Incentivized**		**Unincentivized**
Opt. Choice Rate	$8.25\%^G$	$7.93\%^G$	$<$	$8.64\%^G$	$8.62\%^G$	$>$	$7.90\%^G$
Focal Mis.	5.87%	6.39%	$>$	5.29%	6.27%	$>$	5.45%
Observations	750	391		359	383		367

Notes: **All**: All Data. **Standard**: Data from the treatments with standard instructions. **Multi-Media**: Data from the treatments with Multi-Media Learning instructions. **Incentivized**: Data from the treatments with an Incentivized quiz. **Unincentivized**: Data from the treatments with an Unincentivized quiz. Significance determined using a Wilcoxon Rank Sum test with the Bonferroni Multiple Hypotheses Correction. G: Statistically significantly different from the general results using all of the relevant data.

mistake more often than subjects in the general data set, though this difference is not significant.[21] While there are no significant differences between these subjects and those that committed known mistakes in the previous market,[22] this additional drop from the general market does imply some value to knowing a mistake was made in the previous market. To further verify this, and to determine if the different treatments had an influence in behavior amongst these subjects, regressions were run using the understanding measure as endogenous variables. The results of these regressions can be found on Table 3.10.

For the Optimal Choice regressions, none of the treatments have a significant impact on the choices subjects made. This is true even for Model 3, with full controls. This follows from what we say in the raw data, with no significant differences in the Optimal Choice Rate between the two groups. What is driving Optimal Choices among this group of subjects is doing well on Question Types 2 and 3. Subjects that did well on these questions continue to choose the optimal offer price more often, despite their previous mistakes. Additionally, the marginal effects for the interaction terms between the treatments and the Question Type 2 scores is neg-

[21] P-values in table order after applying the Bonferroni Multiple Hypotheses Correction: 1, 1, 1, 1, 1.

[22] P-values in table order after applying the Bonferroni Multiple Hypotheses Correction: Optimal Choice: 1, 1, 1, 1, 1. Focal Mistake: 1, 1, 0.835, 1, 1.

Table 3.10: Mistake Regressions

	Optimal Choice Dummy			Focal Mistake	
Model	**Model 1**	**Model 2**	**Model 3**	**Model 1**	**Model 2**
Multi-Media Learning (MM)	-0.017	-0.027	-0.006	0.097^{*}	0.127^{**}
	(0.067)	(0.067)	(0.050)	(0.051)	(0.063)
Incentivized (I)	0.067	0.025	0.008	-0.082^{*}	-0.067
	(0.053)	(0.053)	(0.052)	(0.048)	(0.048)
Valuation	0.039^{***}	0.040^{***}	0.029^{***}	0.011^{**}	0.008
	(0.009)	(0.009)	(0.010)	(0.006)	(0.006)
Market	-0.010^{***}	-0.009^{***}	-0.006^{***}	-0.001	-0.000
	(0.002)	(0.003)	(0.002)	(0.003)	(0.003)
Ques. Type 1 (Q1)	0.047	0.054	0.049	0.053^{*}	0.036
	(0.048)	(0.043)	(0.048)	(0.030)	(0.024)
Ques. Type 2 (Q2)	0.068	0.065	0.078^{**}	-0.114^{*}	-0.076
	(0.053)	(0.052)	(0.037)	(0.063)	(0.047)
Ques. Type 3 (Q3)	0.037	0.039	0.055^{***}	0.014	0.048
	(0.037)	(0.030)	(0.016)	(0.016)	(0.031)
MM * Q1	-0.006	0.012	0.021	-0.077^{**}	-0.089^{***}
	(0.053)	(0.065)	(0.064)	(0.032)	(0.034)
MM * Q2	-0.032	-0.036	-0.092^{*}	0.001	-0.022
	(0.057)	(0.064)	(0.049)	(0.071)	(0.065)
MM * Q3	0.055	0.059	0.022	-0.106^{**}	-0.100^{**}
	(0.057)	(0.049)	(0.052)	(0.049)	(0.047)
I * Q1	-0.049	-0.033	-0.033	0.065^{**}	0.076^{**}
	(0.048)	(0.096)	(0.051)	(0.027)	(0.034)
I * Q2	-0.148^{***}	-0.134^{**}	-0.061^{**}	0.143^{*}	0.116^{*}
	(0.051)	(0.058)	(0.029)	(0.077)	(0.060)
MM * Known Mistake	-0.045	-0.038	-0.030	-0.038^{**}	-0.033^{**}
	(0.029)	(0.030)	(0.030)	(0.017)	(0.015)
MM * ECRT2		-0.017	-0.012		0.011^{**}
		(0.013)	(0.012)		(0.005)
I * ECRT2		0.036^{***}	0.022		0.003
		(0.010)	(0.014)		(0.005)
Full Controls	**No**	**No**	**Yes**	**No**	**No**
Observations	906	842	743	906	842

Notes: *: Significant at the 10% level. **: Significant at the 5% level. ***: Significant at the 1% level. Three subjects declined to participate in the CRT task resulting in a decreased sample size using that control. Eight subjects CRT forms were misfiled after the CRT scores were recorded, preventing the formulation of their ECRT2 scores. Additionally, due to a coding error, control questions were not answered by one session of the SI treatment, resulting in the smaller sample size with full controls. Probit regressions were used for all models. Marginal Effects are reported on the table. Standard errors were clustered by subject and session.

ative, providing some evidence that the pre-experiment tasks could substitute for doing well on the quiz.

Result 3.6. *Neither Multi-Media Learning instructions, nor an Incentivized pre-experiment quiz significantly impact the rate at which subjects choose the optimal offer price following a mistake, either known or unknown, in the previous market.*

For the Focal Mistakes regressions, only two models are included. This is for the same reason only two models were included for the Optimal Choice Rate regressions in section 3.3.1. Looking at the regression results, as with the All Markets regressions, receiving Multi-Media Learning instructions increases the rate at which these subjects make the Focal Mistake. However, several interactions between the Multi-Media Learning instructions and other exogenous variables do have a significant impact on reducing this rate. These subjects are less likely to make the Focal Mistake following making a known mistake, and if they scored higher on Question Types 1 and 3.

Result 3.7. *Receiving Multi-Media Learning instructions impacts the rate subjects make Focal Mistakes in markets following a mistake through high scores on Question Types 1 and 3, and whether or not the previous mistake was known. The net effect of these is to lower the rate at which these subjects make the Focal Mistake.*

3.3.4 Grouped by CRT Score

One common feature of the regressions on tables 3.5, 3.7, and 3.10 was that the marginal effect of subjects ECRT2 scores was significant. This implies differences in behavior among subjects with different levels of ability. To further determine whether or not these pre-experiment tasks impact different subjects in different

Table 3.11: Summary Statistics, Understanding Measures, First Market, CRT 0 and 1.

	All	Standard		Multi-Media Learning	Incentivized		Unincentivized
Quiz Correct	2.59	2.60	$\approx$	2.59	2.82	$>^{**}$	2.30
Observations	106	47		59	60		46
Opt. Choice Rate	13.41%	8.11%	$<$	17.78%	7.32%	$<$	19.51%
Focal Mis.	4.88%	5.41%	$>$	4.44%	7.32%	$>$	2.44%
Observations	82	37		45	41		41

Notes: **All**: All Data. **Standard**: Data from the treatments with standard instructions. **Multi-Media**: Data from the treatments with Multi-Media Learning instructions. **Incentivized**: Data from the treatments with an Incentivized quiz. **Unincentivized**: Data from the treatments with an Unincentivized quiz. Significance determined using a Wilcoxon Rank Sum test with the Bonferroni Multiple Hypotheses Correction. *: Statistically significant at the 10% level.

ways, the data set was split into two parts, subjects that scored either a 0 or a 1 on the CRT and subjects that scored either a 2 or a 3 on the CRT. The behavior of these subjects was analyzed in the same manner the entire data set was above, including the removal of markets where subjects received a valuation of two. Therefor examining differences in behavior will begin with the first market.

First Market

Tables 3.11 and 3.12 show the summary statistics for the first market subjects with CRT scores of 0 or 1, or 2 or 3, respectively. Starting with Table 3.11, the summary statistics for these relatively low ability subjects looks similar to the general results. On the Instruction Type side, subjects that received Multi-Media Learning instructions again make the choose the optimal offer price more often and make the focal mistake less often than subjects that received Standard instructions, though these differences are no significant. The quiz scores are also closer together than they were before, almost equal between the two Instruction Types.

On the Quiz Structure side, as with the Instruction Type, many of the general

Table 3.12: Summary Statistics, Understanding Measures, First Market, CRT 2 and 3.

	All	**Standard**		**Multi-Media Learning**	**Incentivized**		**Unincentivized**
Quiz Correct	2.70	2.61	$<$	2.83	2.70	$\approx$	2.71
Observations	54	31		23	20		34
Opt. Choice Rate	35.90%^{C}	27.27%	$<$	47.06%^{C}	23.08%	$<$	42.31%
Focal Mis.	0%	0%	$=$	0%	0%	$=$	0%
Observations	39	22		17	13		26

Notes: **All**: All Data. **Standard**: Data from the treatments with standard instructions. **Multi-Media**: Data from the treatments with Multi-Media Learning instructions. **Incentivized**: Data from the treatments with an Incentivized quiz. **Unincentivized**: Data from the treatments with an Unincentivized quiz. Significance determined using a Wilcoxon Rank Sum test with the Bonferroni Multiple Hypotheses Correction. *: Statistically significant at the 10% level. C: Statistically significantly different from the results using data from CRT 0 and 1 subjects.

patterns repeat themselves. Subjects that received an Unincentivized quiz choose the optimal offer price more often, though this gap is no longer significant. More interestingly the gap in quiz scores between subjects that received an Incentivized quiz that those that received an Unincentivized quiz increased from the general result. This indicates that an Incentivized quiz maybe more useful for subjects of lower ability than for a general group of subjects.

For the high ability subjects that had CRT scores of 2 or 3, on Table 3.12, there are several noticeable differences between these subjects choices and those of the general population and the low ability subjects. None of these subjects made the focal mistake in the first market. All of the subjects did better on the quiz, with the exception of the subjects that received an Incentivized quiz. Additionally the gap between the quiz scores of the subjects that received and Incentivized quiz and those that received an Unincentivized quiz is gone, further suggesting that low ability subjects are the ones that most benefit from quiz incentivization.

The most interesting difference between these subjects and others however, is the rate at which they choose the optimal offer price. On average the rate at which these subjects choose the optimal offer price is 15 percentage points higher than the general data and over 20 percentage points higher than the low ability subjects. This gap is even larger for the high ability subjects that received Multi-Media Learning instructions, who choose the optimal offer price almost half the time.

Do to the small sample sizes of these data sets, regressions for the Optimal Choice Rate and Focal Mistakes could not be run. Thus the results for these understanding measures comes from the raw data.

Result 3.8. *High ability subjects that received Multi-Media Learning instructions are significantly more likely to choose the optimal offer price in the first market compared to low ability subjects that received Multi-Media Learning instructions.*

Result 3.9. *High ability subjects are significantly less likely to make the focal mistake than low ability subjects, regardless of which pre-experiment tasks they participate in.*

The regression results for the quiz regressions dividing the subject pool by CRT score are on Table 3.13. As before multiple different models are presented. The difference here is that only two different models, one with the survey question data and one without, were run instead of three. The reason for this is because the data sets were created using CRT data, then control measures derived using CRT data should not be included in the regressions. Thus the models that only added CRT or ECRT data were removed, and those scores were removed from the model with full controls as well.

As was the case for the models on Table 3.4 neither treatment variable, nor their interaction significantly impacted subjects quiz scores. The only difference between the models is that the intercepts for the high ability subjects is higher than those for the low ability subjects. This follows from the raw data, where high ability subjects

Table 3.13: Quiz Regressions, CRT Split

Model	CRT 0 and 1 Model 1	CRT 0 and 1 Model 2	CRT 2 and 3 Model 1	CRT 2 and 3 Model 2
Intercept	2.278	1.915	2.800	3.725
	(0.302)	(1.068)	(0.233)	(1.703)
Multi-Media Learning (MM)	0.044	0.009	-0.229	-0.376
	(0.331)	(0.336)	(0.270)	(0.347)
Incentivized (I)	0.515	0.208	-0.527	-0.753
	(0.463)	(0.428)	(0.439)	(0.568)
MM * I	0.002	0.372	1.178	1.526
	(0.572)	(0.518)	(0.749)	(0.932)
Observations	106	93	54	50

Notes: Three subjects declined to participate in the CRT task resulting in a decreased sample size using that control. Eights subjects CRT forms were collected by other researchers and were misplaced, so their ECRT1 and ECRT2 scores could not be determined. Additionally, due to a coding error, control questions were not answered by one session of the SI treatment, resulting in the smaller sample size with full controls. OLS regressions were used for all models. Standard errors clustered at the session level.

scored better than low ability subjects.

Result 3.10. *Neither treatment type significantly impacts subjects quiz scores. This is true for both high ability and low ability subjects.*

All Markets

Understanding measure for the all market outcomes can be found on Table 3.14 for the low ability subjects and Table 3.15 for the high ability subjects. Starting with the low ability subjects on Table 3.14, these subjects were less likely to both make a focal mistake and choose the optimal offer price. Additionally the gaps in optimal choice rate between Multi-Media Learning and Standard instructions and the one between Unincentivized and Incentivized quizzes are greater and significant compared to the general data set, though the comparative differences are still the

Table 3.14: Understanding Measures, All Markets, CRT 0 and 1

	All	Standard		Multi-Media Learning	Incentivized		Unincentivized
Opt. Choice Rate	12.63%	9.47%	$<^{*}$	15.24%	9.93%	$<^{**}$	16.05%
Focal Mis.	4.80%	6.13%	>	3.70%	5.87%	>	3.44%
Observations	792	359		433	443		349

Notes: **All**: All Data. **Standard**: Data from the treatments with standard instructions. **Multi-Media**: Data from the treatments with Multi-Media Learning instructions. **Incentivized**: Data from the treatments with an Incentivized quiz. **Unincentivized**: Data from the treatments with an Unincentivized quiz. Significance determined using a Wilcoxon Rank Sum test with the Bonferroni Multiple Hypotheses Correction.

Table 3.15: Understanding Measures, All Markets, CRT 2 and 3

	All	Standard		Multi-Media Learning	Incentivized		Unincentivized
Opt. Choice Rate	26.24%	25.11%	<	27.81%	30.43%	>	23.46%
Focal Mis.	4.70%	3.83%	<	5.92%	5.59%	>	4.12%
Observations	404	235		169	161		243

Notes: **All**: All Data. **Standard**: Data from the treatments with standard instructions. **Multi-Media**: Data from the treatments with Multi-Media Learning instructions. **Incentivized**: Data from the treatments with an Incentivized quiz. **Unincentivized**: Data from the treatments with an Unincentivized quiz. Significance determined using a Wilcoxon Rank Sum test with the Bonferroni Multiple Hypotheses Correction.

same.

For the high ability subjects on Table 3.15, the optimal choice rate is again higher than it was for the general data set or for the low ability subjects. Additionally, the relationship be the between the Incentivized and Unincentivized treatments has flipped, with the subjects that received an Incentivized pre-experiment quiz choosing the optimal offer price more often those that received an Unincentivized quiz. None of these differences are significant however. These subjects also make the focal mistake less often the general data set, though again, this difference is not significant.[23]

As was the case with the general data set, regressions were run to confirm the insights from the raw data are correct. The results of these regressions can be found

[23] As was stated before, 3 subjects declined to participate in the CRT, so they are dropped from the data set in the move from the general data to the high and low ability groups.

on Tables 3.16 and 3.17.

Table 3.16: All Market Regressions, CRT 0 and 1

Model	Optimal Choice Dummy Model 1	Model 2	Focal Mistake Model 1
Multi-Media Learning (MM)	0.006	-0.012	0.067**
	(0.094)	(0.107)	(0.034)
Incentivized (I)	-0.114	-0.067	-0.249**
	(0.93)	(0.049)	(0.114)
Valuation	0.054***	0.041***	0.004
	(0.010)	(0.012)	(0.005)
Ques. Type 1 (Q1)	-0.069	-0.051	0.022
	(0.071)	(0.076)	(0.016)
Ques. Type 2 (Q2)	0.120	0.166**	-0.021
	(0.078)	(0.071)	(0.059)
Ques. Type 3 (Q3)	0.007	0.079	0.086***
	(0.085)	(0.051)	(0.027)
MM * Q1	0.011	0.068	-0.050**
	(0.057)	(0.085)	(0.022)
MM * Q2	0.052	-0.102*	-0.048
	(0.091)	(0.059)	(0.052)
MM * Q3	0.075	0.038	-0.103***
	(0.114)	(0.100)	(0.033)
I * Q1	0.110	0.039	0.138***
	(0.075)	(0.076)	(0.041)
I * Q2	-0.206**	-0.110	0.084
	(0.089)	(0.080)	(0.068)
observations	792	701	792

Notes: ***: Significant at the 1% level. **: Significant at the 5% level. *: Significant at the 10% level. Due to a coding error, control questions were not answered by one session of the SI treatment, resulting in the smaller sample size with full controls, Probit regressions were used for all models. Marginal Effects are reported on the table. Standard errors were two-way clustered by subject and session.

As one can see on Table 3.16, neither treatment had a significant impact on the Optimal Choice Rate of low ability subjects. The only strong influences on these subjects ability to make the optimal choice was their valuation in the given market, and their ability to successfully answer the Question Type 2 quiz questions. On

Table 3.17: All Market Regressions, CRT 2 and 3

Model	Optimal Choice Dummy Model 1	Model 2
Multi-Media Learning (MM)	0.687***	0.981***
	(0.204)	(0.100)
Incentivized (I)	0.128	0.369
	(0.290)	(0.798)
MM * I	-0.185*	-0.228**
	(0.102)	(0.109)
Ques. Type 1 (Q1)	0.804***	0.518
	(0.187)	(0.470)
Ques. Type 3 (Q3)	0.331***	0.676***
	(0.104)	(0.197)
MM * Q1	-0.796***	-0.737
	(0.251)	(0.518)
MM * Q3	0.117	-0.763***
	(0.145)	(0.187)
Observations	404	370

Notes: ***: Significant at the 1% level. **: Significant at the 5% level. *: Significant at the 10% level. Due to a coding error, control questions were not answered by one session of the SI treatment, resulting in the smaller sample size with full controls, Probit regressions were used for all models. Marginal Effects are reported on the table. Standard errors were two-way clustered by subject and session.

the Focal Mistake side, only Model 1 was able to be run without any collinearity issues. For these subjects, receiving an Incentivized quiz significantly reduced the chances that they made the Focal Mistake. Additionally, while the marginal effect for Multi-Media Learning in positive, it is counterbalanced by the fact that marginal effects for the interactions between it and Question Types 1 and 3 and significant and negative. This follows Result 3.5 which also showed that subjects who received Multi-Media Learning instructions were less likely to make the Focal Mistake if they did well on the quiz.

Table 3.17 only includes Optimal Choice regressions. This is because these

Table 3.18: Understanding Measures, Known Mistakes, CRT 0 and 1

	All	**Standard**		**Multi-Media Learning**	**Incentivized**		**Unincentivized**
Opt. Choice Rate	9.17%	9.09%	<	9.23%	12.07%	>	3.92%
Focal Mis.	4.59%	9.09%	>	1.54%	5.17%	>	5.45%
Observations	109	44		65	58		51

Notes: **All**: All Data. **Standard**: Data from the treatments with standard instructions. **Multi-Media**: Data from the treatments with Multi-Media Learning instructions. **Incentivized**: Data from the treatments with an Incentivized quiz. **Unincentivized**: Data from the treatments with an Unincentivized quiz. Significance determined using a Wilcoxon Rank Sum test with the Bonferroni Multiple Hypotheses Correction.

subjects did not make the focal mistake frequently enough for the probit regressions to converge. The results of these regressions show that high ability subjects are significantly more likely to make the optimal choice if they received Multi-Media Learning instructions prior to the experiment. This benefit replaces the gains they would otherwise receive from doing well on the quiz.

Result 3.11. *High ability subjects are more likely to make the optimal choice if they received Multi-Media Learning instructions prior to the experiment. Neither treatment significantly impacts low ability subjects optimal choice rate.*

Result 3.12. *Low ability subjects are less likely to make the focal mistake if they received an Incentivized pre-experiment quiz, or if they received Multi-Media Learning instructions and did well on Quiz Question Types 1 and 3.*

Learning from Mistakes

This final section looks at the impact of learning between high and low ability subjects. The summary statistics for understanding measures following a known mistake can be found on Table 3.18 for low ability subjects and on Table 3.19 for high ability subjects. Low ability subjects look similar to the general data with optimal choice rates around 10%. The focal mistake rates are slightly lower than the gen-

Table 3.19: Understanding Measures, Known Mistakes, CRT 2 and 3

	All	Standard		Multi-Media Learning	Incentivized		Unincentivized
Opt. Choice Rate	13.33%	20.00%	>	8.00%	10.00%	<	16.00%
Focal Mis.	0%	0%	=	0%	0%	=	0%
Observations	45	20		25	20		25

Notes: **All**: All Data. **Standard**: Data from the treatments with standard instructions. **Multi-Media**: Data from the treatments with Multi-Media Learning instructions. **Incentivized**: Data from the treatments with an Incentivized quiz. **Unincentivized**: Data from the treatments with an Unincentivized quiz. Significance determined using a Wilcoxon Rank Sum test with the Bonferroni Multiple Hypotheses Correction.

Table 3.20: Understanding Measures, Unknown Mistakes, CRT 0 and 1

	All	Standard		Multi-Media Learning	Incentivized		Unincentivized
Opt. Choice Rate	8.61%	7.69%	<	9.47%	7.41%	<	10.28%
Focal Mis.	4.70%	6.07%	>	3.41%	5.72%	>	3.27%
Observations	511	247		264	297		214

Notes: **All**: All Data. **Standard**: Data from the treatments with standard instructions. **Multi-Media**: Data from the treatments with Multi-Media Learning instructions. **Incentivized**: Data from the treatments with an Incentivized quiz. **Unincentivized**: Data from the treatments with an Unincentivized quiz. Significance determined using a Wilcoxon Rank Sum test with the Bonferroni Multiple Hypotheses Correction.

eral data, though the gap between Standard and Multi-Media Learning instructions is even wider, though this difference is not significant.

For high ability subjects, the optimal choice rate is slightly higher than it was for the general data set, though not significantly different. More interesting is the fact that none of the high ability subjects made a focal mistake in a market following a known mistake. This implies that these subjects do learn something from the known mistake that the low ability subjects are not.

Tables 3.20 and 3.21 show the understanding measure for low and high ability subjects following an unknown mistake respectively. Low ability subjects look like the general results, with optimal choice rates between 7% and 10% and focal mistake rates around 5%.

High ability subjects, on the other hand, look quite different from the general

Table 3.21: Understanding Measures, Unknown Mistakes, CRT 2 and 3

	All	**Standard**		**Multi-Media Learning**	**Incentivized**		**Unincentivized**
Opt. Choice Rate	8.11%	8.96%	$<$	6.82%	13.58%	$>^*$	4.96%
Focal Mis.	7.21%	4.48%	$<$	11.36%^{C}	7.40%	$>$	7.09%
Observations	222	134		88	81		141

Notes: **All**: All Data. **Standard**: Data from the treatments with standard instructions. **Multi-Media**: Data from the treatments with Multi-Media Learning instructions. **Incentivized**: Data from the treatments with an Incentivized quiz. **Unincentivized**: Data from the treatments with an Unincentivized quiz. Significance determined using a Wilcoxon Rank Sum test with the Bonferroni Multiple Hypotheses Correction. C: Statistically significantly different from the results using data from CRT 0 and 1 subjects.

results. These subjects choose the optimal offer price less often than than low ability subjects. This difference is not significant. They also make the focal mistake more often than low ability subjects, with the gap between subjects who received Multi-Media Learning instructions being significant. This is further evidence that high ability subjects learn from known mistakes, but receiving no additional help from unknown mistakes, while low ability subjects fail to learn from either type of mistake.

To help confirm these findings, regressions were run including a dummy variable for whether or not the subjects committed a known mistake in the previous market. Due to sample size issue, the focal mistake regressions are not presented here. Thus the result for focal mistakes comes from the raw data.

Result 3.13. *High ability subjects do not commit the focal mistake following a market where they committed a known mistake. They do commit the focal mistake more often than low ability subjects or the general data following an unknown mistake.*

Tables 3.22 and 3.23 have the optimal choice regressions for subjects following a market where they made a mistake. Beginning with the low ability subjects on Table 3.22, one can see that neither treatment impacts decision making in this envi-

Table 3.22: Mistake Regressions, CRT 0 and 1

Model	Optimal Choice Dummy **Model 1**	 **Model 2**
Multi-Media Learning (MM)	-0.047	-0.064
	(0.062)	(0.071)
Incentivized (I)	-0.034	-0.044
	(0.065)	(0.049)
Valuation	0.040***	0.029***
	(0.010)	(0.009)
Market	-0.009***	-0.007***
	(0.003)	(0.002)
Ques. Type 2 (Q2)	0.089	0.103**
	(0.059)	(0.052)
MM * Q2	-0.040	-0.123*
	(0.078)	(0.068)
I * Q2	-0.141**	-0.058
	(0.057)	(0.039)
Observations	620	546

Notes: ***: Significant at the 1% level. **: Significant at the 5% level. *: Significant at the 10% level. Due to a coding error, control questions were not answered by one session of the SI treatment, resulting in the smaller sample size with full controls, Probit regressions were used for all models. Marginal Effects are reported on the table. Standard errors were two-way clustered by subject and session.

ronment. This continues from the Result 3.6. Only the Market Subjects are in and their valuation in that market have a significant impact on their ability to choose the optimal offer price.

For the high ability subjects on Table 3.23, Model 1 shows several of the features observed in the raw data. Having the mistake made in the previous market be known increased these subjects odds of choosing the optimal offer price. This is somewhat tempered by the interaction terms between Known Mistake and the treatments, suggesting that Multi-Media Learning instructions and/or an Incentivized pre-experiment quiz also help with learning in a similar manner. The magnitude

Table 3.23: Mistake Regressions, CRT 2 and 3

Model	Optimal Choice Dummy **Model 1**	 **Model 2**
Multi-Media Learning (MM)	-0.031	0.010
	(0.115)	(0.030)
Incentivized (I)	0.287	0.070
	(0.292)	(0.184)
Valuation	0.033**	0.001
	(0.014)	(0.001)
Known Mistake (KM)	0.244*	0.043
	(0.132)	(0.053)
Ques. Type 3 (Q3)	0.096***	0.010
	(0.017)	(0.010)
MM * KM	-0.50**	-0.002
	(0.023)	(0.002)
I * KM	-0.051**	-0.002
	(0.021)	(0.002)
MM * Q3	0.016	-0.008*
	(0.060)	(0.005)
I * Q2	-0.042*	0.016
	(0.026)	(0.040)
I * Q3	-0.121***	-0.011
	(0.039)	(0.009)
Observations	267	242

Notes: ***: Significant at the 1% level. **: Significant at the 5% level. *: Significant at the 10% level. Due to a coding error, control questions were not answered by one session of the SI treatment, resulting in the smaller sample size with full controls, Probit regressions were used for all models. Marginal Effects are reported on the table. Standard errors were two-way clustered by subject and session.

on the marginal effects for the model that includes survey data is much smaller though. To the point where virtually none of the marginal effects are significant. This is because the subject specific characteristics do a better job of explaining subject improvement than knowing they made a mistake in the previous mistake.

Result 3.14. *Neither treatment significantly impacts the rate at which subjects choose the optimal offer price in markets following a mistake. This is true for both high and low ability subjects.*

3.4 Conclusion

This paper examines the impact of different pre-experiment tasks on subject understanding of the experimental environment. Specifically it looks to see if instructions written using insights from the Multi-Media Learning literature, and/or an Incentivized pre-experiment quiz cause subjects to better understand the BDM mechanism. To test if these tasks do improve understanding, a 2x2 design was used with subjects receiving either Multi-Media Learning or Standard instructions, followed by either an Incentivized or an Unincentivized quiz to check their understanding of the instructions. After these pre-experiment tasks, subjects participated in 10 markets of a modified BDM market as a seller with induced valuations. The mechanism was modified by having half the offer price subjects chose be the threshold for determining whether or not subjects received the posted price, or their valuation as payment. This modified threshold was used to prevent to the optimal offer price from also being the most focal one. Induced valuations were used to ensure that each market had a unique, optimal offer price.

High ability subjects in this study that received Multi-Media Learning instructions were more likely to choose the optimal offer price across all markets they participated in, after controlling for a variety of subject specific factors. These subjects who received Multi-Media Learning instructions were also more likely to choose the optimal offer price in the first market they participated in. Multi-Media Learning instructions had different effects when it came to helping subjects learn the BDM mechanism after they had some experience with it. While they did not help subjects choose the optimal offer price more often after mistakes, Multi-Media Learning instructions did help subjects to make fewer focal mistakes following a market where they failed to take the optimal action. This is true for both high and low ability subjects.

On the quiz side, Incentivized pre-experiment quizzes were not found to have a significant impact on the rate at which subjects chose the optimal offer price. This is true in both the first market, and across all markets. This is in contrast to the Freeman et al. (2018) result that salience of mistakes from Incentivized quizzes improves understanding. Incentivized quizzes had little impact on improving understanding following known mistakes.

Going forward, I recommend researchers use techniques from the Multi-Media Learning literature when writing their experimental instructions. These techniques have been shown to improve the decision making of high ability subjects, while not adversely impacting the performance of low ability subjects. I also recommend that researchers not use an Incentivized quiz at the start of their experiment. While a pre-experiment quiz is useful, and does help to improve subject understanding, incentivizing that quiz provides no further benefit. This is true even when it comes to improving subjects' quiz scores, after controlling for various subject specific characteristics. Thus, all the Incentivized quiz does is increase the amount researchers are paying to subjects without improving the quality of the data.

As a line of further research, it would be interesting to see the type of impact

Multi-Media Learning instructions could have in other experimental economics settings. One environment where they could have their greatest impact is in games well described by level-k models, such as the p-beauty contest. In these settings, if subjects have a greater understanding of the game, they maybe more likely to play higher level strategies, potentially even strategies not previously seen played in the literature. These types of instructions would also be useful for experiments involving complex settings, where a high level of understanding is needed to even approach equilibrium play.

Appendix A

A.1 Full SFEM Results

Table A.1 shows the full results from SFEM estimation. Standard errors generated from 1000 bootstraps samples are provided in parentheses next to each estimate. Table A.2 groups these estimates into a few macroarchetypes to get a better overall view about how subjects are approaching these supergames.

A.2 Risk Attitudes

There was little difference between the results of the risk elicitation tasks between the subjects that incorrectly answered one of more quiz questions and those that correctly answered all three questions. The only statistically significant difference between the to groups was that those that failed the quiz were more likely to make irrational choices in the Holt-Laury multiple price list task. An irrational choice is either choosing option A for all 10 rows, or switching multiple times, i.e. choosing A from rows 1-4, B in row 5, A in rows 6-7 and then B in rows 8-10. Using the results of all three tasks, CRRA parameters were estimated for all 50 subjects. Bounds on the task specific estimates are provided on the graphs below. Addition-

Table A.1: Full SFEM Results

	Disc	RT
All C	$0.043\ (0.030)$	$0.057\ (0.039)$
All D	$0.155\ (0.049)^{***}$	$0.060\ (0.034)^{*}$
TFT	$0.194\ (0.060)^{***}$	$0.123\ (0.071)^{*}$
TF2T	$0.108\ (0.050)^{**}$	$0.051\ (0.052)$
TF3T	$2.07E^{-15}\ (9.48E^{-3})$	$0.082\ (0.057)$
T2FT	$0.021\ (0.029)$	$0.101\ (0.070)$
T2F2T	$1.27E^{-14}\ (1.78E^{-10})$	$0.098\ (0.068)$
STFT	$0.249\ (0.061)^{***}$	$0.082\ (0.048)^{*}$
STF2T	$1.53E^{-14}\ (1.24E^{-12})$	$1.22E^{-12}\ (6.74E^{-8})$
STF3T	$1.38E^{-14}\ (1.83E^{-13})$	$0.020\ (6.25E^{-4})$
Grim	$0.036\ (0.024)$	$0.103\ (0.059)^{*}$
Grim2	$5.73E^{-15}\ (3.38E^{-3})$	$0.198\ (0.087)^{**}$
Grim3	$0.089\ (0.044)^{**}$	$2.73E^{-13}\ (0.021)$
SGrim2	$0.036\ (9.68E^{-4})$	$1.06E^{-12}\ (2.01E^{-9})$
SGrim3	$0.020\ (0.020)$	$1.44E^{-12}\ (1.08E^{-7})$
WSLS	$0.031\ (0.023)$	$1.03E^{-12}\ (1.18E^{-7})$
WSLS2	$1.76E^{-14}\ (1.32E^{-12})$	$0.025\ (0.025)$
T2	$1.73E^{-14}\ (3.38E^{-12})$	$7.27E^{-13}\ (7.05E^{-3})$
False C	$0.017\ (0.018)$	$1.33E^{-13}\ (1.97E^{-3})$
ALT	$1.76E^{-14}\ (1.92E^{-12})$	$1.34E^{-12}\ (1.08E^{-7})$
β	$0.863\ (4.48E^{-4})$	$0.914\ (4.55E^{-4})$

Notes: ***: Significant at the 1% level. **: Significant at the 5% level. *: Significant at the 10% level. Disc is the estimation using data from the discounted treatment. RT is the estimation using the data from the random termination treatment. The first 20 periods (or every period is the supergame lasted less than 20 periods) of each of the last four supergames were used to generate these estimates. Standard errors generated using 1000 bootstrap samples.

Table A.2: Grouped Results

	Disc	RT
All C	0.043	0.057
All D	0.155	0.060
TFTs	0.323	0.455
STFTs	0.249	0.102
Grims	0.181	0.301
Others	0.048	0.025
β	0.863	0.914

Notes: Disc is the estimation using data from the discounted treatment. RT is the estimation using the data from the random termination treatment. All C and All D are their respective estimates. TFTs is the sum of all non-suspicious Tit-for-Tats (TFT, TF2T, TF3T, T2FT, and T2F2T). STFTs is the sum of all the suspicious Tit-for-Tats (STFT, STF2T, and STF3T). Grims is the sum of all the Grims (Grim, Grim2, Grim3, SGrim2 and SGrim3). Others is the sum of all the remaining strategies that do not fit into one of the above categories (WSLS, WSLS2, T2, False C, and ALT).

Table A.3: Summary Statistics, Understanding Measures, First Market.

	All	**Standard**		**Multi-Media Learning**	**Incentivized**		**Unincentivized**
Observations	163	80		83	81		82
Quiz Correct	2.63	2.59	$<$	2.66	2.78	$>$	2.48
Opt. Choice Rate	19.63%	18.75%	$<$	20.48%	11.11%	$<^{*}$	28.05%
Focal Mis.	3.68%	3.75%	$>$	3.61%	3.70%	$>$	3.66%

Notes: **All**: All Data. **Standard**: Data from the treatments with standard instructions. **Mulit-Media**: Data from the treatments with Multi-Media Learning instructions. **Incentivized**: Data from the treatments with an Incentivized quiz. **Uninventivized**: Data from the treatments with an Unincentivized quiz. Significance determined using a Wilcoxon Rank Sum test with the Bonferroni Multiple Hypotheses Correction. *: Statistically significant at the 10% level. ***: Statistically significant at the 1% level.

ally bounds on the whole range of estimates from all three tasks are provided for each subject.

A.3 Regressions with the Full Data Set

This section of the appendix includes all the summary statistics for understanding measures, and regressions run using the full data set. This includes subjects that had a valuation of 2 in the given market.

Figure A.1: Risk Attitudes Paid

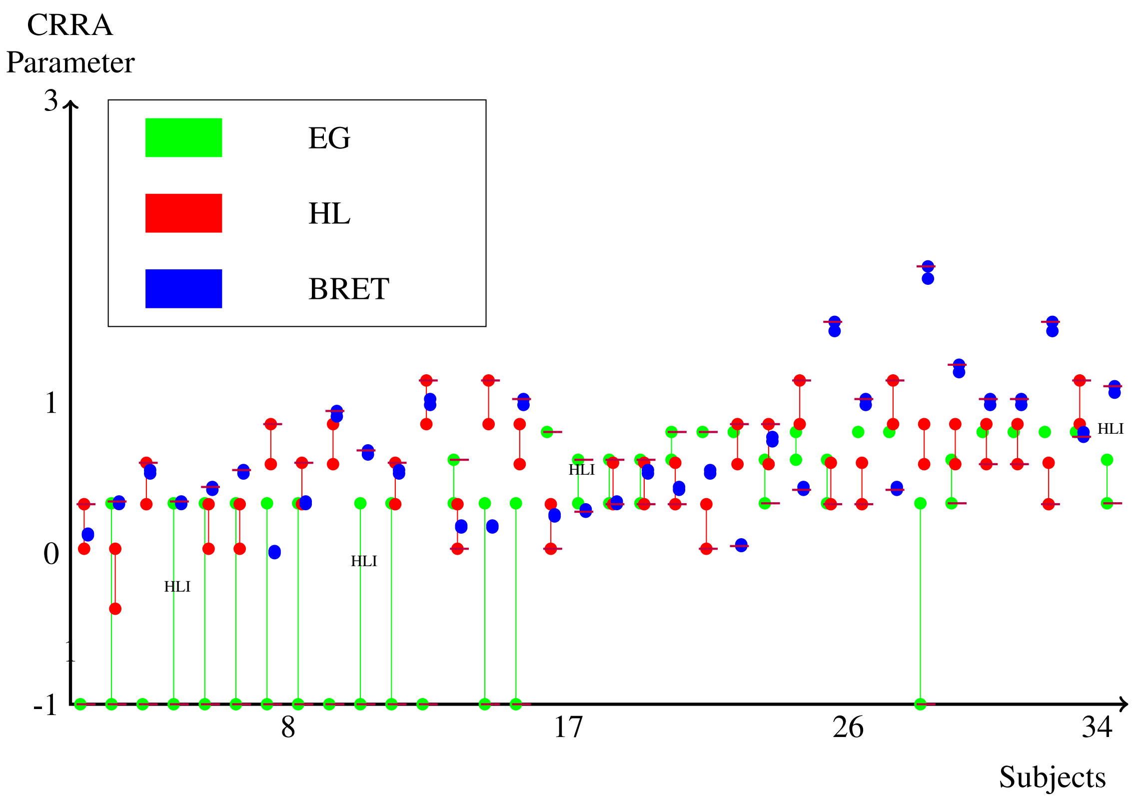

Notes: This graph shows the range of possible CRRA parameter values for the subjects that successfully passed the quiz on the risk elicitation tasks. The dots represent the upper and lower bounds of this estimate for the task of the matching color. The purple bars represent the upper and lower bounds from all three tasks. HLI means that that subjects response in the Holt-Laury multiple price list was irrational, i.e. multiple switching points or choosing option A for all 10 rows.

Figure A.2: Risk Attitudes Not Paid

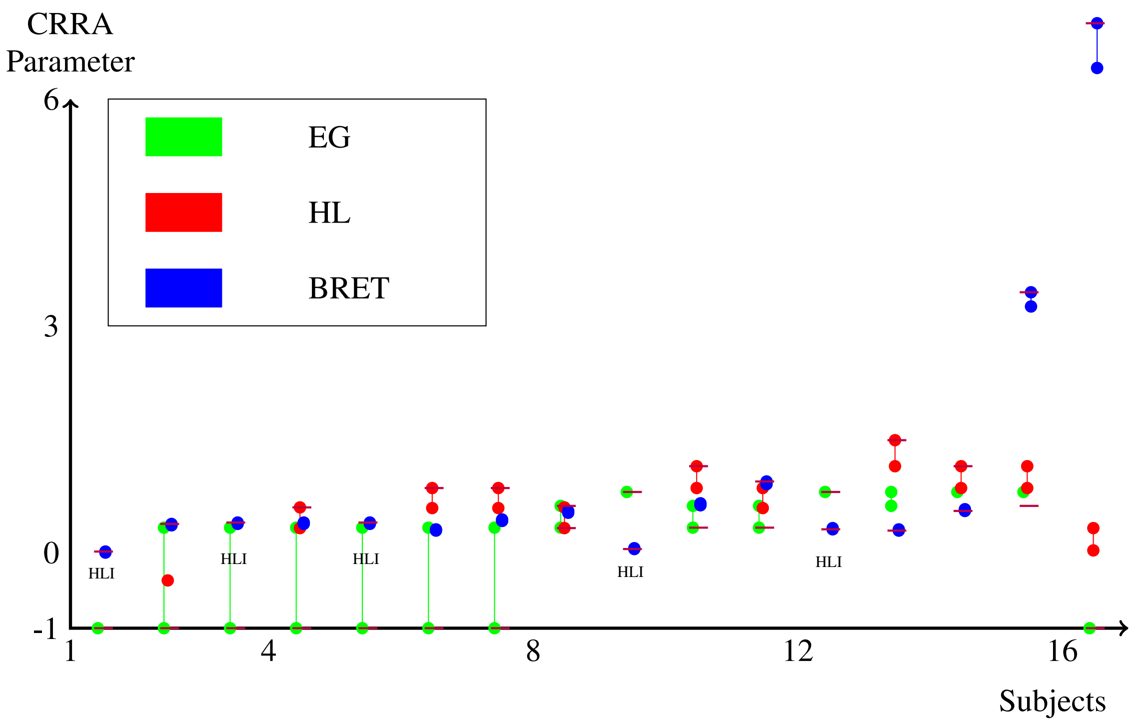

Notes: This graph shows the range of possible CRRA parameter values for the subjects that incorrectly answered one or more quiz questions on the risk elicitation tasks. The dots represent the upper and lower bounds of this estimate for the task of the matching color. The purple bars represent the upper and lower bounds from all three tasks. HLI means that that subjects response in the Holt-Laury multiple price list was irrational, i.e. multiple switching points or choosing option A for all 10 rows.

Table A.4: Summary Statistics, Understanding Measures, All Markets.

	All	Standard		Multi-Media Learning	Incentivized		Unincentivized
Observations	1630	800		830	810		820
Opt. Choice Rate	16.01%	14.50%	$<$	17.47%	14.44%	$<$	17.56%
Focal Mis.	6.01%	7.50%	$>$	4.58%	6.17%	$>$	5.85%

Notes: **All**: All Data. **Standard**: Data from the treatments with standard instructions. **Mulit-Media**: Data from the treatments with Multi-Media Learning instructions. **Incentivized**: Data from the treatments with an Incentivized quiz. **Uninventivized**: Data from the treatments with an Unincentivized quiz. Significance determined using a Wilcoxon Rank Sum test with the Bonferroni Multiple Hypotheses Correction. ***: Statistically significant at the 1% level.

Table A.5: Understanding Measures Following Known Mistakes

	All	Standard		Multi-Media Learning	Incentivized		Unincentivized
Observations	208	87		121	106		102
Opt. Choice Rate	10.10%	11.49%	$>$	9.09%	10.38%	$>$	9.80%
Focal Mis.	5.77%	11.49%	$>^{**}$	1.65%	6.60%	$>$	4.90%

Notes: **All**: All Data. **Standard**: Data from the treatments with standard instructions. **Mulit-Media**: Data from the treatments with Multi-Media Learning instructions. **Incentivized**: Data from the treatments with an Incentivized quiz. **Uninventivized**: Data from the treatments with an Unincentivized quiz. Significance determined using a Wilcoxon Rank Sum test with the Bonferroni Multiple Hypotheses Correction. G: Statistically significantly different from the general results using all of the relevant data. **: Statistically significant at the 5% level. ***: Statistically significant at the 1% level.

Table A.6: Understanding Measures Following Unknown Mistakes

	All	Standard		Multi-Media Learning	Incentivized		Unincentivized
Observations	1018	526		492	514		504
Opt. Choice Rate	8.25%^{G}	7.98%^{G}	$<$	8.53%^{G}	9.14%^{G}	$>$	7.34%^{G}
Focal Mis.	6.68%	7.60%	$>$	5.69%	6.42%	$<$	6.94%

Notes: **All**: All Data. **Standard**: Data from the treatments with standard instructions. **Mulit-Media**: Data from the treatments with Multi-Media Learning instructions. **Incentivized**: Data from the treatments with an Incentivized quiz. **Uninventivized**: Data from the treatments with an Unincentivized quiz. Significance determined using a Wilcoxon Rank Sum test with the Bonferroni Multiple Hypotheses Correction. G: Statistically significantly different from the general results using all of the relevant data. K: Statistically significantly different from Known Mistake. *: Statistically significant at the 10% level. ***: Statistically significant at the 1% level.

Table A.7: Quiz Correct Regressions

Model	**Model 1**	**Model 2**	**Model 3**
Intercept	2.54	2.51	2.47
	(0.058)	(0.075)	(0.742)
Multi-Media (MM)	-0.120	-0.243	-0.231
	(0.144)	(0.141)	(0.205)
Incentivized (I)	0.096	0.200	0.142
	(0.190)	(0.214)	(0.224)
MM * I	0.411	0.567	0.590
	(0.413)	(0.473)	(0.420)
ECRT2		0.038	-0.008
		(0.094)	(0.104)
MM * ECRT2		0.222*	0.221*
		(0.113)	(0.104)
I * ECRT2		-0.142	-0.153
		(0.114)	(0.097)
Full Controls	**No**	**No**	**Yes**
Observations	163	152	135

Notes: *: Significant at the 10% level. Three subjects declined to participate in the CRT task resulting in a decreased sample size using that control. Eights subjects CRT forms were collected by other researchers and were misplaced, so their ECRT1 and ECRT2 scores could not be determined. Additionally, due to a coding error, control questions were not answered by one session of the SI treatment, resulting in the smaller sample size with full controls. OLS regressions were used for all three models. Standard errors clustered at the session level.

Table A.8: First Market Regressions

	Optimal Choice Dummy		
Model	**Model 1**	**Model 2**	**Model 3**
Multi-Media Learning (MM)	0.227*	0.453**	0.497***
	(0.129)	(0.193)	(0.161)
Incentivized (I)	-0.253	-0.132	-0.172
	(0.161)	(0.201)	(0.198)
MM * I	0.052	0.010	-0.074
	(0.132)	(0.123)	(0.067)
Valuation	0.047*	0.050**	0.044*
	(0.028)	(0.024)	(0.027)
Ques. Type 1 (Q1)	0.209	0.248	0.157
	(0.137)	(0.176)	(0.153)
Ques. Type 2 (Q2)	0.138	0.257**	0.351**
	(0.117)	(0.115)	(0.176)
ECRT2		0.137***	0.134**
		(0.066)	(0.052)
MM * Q1	-0.375**	-0.437**	-0.325*
	(0.159)	(0.180)	(0.183)
MM * Q2	-0.042	-0.189	-0.346*
	(0.176)	(0.154)	(0.206)
I * Q3	0.100	0.253	0.339*
	(0.228)	(0.188)	(0.028)
MM * ECRT2		-0.086*	-0.116*
		(0.043)	(0.049)
Full Controls	**No**	**No**	**Yes**
Observations	163	152	135

Notes: *: Significant at the 10% level. **: Significant at the 5% level. ***: Significant at the 1% level. Three subjects declined to participate in the CRT task resulting in a decreased sample size using that control. Eight subjects CRT forms were misfiled after the CRT scores were recorded, preventing the formulation of their ECRT2 scores. Additionally, due to a coding error, control questions were not answered by one session of the SI treatment, resulting in the smaller sample size with full controls. Probit regressions were used for all models. Marginal effects reported on the table. Standard errors clustered at the session level.

Table A.9: All Market Regressions

	Optimal Choice Dummy			Focal Mistake		
Model	**Model 1**	**Model 2**	**Model 3**	**Model 1**	**Model 2**	**Model 3**
Multi-Media Learning (MM)	0.084	0.191*	0.198*	0.060	0.091**	0.049**
	(0.110)	(0.107)	(0.112)	(0.037)	(0.045)	(0.019)
Incentivized (I)	-0.010	-0.061	-0.072	-0.070	-0.042	-0.005
	(0.085)	(0.079)	(0.110)	(0.042)	(0.041)	(0.022)
MM * I	-0.042	-0.037	-0.078	0.002	-0.005	0.018
	(0.079)	(0.079)	(0.068)	(0.037)	(0.034)	(0.027)
Valuation	0.026***	0.026**	0.018	-0.003	-0.005	-0.004
	(0.010)	(0.011)	(0.012)	(0.005)	(0.005)	(0.004)
Market	-0.006**	-0.005**	-0.005**	0.001	0.001	0.001
	(0.002)	(0.002)	(0.003)	(0.002)	(0.002)	(0.001)
Ques. Type 1 (Q1)	0.189***	0.194*	0.163***	0.025	0.018	-0.003
	(0.072)	(0.072)	(0.097)	(0.050)	(0.038)	(0.032)
Ques. Type 2 (Q2)	0.066	0.120*	0.157***	-0.076*	-0.052*	-0.038
	(0.077)	(0.062)	(0.029)	(0.044)	(0.032)	(0.032)
Ques. Type 3 (Q3)	0.072	0.040	0.107*	0.041	0.079***	0.055***
	(0.082)	(0.081)	(0.064)	(0.029)	(0.027)	(0.017)
ECRT2		0.050**	0.038**		-0.001	-0.003
		(0.022)	(0.017)		(0.002)	(0.006)
MM * Q1	-0.146*	-0.197**	-0.187 -0.052	-0.066	-0.038	
	(0.077)	(0.087)	(0.125)	(0.048)	(0.046)	(0.027)
MM * Q2	0.056	-0.048	-0.143**	-0.036	-0.055	-0.003
	(0.092)	(0.094)	(0.066)	(0.060)	(0.056)	(0.035)
MM * Q3	0.138	0.154	0.055	-0.116**	-0.129***	-0.110***
	(0.112)	(0.104)	(0.122)	(0.047)	(0.044)	(0.032)
I * Q2	-0.087	-0.106	-0.048	0.138**	0.122**	0.041
	(0.096)	(0.083)	(0.050)	(0.011)	(0.053)	(0.031)
Full Controls	**No**	**No**	**Yes**	**No**	**No**	**Yes**
Observations	1630	1520	1350	1630	1520	1350

Notes: *: Significant at the 10% level. **: Significant at the 5% level. ***: Significant at the 1% level. Three subjects declined to participate in the CRT task resulting in a decreased sample size using that control. Eight subjects CRT forms were misfiled after the CRT scores were recorded, preventing the formulation of their ECRT2 scores. Additionally, due to a coding error, control questions were not answered by one session of the SI treatment, resulting in the smaller sample size with full controls. Probit regressions were used for all models. Marginal effects reported above. Standard errors were clustered by subject and session.

Table A.10: Mistakes Regressions

	Optimal Choice Dummy			Focal Mistake		
Model	**Model 1**	**Model 2**	**Model 3**	**Model 1**	**Model 2**	**Model 3**
Multi-Media Learning (MM)	-0.022	-0.007	0.010	0.131**	0.163**	0.082**
	(0.065)	(0.068)	(0.069)	(0.061)	(0.073)	(0.033)
Incentivized (I)	0.063	0.032	0.013	-0.109**	-0.082	-0.032
	(0.056)	(0.055)	(0.066)	(0.055)	(0.054)	(0.028)
Ques. Type 1 (Q1)	0.045	0.056	0.063	0.064	0.050	0.010
	(0.046)	(0.052)	(0.064)	(0.051)	(0.042)	(0.030)
Ques. Type 2 (Q2)	0.065	0.076	0.104**	-0.062	-0.037	-0.025
	(0.057)	(0.055)	(0.046)	(0.035)	(0.031)	(0.024)
Ques. Type 3 (Q3)	0.015	0.028	0.058**	0.032	0.063*	0.036**
	(0.041)	(0.037)	(0.026)	(0.032)	(0.034)	(0.016)
ECRT2		0.001	0.007		0.002	0.003
		(0.008)	(0.010)		(0.007)	(0.005)
MM * Q1	-0.003	0.004	0.027	-0.108*	-0.122**	-0.064**
	(0.052)	(0.070)	(0.084)	(0.055)	(0.057)	(0.031)
MM * Q2	- 0.017	-0.023	-0.091*	-0.034	-0.061	-0.001
	(0.061)	(0.065)	(0.055)	(0.081)	(0.075)	(0.035)
MM * Q3	0.063	0.055	-0.005	-0.136**	-0.129**	-0.110***
	(0.068)	(0.064)	(0.061)	(0.067)	(0.060)	(0.025)
I * Q1	-0.043	-0.037	-0.036	0.071*	0.073*	0.025
	(0.051)	(0.058)	(0.068)	(0.040)	(0.044)	(0.027)
I * Q2	-0.142***	-0.135**	-0.069**	0.144**	0.129**	0.045*
	(0.055)	(0.055)	(0.031)	(0.056)	(0.050)	(0.025)
MM * ECRT2		-0.030**	-0.030**		0.013	0.007
		(0.015)	(0.013)		(0.008)	(0.007)
I * ECRT2		0.040***	0.033**		0.001	-0.000
		(0.011)	(0.015)		(0.007)	(0.007)
MM * Known Mistake	-0.028	-0.027	-0.015	-0.048***	-0.045***	-0.019***
	(0.029)	(0.028)	(0.031)	(0.016)	(0.014)	(0.006)
Full Controls	**No**	**No**	**Yes**	**No**	**No**	**Yes**
Observations	1226	1136	998	1226	1136	998

Notes: *: Significant at the 10% level. **: Significant at the 5% level. ***: Significant at the 1% level. Three subjects declined to participate in the CRT task resulting in a decreased sample size using that control. Eight subjects CRT forms were misfiled after the CRT scores were recorded, preventing the formulation of their ECRT2 scores. Additionally, due to a coding error, control questions were not answered by one session of the SI treatment, resulting in the smaller sample size with full controls. Probit regressions were used for all models. Marginal effects reported above. Standard errors were clustered by subject and session.

Bibliography

[1] Abreu, Dilip and Rubinstein, Ariel. "The Structure of Nash Equilibrium in Repeated Games with Finite Automata". In: *Econometrica* (1988).

[2] Alekseev, Aleksandr, Charness, Gary, and Gneezy, Uri. "Experimental Methods, When and Why Contextual Instructions are Important". In: *Journal of Economic Behavior & Organization* (2017).

[3] Becker, Gordon, DeGroot, Morris, and Marschak, Jacob. "Measuring Utility by a Single-Response Sequential Method". In: *Behavioral Science* (1964).

[4] Bigoni, M. and Dragone, D. "Effective and Efficient Experimental Instructions". In: *Economics Letters* (2012).

[5] Blonski, Matthias, Ockenfels, Peter, and Spagnolo, Giancarlo. "Equilibrium Selection in the Repeated Prisoner's Dilemma: Axiomatic Approach and Experimental Evidence". In: *American Economic Journal: Microeconomics* (2011).

[6] Bruttel, Lisa and Kemecke, Ulrich. "Infinity in the Lab: How do People Play Repeated Games?" In: *Theory and Decision* (2012).

[7] Burfurd, Ingrid and Wilkening, Tom. "Experimental Guidance for Eliciting Beliefs with the Stochastic Becker-DeGroot-Marschak Mechanism". In: *Journal of the Economic Science Association* (2018).

[8] Cason, Timothy N. and Plott, Charles R. “Misconceptions and Game Form Recognition: Callenges to Theories of Revealed Preference and Framing”. In: *Journal of Political Economy* (2014).

[9] Chandler, Paul and Sweller, John. “Cognitive Load Theory and the Format of Instruction”. In: *Cognition and Instruction* (1991).

[10] Cox, James C., Smith, Vernon L., and Walker, James M. “Experimental Development of Sealed-Bid Auction Theory: Calibrating Controls for Risk Aversion”. In: *The American Economic Review* (1985).

[11] Crosetto, Paolo and Filippin, Antonio. “A Theoretical and Experimental Appraisal of Four Risk Elicitation Methods”. In: *Experimental Economics* (2016).

[12] Crosetto, Paolo and Filippin, Antonio. “The ’Bomb’ Risk Elicitation Task”. In: *Journal of Risk and Uncertainty* (2013).

[13] Dal Bó, Pedro. “Cooperation under the Shadow of the Future: Experimental Evidence from Infinitely Repeated Games”. In: *American Economic Review* (2005).

[14] Dal Bó, Pedro and Fréchette, Guillame R. “The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence”. In: *The American Economic Review* (2011).

[15] Dreber, Anna, Fudenberg, Drew, and Rand, David. “Whoo Cooperates in Repeated Games: The Role of Altruism, Inequity Aversion, and Demographics”. In: *Journal of Economic Behavior & Organization* (2014).

[16] Eckel, Catherine C. and Grossman, Philip J. “forecasting Risk Attitudes: An Experimental Study Using Actual and Forecast Gamble Choices”. In: *Journal of Economic Behavior & Organization* (2008).

[17] Farrell, Joseph and Maskin, Eric. “Renegotiation in Repeated Games”. In: *Games and Economic Behavior* (1989).

[18] Fréchette, Guillaume R. and Yuksel, Sevgi. “Infinitely Repeated Game in the Laboratory: Four Perspectives on Discounting and Random Termination”. In: *Experimental Economics* (2017).

[19] Frederick, Shane. “Cognitive Reflection and Decision Making”. In: *The Journal of Economic Perspectives* (2005).

[20] Freeman, David J. et al. “Instructions”. In: *Journal of the Economic Science Association* (2018).

[21] Fudenberg, Drew and Maskin, Eric. “The Folk Theorem in Repeated Games with Discounting or with Incomplete Information”. In: *Econometrica* (1986).

[22] Fudenberg, Drew, Rand, David, and Dreber, Anna. “Slow to Anger and Fast to Forgive: Cooperation in an Unvertain World”. In: *The American Economic Review* (2012).

[23] Healy, Paul J. “Explaining the BDM- or Any Random Binary Choice Elicitation Mechanism - To Subjects”. Working Paper. 2018.

[24] Holt, Charles A. and Laury, Susan K. “Risk Aversion and Incentive Effects”. In: *The American Economic Review* (2002).

[25] Ioannou, Christos A. and Romero, Julian. “A Generalized Approach to Belief Learning in Repeated Games”. In: *Games and Economic Behavior* (2014).

[26] Kim, Young-Gwan. “Evolutionarily Stable Strategies in the Repeated Prisoner’s Dilemma”. In: *Mathematical Social Sciences* (1994).

[27] Mayer, Richard E. and Fiorella, Logan. “Principles for Reducing Extraneous Processing in Multimedia Learning: Coherence, Signaling, Redundancy, Spatial Contiguity, and Temporal Contiguity Principles”. In: *The Cambridge Handbook of Multimedia Learning*. 2014. Chap. 12, pp. 279–315.

[28] Normann, Hans-Theo and Wallace, Brian. “The Impact of the Termination Rule on Cooperation in a Prisoner’s Dilemma Experiment”. In: *International Journal of Game Theory* (2012).

[29] Noussair, Charles N., Tucker, Steven, and Xu, Yilong. “Futures Markets, Cognitive Ability, and Mispricing in Experimental Asset Markets”. In: *Journal of Economic Behavior & Organization* (2016).

[30] Ramalingam, Abhijit, Morales, Antonio, and Walker, James. “Varying Experimental Instructions to Improve Comprehension: Punishment in Public Goods Games”. In: *Journal of Behavioral and Experimental Economics* (2018).

[31] Rapoport, Anatol and Chammah, Albert. 1965.

[32] Romero, Julian and Roskha, Yaroslva. “Constructing Strategies in the Indefinitely Repeated Prisoner’s Dilemma Game.” In: *European Economic Review* (2018).

[33] Roth, Alvin E. and Murnighan, J. Keith. “Equilibrium Behavior and Repeated Play of the Prisoner’s Dilemma”. In: *Journal of Methematical Psychology* (1978).

[34] Sabater-Grande, Gerardo and Georgantzis, Nikolaos. “Accounting for Risk Aversion in Repeated Prisoners’ Dilemma Games: An Experimental Test”. In: *Journal of Economic Behavior & Organization* (2002).

[35] Schwitzer, Sascha. “Cognitive-based Design Principles in Experimental Economics: How to Devise More Effective Instructions, Software, and Comprehension Tests”. Working Paper. 2017.

[36] Smith, Vernon L. “Microeconomic systems as an Experimental Science”. In: *The American Economic Review* (1982).

[37] Van Assen, M.A.L.M. and Snijders, C.C.P. “The Effect of Nonlinear Utility on Behavior in Repeated Prisoner’s Dilemmas”. In: *Rationality and Society* (2010).

[38] Zeng, Weijun et al. “Risk Consideration and Cooperation in the Iterated Prisoner’s Dilemma”. In: *Soft Computing* (2016).