NOISY FORESIGHT

Anujit Chakraborty
Chad W. Kendall

Working Paper 30333
http://www.nber.org/papers/w30333

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2022

Noisy Foresight
Anujit Chakraborty and Chad W. Kendall
NBER Working Paper No. 30333
August 2022
JEL No. D03,D90

ABSTRACT

Rational agents must perform backwards induction by thinking contingently about future states and actions, but failures of backwards induction and contingent reasoning are ubiquitous. How do boundedly-rational agents make decisions when they fail to correctly forecast actions in the future? We construct an individual decision-making experiment to collect a rich dataset in which subjects must reason only about their own future actions. We demonstrate substantial mistakes relative to the rational benchmark, and use the rich dataset to estimate several possible models of boundedly-rational foresight. We find that a model in which subjects expect to make more mistakes when the payoff consequences of their future actions are more similar best explains behavior.

Anujit Chakraborty
Department of Economics,
University of California, Davis
Davis, CA
United States
chakraborty@ucdavis.edu

Chad W. Kendall
Marshall School of Business
University of Southern California
701 Exposition Blvd, Ste. 231
Los Angeles, CA 90089
and NBER
chadkend@marshall.usc.edu

NOISY FORESIGHT

ANUJIT CHAKRABORTY* AND CHAD KENDALL**

ABSTRACT. Rational agents must perform backwards induction by thinking contingently about future states and actions, but failures of backwards induction and contingent reasoning are ubiquitous. How do boundedly-rational agents make decisions when they fail to correctly forecast actions in the future? We construct an individual decision-making experiment to collect a rich dataset in which subjects must reason only about their *own* future actions. We demonstrate substantial mistakes relative to the rational benchmark, and use the rich dataset to estimate several possible models of boundedly-rational foresight. We find that a model in which subjects expect to make more mistakes when the payoff consequences of their future actions are more similar best explains behavior.

1. Introduction

Economic theory for multi-stage games relies on backwards induction - systematic reasoning about the optimal actions at all possible future events. Economic theory for individual decision-making problems, such as job search or partner search problems, as well as sequential portfolio choice, similarly require decision-makers to contingently reason about how they will act in all future states of the world. Yet, we have evidence that players in multi-stage games pay little attention to the payoff consequences at future stages [Johnson et al., 2002] and fail to perform backwards induction reasoning [Güth and Tietz, 1990, McKelvey and Palfrey, 1992, Binmore et al., 2002]. Similarly, most people cannot reason contingently about others' actions taken at hypothetical events, be it co-participants in auctions [Kagel and Levin, 1986], or robots who play pre-defined strategies [Charness and Levin, 2009, Esponda and Vespa, 2014].

How should we think about the decisions of people that don't perform backwards induction? Do they completely ignore the payoffs beyond some horizon, or do they respond to changes in these payoffs in predictable ways albeit while making mistakes? To date, we have few theoretical models that consider this problem (we

* Department of Economics, University of California, Davis. (e-mail chakraborty@ucdavis.edu).
** Department of Finance and Business Economics, Marshall School of Business, University of Southern California; National Bureau of Economic Research (e-mail chadkend@marshall.usc.edu).
We thank Ryan Oprea for valuable discussions.

discuss the exceptions below), likely because most tests of forward-looking behavior have been well-designed to document reasoning failures, but not as well-designed to elicit comprehensive patterns of behavior on which new theories can be built.

In this paper, we provide rich empirical evidence about how boundedly-rational agents make decisions when they fail to perform backwards induction. We design an *individual* decision problem that requires subjects to understand only *their own decisions* at hypothetical future events and to respond accordingly. Being an individual decision problem, without any other real players, automated players, or complicated mechanism, we are able to rule out many possible reasons for departures from rational predictions, such as other-regarding preferences, incorrect beliefs about others' play, or misunderstanding the rules of a mechanism. Instead, observed mistakes can cleanly be linked to a failure to think through one's own actions in the future.

Further, subjects make decisions in many versions of this problem, providing us with comprehensive comparative statics with respect to all of the relevant parameters of the decision. We leverage this detailed empirical data to take some first steps in understanding what types of models can explain the patterns we observe: we build several variations of models proposed in the literature, and structurally estimate them to determine which best fits the data.

The two-period version of our decision problem (Figure 1.1) is identical to the Gneezy and Potters [1997] investment task, except for one additional feature. As in the standard investment task, in the first period, subjects choose to withdraw a fraction $1-x$ of one dollar at some known return R_1. The remaining fraction, x, remains invested in an uncertain project which will provide the opportunity to withdraw at a low or a high return (R_F^- or R_F^+, respectively). The critical additional feature is that *after subjects learn the realized return* (R_F^- or R_F^+) in the second period, we allow them to choose *not* to withdraw, instead earning an outside option with intermediate return rate $R_F^O \in [R_F^-, R_F^+)$. Subjects know about the outside option before their first-period decision, so that a rational decision-maker would perform backwards induction, reasoning that she will withdraw in the hypothetical event that R_F^+ realizes, but not withdraw when R_F^- realizes. Conditional on understanding her future actions, a subject then faces a standard Gneezy-Potters task: she decides how much to invest in a risky investment that returns either R_F^+ or R_F^O. We have subjects participate in twenty versions of this sequential decision problem while varying the returns (R_F^-, R_F^+, R_F^O) and the probability of R_F^+ versus R_F^-, thereby constructing a rich dataset from which to identify *individual* behavior.

To measure and account for the role of risk preferences, we also have subjects make investment decisions in five lottery tasks (see the right panel of Figure 1.1). These 'reduced' lottery tasks are specially chosen to correspond exactly to the

FIGURE 1.1. Two-Period Decision Problem

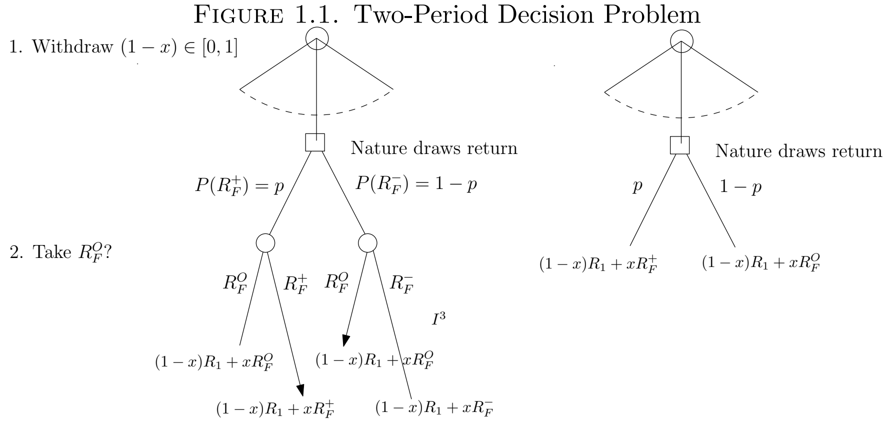

Notes: The two-period decision problem is on the left. The decision-maker decides once in each period, at the circular decision nodes. First, she chooses the amount to withdraw, $1-x$. Next, Nature, at the square node, chooses the return (R_F^+ or R_F^-). Then, in the second period, the decision-maker chooses between the realized return and the outside option return R_F^O. With $R_F^+ > R_F^O > R_F^-$, the decision-maker's correct second period choices are marked with an arrow. If the second-period choices are forecasted correctly, the two-period problem reduces to the simpler one-period lottery task on the right.

decision problems conditional on optimal decisions being made in the second period. Subjects' choices in the lottery tasks therefore serve as a benchmark for behavior in the decision problems - a rational subject would invest identically in a decision problem and its corresponding lottery task.

We document widespread failures of backwards induction in this simplest of settings: subjects invest systematically less in the decision problems than the corresponding lottery tasks. When we vary R_F^-, always keeping it less than the outside option, it should have no impact on decisions, but about 60% of subjects respond to the changes. Critically, these responses to an irrelevant parameter change do not reflect subjects actually making mistakes in the second period - mistakes in the second period are rare.[1] Perhaps more surprisingly, we find almost no evidence of learning - even after having made rational decisions in the second periods of many problems, subjects still fail to account for this behavior in the first period of the next problem.

Although subjects respond irrationally to the irrelevant second period return, other comparative statics match what we would intuitively expect, indicating that although subjects fail to perform backwards induction, they are otherwise responding to economic forces in predictable ways: they invest more when R_F^+ or the probability of it occurring increases, and less when the safe return, R_1, increases.

[1]About 10% of subjects frequently make mistakes in the second period, choosing dominated payoffs. We remove these subjects from the analysis because choosing less money over more money almost certainly indicates inattention and violates the most basic tenet of reasonable choice.

We conduct a second experimental treatment in which we extend the decision problems to three periods. In these problems, subjects can withdraw any amount in the first period and then can withdraw the remaining amount in either of two future periods, so that they now have an extra round of hypothetical thinking to perform. Importantly, we designed the three-period decision problems so that, from the perspective of a rational decision-maker, each is identical to a decision problem from the two-period treatment (and therefore also identical to one of the lottery tasks). Furthermore, we designed four different versions of the three-period problems that are all identical to the *same* two-period problem for a rational decision-maker, such that any observed variation provides additional evidence of what drives decision-making.

One simple hypothesis is that, because the three-period decision problems require an additional round of backwards induction, subjects perceive them as more complex, and hence invest less than in their two-period counterparts. Instead, we find more nuanced patterns of behavior. As in the two-period decision problems, subjects respond to the irrelevant final period return even though they rarely choose to withdraw at this return. However, they don't respond as strongly to this parameter as they do in the two-period problems. Subjects also invest more when an irrelevant (to a rational decision-maker) return in the second period increases, and invest differently when the probabilities of the highest possible return are irrelevantly swapped between the second and third periods.

The reduced-form evidence overall points to two regularities in the data: (i) subjects respond to irrelevant payoffs and changes in probabilities, but (ii) also respond to changes in other parameters as economic theory would predict. In addition to ruling out rational behavior, this evidence rules out several other models that have been put forth in the literature. It rules out completely ignoring payoffs beyond some horizon [Jéheil, 1995], as well treating the payoffs beyond some finite horizon as being random [Jehiel, 2001]. It also rules out any model in which a decision-maker assumes a fixed distribution over future payoffs (e.g. with a Laplacian prior or taking the average of the maximum and minimum payoffs as proposed in Rampal [2018]) because of the fact that decisions respond to the probabilities with which these payoffs are received.

We are led therefore to models in which subjects aggregate payoffs in the future in some reasonable, but noisy way. We build five potential models of behavior, leveraging ideas from the literature. Each model differs in how subjects aggregate future payoffs. The 'naïve' model assigns zero probability to the outside option (i.e. completely ignores it). The 'tremble' model assigns a constant probability of making a mistake (trembling) at every future event. The 'cursed' model, à la Eyster and Rabin [2005], assigns correct beliefs on average, but ignores the correlation between future actions and the realized contingencies. The 'noisy self' model is a version of quantal response equilibrium but relaxes the assumption that beliefs have to be consistent with actual future actions - correct actions are believed to be taken proportionally to the difference in payoffs across future actions. Finally, the 'generalized mean' model builds on the ideas in Ke [2019] - subjects aggregate future payoffs using a generalized mean function.

The dataset we collected allows us to structurally estimate all five models, in addition to the rational model, at the individual level. Doing so, we find that the noisy self model is strongly favored by model selection tests, both at the individual level and in the aggregate. The closely-related tremble model performs second best, suggesting that models in which subjects forecast their future actions with noise most accurately represent behavior.

Finally, we consider the possibility that subjects solve problems forward, placing more weight on immediate payoffs than future payoffs - they behave *as if* they discount payoffs received in future periods relative to those received in the current period. Of course, any form of actual time preference is unlikely because all payoffs are received at the same time, so we interpret this possibility as *cognitive discounting,* the discounting originating from the cognitive noise present in the mental simulation of future monetary amounts or utils [Gabaix and Laibson, 2017].[2] Re-estimating each of the five behavioral models and the rational model allowing for discounting, we find that models with cognitive discounting are favored over those without in 80 percent of cases. Estimation results with the noisy self model with discounting suggest that about 90 percent of subjects imperfectly forecast future actions, and about 70 percent cognitively discount future payoffs, so that a majority of subjects are best explained by a combination of the two behavioral factors.

Our paper contributes to the experimental literature on failures of backwards induction [Güth and Tietz, 1990, McKelvey and Palfrey, 1992, Johnson et al., 2002, Binmore et al., 2002, Mantovani et al., 2014, Rampal, 2018] and contingent reasoning [Esponda and Vespa, 2014, 2019, Martínez-Marquina et al., 2019]. A closely related companion paper by the co-authors shows people don't reason about their future actions in dynamic elicitation mechanisms [Kendall and Chakraborty, 2022].[3] One of our contributions is to demonstrate contingent reasoning failures in a very simple (two-period) setting when subjects only have to reason about their own actions (and not the rules of some mechanism or the actions of other players). We

[2]In an experimental study of intertemporal preferences, Enke and Graeber [2021] find that self-reported cognitive noise is highly correlated with behavioral patterns like high impatience over short horizons, hyperbolic discounting, and transitivity violations

[3]Fragiadakis et al. [2016] and Bosch-Rosa and Meissner [2020] show that people also don't understand their own actions when taken in the past or concurrently, respectively.

also contribute to the literature on decision-making by agents who use partially myopic calculations [Gabaix and Laibson, 2000, Gabaix et al., 2006].

A small theoretical literature models players with limited foresight in dynamic games. These models utilize different assumptions about how payoffs realized beyond a player's foresight horizon influence her decisions. For example, Jéheil [1995] and Mengel [2014] assume full discounting of all payoffs beyond the foresight horizon. Jehiel [2001] assumes payoffs outside of the foresight horizon are perceived as random. Rampal [2018] uses a curtailment rule where players with limited foresight perceive payoffs beyond their foresight horizon as the average of the maximum and minimum possible payoffs.[4] None of these assumptions about payoffs beyond the horizon can explain the behavior we observe but perhaps because with at most three periods, all payoffs are within subjects' horizons.

Ke [2019] uses a decision maker's preferences over pairs of individual decision trees to characterize how she evaluates what lies beyond her limited foresight. Our investment games do not belong to the class of decision trees studied in Ke [2019], because, in his model, the choice sets the decision-maker faces in every step of the tree depend deterministically on past choices. However, we adapt our generalized mean model from the functional form he axiomatically characterizes.

The noisy self model we construct is directly related to the concept of quantal response equilibrium, QRE [McKelvey and Palfrey, 1995, Goeree et al., 2020], in which players are more likely to choose best responses, but do not play them with probability one. QRE and its extension to extensive form games (AQRE) have been greatly successful in explaining departures from standard predictions in multi-player games [McKelvey and Palfrey, 1998]. As an internally consistent equilibrium model, AQRE requires players' beliefs to match the equilibrium probability distribution of actions. Herein lies the crucial difference between the noisy self model we construct and AQRE: the noisy self model allows probabilistic beliefs about future actions even when future selves play best responses *with probability one,* so that beliefs need not match actions. This relaxation of rational expectations is critical for explaining the data.

2. The Decision Problems

Figure 1.1 shows the two-period decision and its equivalent (under rationality) lottery. We restrict the parameters to make the problem non-trivial: $R_F^+ > R_1 \geq R_F^O \geq R_F^-$. $R_F^+ > R_1$ guarantees that investing is not dominated and $R_F^O \geq R_F^-$

[4] Rampal [2022] uses an open-ended curtailment rule that is common knowledge among players to define an equilibrium concept where subjects can learn and best respond to their beliefs about opponent's limited foresight. See also Heller [2015] in which players choose their foresight horizon subject to a cost.

FIGURE 2.1. Three-Period Decision Problem

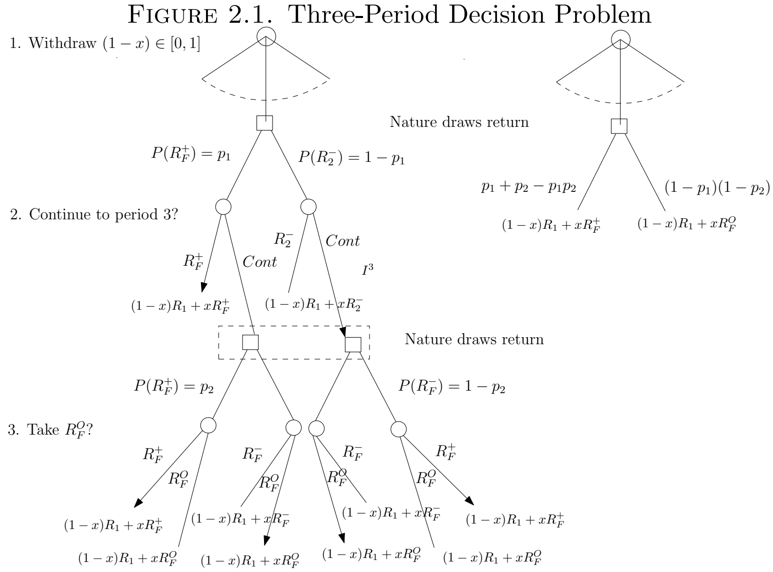

Notes: The first and last periods of the three-period decision problem resemble those of the two-period problem. In the intermediate second period, the subject can choose to withdraw everything remaining in the investment, x, after learning the rate that withdrawals earn (R_F^+ or R_2^-).

guarantees that a rational player recognizes that withdrawing at R_F^- is always (weakly) dominated.

A rational decision-maker that is deciding how much to withdraw in the first period, $1-x$, must use backwards induction to first determine the possible returns for the amount invested, x. Under the parameter restrictions given above, a rational decision-maker will realize that they will withdraw when R_F^+ realizes and not withdraw (take the outside option) otherwise. The first-period maximization problem is then

$$\max_x pu\left(xR_F^+ + (1-x)R_1\right) + (1-p)u\left(xR_F^O + (1-x)R_1\right) \tag{2.1}$$

assuming the decision-maker maximizes expected utility with utility function, $u()$. Critically, R_F^- does not enter the problem - it is irrelevant because the decision-maker never withdraws when it realizes. Varying R_F^- to see whether or not it affects initial investment is a key comparative static of interest.

In the three-period decision problem shown in Figure 2.1, the first and final periods are identical to those in the the two-period problem. In the first period, subjects start by deciding the amount $1-x$ to withdraw at a return R_1. In the final third period, subjects can either withdraw at the realized rate (R_F^-or R_F^+), or take the outside option R_F^O. We add an intermediate second period, where the subject can choose to withdraw *all* of x.[5] The withdrawal return in the second period is either R_F^+ or R_2^-,and is revealed immediately after the first period choice. We make the highest possible return in the second period the same as in the third period to ensure that we can choose the remaining parameters such that the three-period problem is equivalent to a two-period problem while still requiring non-trivial contingent reasoning in the second period.[6] As before, we constrain the parameters such that $R_F^+ > R_1 \geq R_F^O \geq R_F^-$.

For a three-period problem, a rational decision-maker must apply two steps of backwards induction. The first step is foreseeing choices in the final period. Based on these forecasted choices, in the second period, the decision-maker must decide to withdraw or not, conditional on the realization of the second period withdrawal return, R_F^+ or R_2^-. Because R_F^+ is the highest return possible, withdrawing when R_F^+ is realized first-order stochastically dominates (FOSD) continuing so that a rational decision-maker should withdraw. Under the additional restriction, $R_2^- < R_F^O$, continuing when R_2^- realizes FOSD withdrawing so that a rational decision-maker should continue. Imposing $R_2^- < R_F^O$ guarantees that the optimal second period decisions are independent of a subject's risk-preferences.[7] The rational decision-maker's first-period maximization problem is then

(2.2)
$$\max_x (p_1 + p_2 - p_1 p_2) u \left(x R_F^+ + (1-x) R_1 \right) + (1-p_1)(1-p_2) u \left(x R_F^O + (1-x) R_1 \right)$$

From (2.1) and (2.2), it is clear that, under rationality, the three-period decision problem is equivalent to a two-period decision problem when $p = p_1 + p_2 - p_1 p_2$. Yet, behaviorally, it may be more complex because the decision-maker first has to recognize that R_F^- is irrelevant in the final period in order to then recognize that continuing in the second period when R_2^- realizes dominates withdrawing.

Note also that the expression in (2.2) is invariant to interchanging p_1 and p_2. We deliberately designed this feature into the three-period problems so that we could observe decisions under both versions, interchanging the values of p_1 and p_2. Without loss, assume $p_1 < p_2$. We call the problem *front-loaded* when it offers the highest return, R_F^+, in the second-period with the larger probability p_2. Conversely, when the higher return occurs with the lower probability, p_1, in the

[5] Forcing this decision to be binary greatly simplifies the problem.

[6] There exist three-period problems that are equivalent to two-period problems that require only trivial contingent reasoning in the second period. For example, if both second period withdrawal returns are less than the outside option in the final period, a rational decision-maker only needs to reason contingently about the final period to recognize that the second period is irrelevant (withdrawal is always dominated).

[7] If $R_2^- > R_F^O$, the final period lottery between R_F^+ and R_F^O may or may not be preferred to withdrawing at R_2^-.

second period, we call it *back-loaded*. Figure A.1 in Appendix A illustrates a pair of decision problems, one of which is front-loaded and one of which is back-loaded. Rational decision-makers will find the two problems equivalent, but a boundedly-rational decision-maker that only looks forward one period, would invest more in front-loaded problems than in the corresponding back-loaded problems.

3. Experimental Design

We designed two treatments, one for the two-period decision problems and one for the three-period decision problems in a between subjects design. We refer to the treatments simply as the two-period and three-period treatments, respectively. We varied the withdrawal returns and their associated probabilities using the parameters shown in Table 1.

For the two-period treatment, we designed twenty two-period decision problems, divided into five major groups numbered 1-5 in Table 1. Each problem within a group (for example, problems 1a to 1d), has the same (p, R_1, R_F^+, R_F^O) parameter combination. Only R_F^- varies within the group, the parameter that should be irrelevant to a rational investor's choice. Given this fact, for a rational investor, each decision problem within a group reduces to the same lottery task, one in which the safe return is R_1 and the risky return is R_F^+ with probability p and R_F^O with probability $1 - p$. Subjects in the two-period treatment therefore complete twenty two-period decision problems and five lottery tasks. We placed the lottery tasks after the decision problems, but within each set of tasks (decision problems and lottery tasks), randomized the order across subjects.[8]

The parameters for the two-period decision problems were chosen to satisfy the constraints discussed in Section 2 but also to meet several other goals. First, for groups 1-3, we set the expected value of the lottery to be just higher than R_1. As such, any deviations from rationality for a risk-neutral decision-maker should result in large decreases in investment. Across these three groups, we varied the probability of the highest return and its value jointly, holding the expected value fixed. Decision problems in group 4 are instead such that the expected value of the the lottery is slightly lower than R_1. This group therefore provides variation that allows us to identify risk-neutral subjects as those that withdraw everything in the first period. The decision problems in group 5 are such that investing first order stochastically dominates (FOSD) withdrawing, thereby providing a strong test for rationality - any deviation from investing everything indicates irrationality regardless of risk preferences. Finally, note that the parameters allow for several comparative static tests (e.g. the only difference between groups 2 and 4 is in R_F^+).

[8]We placed the lottery tasks after the decision problems because we did not want to cue subjects that the decision problems might be reducible to a lottery.

In the three-period treatment, we chose six of the two-period investment problems from Table 1 to construct three-period problems. We sacrificed intermediate values of R_F^- in order to generate additional variation along other dimensions. For each of the six problems, we used the parameters of the two-period problem along with $p_1 = .1$, $p_2 = .44$, and $R_2^- = 0$ or 0.9 to create four different three-period decision problems in a 2x2 design: two had $R_2^- = 0$ and two had $R_2^- = 0.9$, and two were such that the probability of R_F^+ in the second period was p_1 (back-loaded) and the other two were such that it was p_2 (front-loaded). As with the two-period decision problems, a rational decision-maker would reduce each decision problem to a simple lottery. Doing so in this case, however, requires subjects to both reason contingently and to reduce compound lotteries. Because subjects may have difficulty doing the latter, in the lottery tasks, we described the lottery over R_F^+ and R_F^O as a compound lottery with probabilities p_1 and p_2. Thus, we had six lottery tasks in total: two for each of the groups of problems labeled 1, 3, and 5. As in the two-period treatment, the order of decision problems was randomized across subjects, as was the order of the lottery tasks, but the lottery tasks were performed after the decision problems.

In the three-period treatments, we vary R_2^- between the two values of 0 and 0.9 because these are (almost) at the two extreme ends of the parameter range that satisfies the constraints discussed in Section 2. We chose p_1 and p_2 such that $p_1 + p_2 - p_1 p_2 = p$ (to a high degree of precision) so that the three-period decision problems are equivalent to the two-period decision problems in the same group for a rational decision-maker. Furthermore, we made the difference between p_1 and p_2 quite large so that we would be able to pick up any difference across the front-loaded and back-loaded problems.[9]

Summarizing, the design allows a rich set of comparisons to be made (on top the parameter comparative statics not listed here):

(1) The two-period decision problems can be compared to their lottery counterparts.
(2) The three-period decision problems can be compared to:
  (a) their two-period decision problem counterparts.
  (b) their compound lottery counterparts.
(3) The lotteries from the two-period treatment can be compared to the compound lotteries from the three-period treatment.

[9]It is for this reason that we chose not to construct three-period problems equivalent to the two-period problems with $p = 0.2$: p_1 and p_2 either have to be very similar or one of them has to be close to zero (in which case it might be ignored completely).

Decision Problem	Two-period						Three-period		
	p	R_1	R_F^+	R_F^-	R_F^O		p_1	p_2	R_2^-
1a	0.5	1.5	2.2	0	1		0.1	.44	0,0.9
1b	0.5	1.5	2.2	0.3	1				
1c	0.5	1.5	2.2	0.7	1				
1d	0.5	1.5	2.2	1	1		0.1	.44	0,0.9
2a	0.2	1.5	4	0	1				
2b	0.2	1.5	4	0.3	1				
2c	0.2	1.5	4	0.7	1				
2d	0.2	1.5	4	1	1				
3a	0.3	1.5	3	0	1		0.05	0.26	0,0.9
3b	0.3	1.5	3	0.3	1				
3c	0.3	1.5	3	0.7	1				
3d	0.3	1.5	3	1	1		0.05	0.26	0,0.9
4a	0.2	1.5	3	0	1				
4b	0.2	1.5	3	0.3	1				
4c	0.2	1.5	3	0.7	1				
4d	0.2	1.5	3	1	1				
5a	0.3	1	3	0	1		0.05	0.26	0,0.9
5b	0.3	1	3	0.3	1				
5c	0.3	1	3	0.7	1				
5d	0.3	1	3	1	1		0.05	0.26	0,0.9

TABLE 1. Decision Problem Parameters

Differences in (1) and (2) identify failures of backwards induction, while differences in (3) indicate a failure to reduce compound lotteries (on average, across subjects only).

3.1. **Implementation.** We recruited gender-balanced groups of subjects from the U.S. population for each treatment using the Prolific platform in April and May of 2022.[10] Subjects were directed to a website which used Qualtrics and custom JavaScript coded by the the authors to run the experiment. Subjects first read instructions (replicated in Appendix C) for the decision problems and then answered a series of comprehension questions which they had to answer correctly (with unlimited attempts) to proceed. After completing the decision problems,they then read further instructions for the lottery tasks, and again answered comprehension questions before proceeding to the lottery tasks. We targeted 250 subjects for each treatment, but ended up with 249 in the two-period treatment and 251 in the three-period treatment. Average earnings were \$4.15 for about 19 minutes of time (\$13.13/hour) in the two-period treatment and \$4.71 for about 25 minutes of time (\$11.56/hour) in the three-period treatment, exceeding the minimum wage rate on Prolific (\$8/hour) by about 50%.

[10]Average age: 39.6 in two-period and 41.1 in three-period.

4. Results

In Section 4.1, we present the results of the two-period decision problems, establishing that subjects fail to correctly predict their own future actions. In Section 4.2, we show similar results for three-period decision problems, finding that investment does not simply decrease in these more complicated problems. Instead, additional interesting comparative static results emerge. Finally, in Section 5, we structurally estimate several potential models of behavior, finding that a model in which subjects noisily forecast their future actions and cognitively discount future payoffs best explains the patterns in the data.

4.1. Two-Period Results.

We begin by testing whether or not subjects understand that they will never choose the low return, R_F^-, in the second period when making first period choices. To simplify the interpretation of these results, we remove all subjects that actually do choose dominated payoffs in the second period more than twice - approximately 15% of subjects.[11] We focus on the resulting dataset of 214 subjects for the remainder of the analysis.

Figure 4.1 summarizes investment choices in the first period of the two period decision problems, plotting investment in the $R_F^- \in \{0, 0.3, 0.7\}$ problems against that for the $R_F^- = 1$ problem. We use investment in the $R_F^- = 1$ decision problem as the benchmark, instead of the lottery, because it holds constant the framing across problems.[12] In Figure A.2 of Appendix A, we provide a similar figure comparing investment in the decision problems to investment in the lottery tasks, showing similar, but somewhat noisier results.

Rational individuals, recognizing that $R_F^- < 1$ will never be chosen in the second period, would invest independently of R_F^- so that all investments would lie on the 45 degree line. Instead, for all five groups of decision problems and all three R_F^- values, we observe less investment when $R_F^- < 1$. The 95 percent confidence intervals indicate that each of the differences is significant at the 5% level, providing our first main result.

Result 1. *Average investment significantly decreases when the irrelevant second-period return is less than that of the outside option, inconsistent with subjects rationally forecasting their own future actions.*

[11] These subjects are likely simply not paying attention because choosing a smaller monetary payment over a larger one is difficult to reconcile with any theory.

[12] We find that investment in the $R_F^- = 1$ problem is actually significantly *greater* than that in the lottery task, on average. This result suggests that framing as a decision problem versus a lottery is consequential.

Figure 4.1. Two-Period Investments

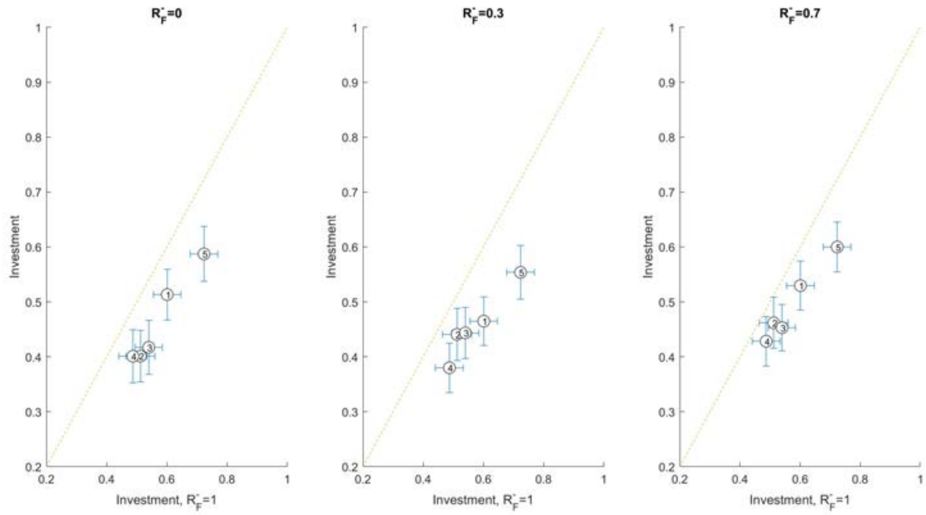

Notes: Each dot plots the investment in decision problems from groups 1-5, averaged across all subjects, under $R_F^- = 1$ on the x-axis, and under $R_F^- \in \{0, 0.3, 0.7\}$ on the y-axis. The error bars indicate 95% confidence intervals in both dimensions.

The results of Figure 4.1 suggest that subjects, at least on average, do not correctly forecast their future actions. One might conjecture that subjects should rapidly learn away from this failure of contingent reasoning: after participating in several problems, a subject should come to understand that they never actually receive R_2^- and thus learn to ignore it. However, we find almost no evidence of learning. Figure A.3 of Appendix A provides a figure corresponding to Figure 4.1, but restricting the data to only the second half of the decision problems, after 10 repetitions have already been completed. There, we see very similar patterns of reduced investment with $R_F^- < 1$.

We also explore learning paramaterically. Regardless of risk preferences, we would expect investment to increase as subjects learn that R_2^- is irrelevant and therefore expect investment to increase over time if learning plays a role. However, when we regress investment at the individual level on an indicator for the order of the task, while controlling for individual and problem fixed effects, we find a very precisely estimated zero coefficient on the task order (-0.007; $p = 0.84$). If we include interactions between the subject fixed effect and task order, only 3% of the interaction coefficients are significantly positive at the 5% level, suggesting that at most a handful of subjects learn to any significant extent.[13] Given little evidence

[13]We also ran regressions at the individual level, finding that only 6% of subjects have significantly increasing investment. We prefer the aggregate regressions because the individual results could be spurious due to the particular order of the investment problems they observed.

of learning, we continue to pool results across the order in which the problems were observed.

Result 2. *Subjects fail to learn to forecast their future actions at hypothetical events, despite repeated play.*

The comparative static with respect to R_2^- suggests some degree of irrationality, but other comparative statics are as standard theories would predict.[14] Specifically, problems from group 5 should have higher investment than those in group 3 because the parameters are identical except for R_1, the outside option is lower in group 5 problems. Using a paired t-test with the average (across R_2^- values) individual investment as an observation, we find a highly statistical difference ($p < 0.001$), as visible in Figure 4.1. Similarly, problems in group 2 should have higher investment than those in group 4 because R_F^+ is larger, and those in group 3 should have higher investment than those in group 4 because p, the probability of R_F^+, is larger. These predictions are both seen to be correct in Figure 4.1, although neither difference is statistically significant at the 5% level. Subjects therefore seem to respond to the parameters of the environment as predicted, apart from irrationally responding to changes in R_F^-. These results immediately rule out theories in which subjects hold constant beliefs about the probability of each return in the second period (e.g. Laplacian beliefs).

Result 3. *Subjects significantly respond to R_1 and respond to R_F^+ and p in directions consistent with standard theory, thus ruling out theories in which subjects hold constant beliefs about the probability of each return.*

Looking at average investments masks individual heterogeneity - what percentage of individual subjects invest less when R_F^- decreases? Figure 4.2 plots the empirical cdf of the average (across decision problems) difference between an indvidual's investment when $R_F^- = 1$ and each of the other three values. For all three values, we observe that roughly 60% of subjects have negative average differences, indicating less investment at lower R_F^- values. Averaging over all R_F^- values less than one (dashed line), we find that the differences across different R_F^- values are correlated - the same subjects appear to invest less across all three R_F^- values. These results provide evidence that the average deviations from rationality are concentrated among a subset of subjects, but those that make up a majority. We explore this heterogeneity further when estimating models in Section 5.

The behavioral theories we consider in Section 5 require subjects to respect monotonicity in payoffs *and* in probabilities. The lottery task corresponding to problem

[14] We formally show that the intuitive comparative statics we consider do indeed hold in a general behavioral model. See Proposition 1 of Appendix B.

FIGURE 4.2. Two-period Investments by Individual

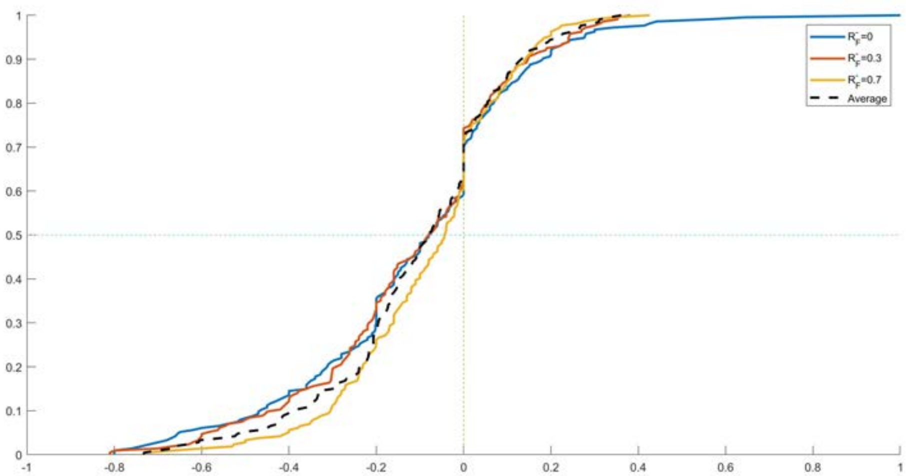

Notes: Empirical cdf of the individual differences between investment with $R_2^- = 1$ and the other three values (solid lines). Each difference is averaged across decision problems. The dashed line represents the individual differences averaged across $R_F^- = 0$, $R_F^- = 0.3$, and $R_F^- = 0.7$.

5 provides a test for monotonicity in probabilities: the investment return first-order stochastically dominates (FOSD) the first-period withdrawal return. 44% of subjects invest one in this lottery task, thus satisfying FOSD. Given that a majority of subjects violate FOSD, one may wonder if the results thus far are being driven by this subset of subjects. Figures 4.3 and 4.4 provide plots corresponding to Figures 4.1 and 4.2 for the 44% of subjects that do not violate FOSD.

Comparing the figures, we find that the results are actually *stronger* among the subset of subjects that satisfy FOSD. All of the average investments in Figure 4.3 with $R_F^- < 1$ remain significantly less than those with $R_F^- = 1$ even though the confidence intervals are wider. Furthermore, Figure 4.4 shows that close to 90% of this subset of subjects invest lower amounts when $R_F^- < 1$ (compared to only 60% in the full dataset). One interpretation of this result is that subjects that violate FOSD are likely inattentive and thus introducing noise.

Result 4. *Even subjects that satisfy FOSD in the lottery tasks do not forecast their own future actions correctly, investing less when the irrelevant second-period return is less than the outside option.*

4.2. **Three-Period Results.** As before, we exclude subjects that choose dominated payoffs in the third period more than twice - 10% of subjects - resulting in a dataset of 226 subjects. Figures 4.5 and 4.6 plot average investments and the empirical cdf

Figure 4.3. Two-Period Investments (Subjects Satisfying FOSD)

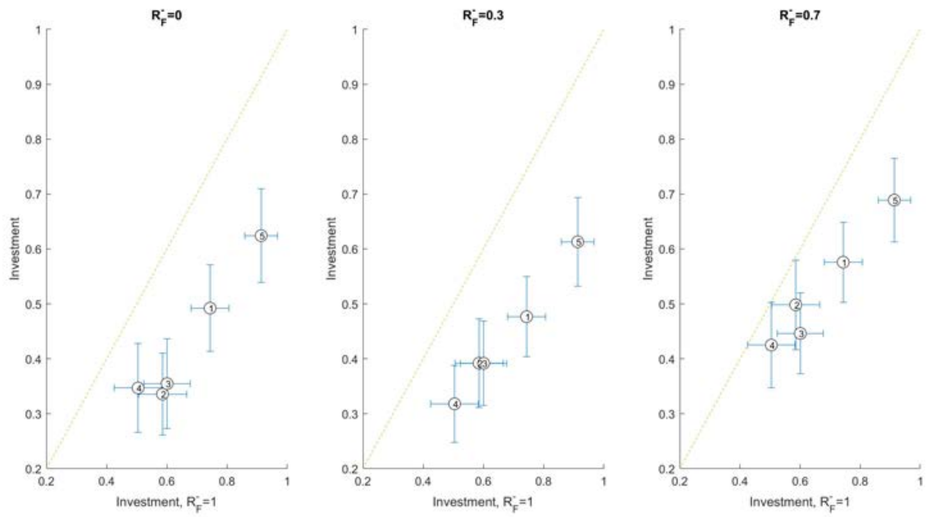

Notes: Each dot plots the investment in decision problems from groups 1-5, averaged across all subjects that satisfy FOSD in the lottery task, under $R_F^- = 1$ on the x-axis, and under $R_F^- \in \{0, 0.3, 0.7\}$ on the y-axis. The error bars indicate 95% confidence intervals in both dimensions.

Figure 4.4. Two-period Investments by Individual (Subjects Satisfying FOSD)

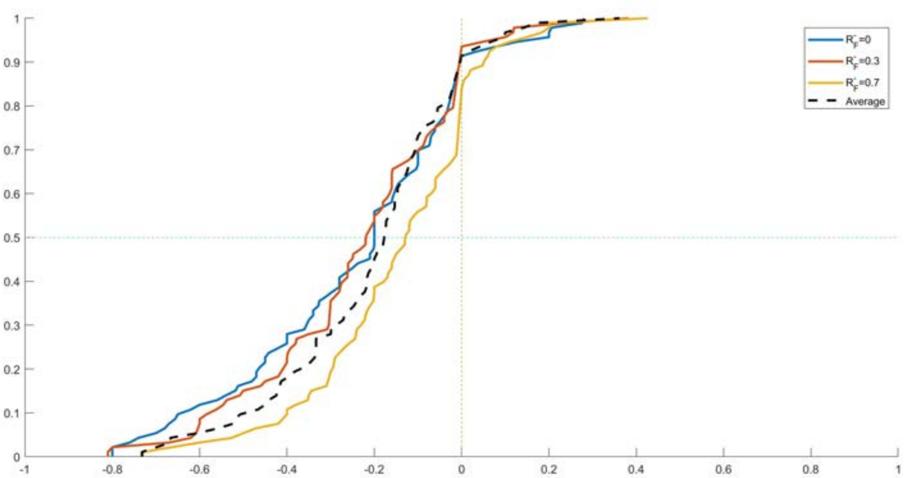

Notes: Empirical cdf of the individual differences between investment with $R_F^- = 1$ and the other three values (solid lines) for the subset of subjects that satisfy FOSD in the lottery task. Each difference is averaged across decision problems. The dashed line represents the individual differences averaged across $R_F^- = 0$, $R_F^- = 0.3$, and $R_F^- = 0.7$.

FIGURE 4.5. Three-Period Investments

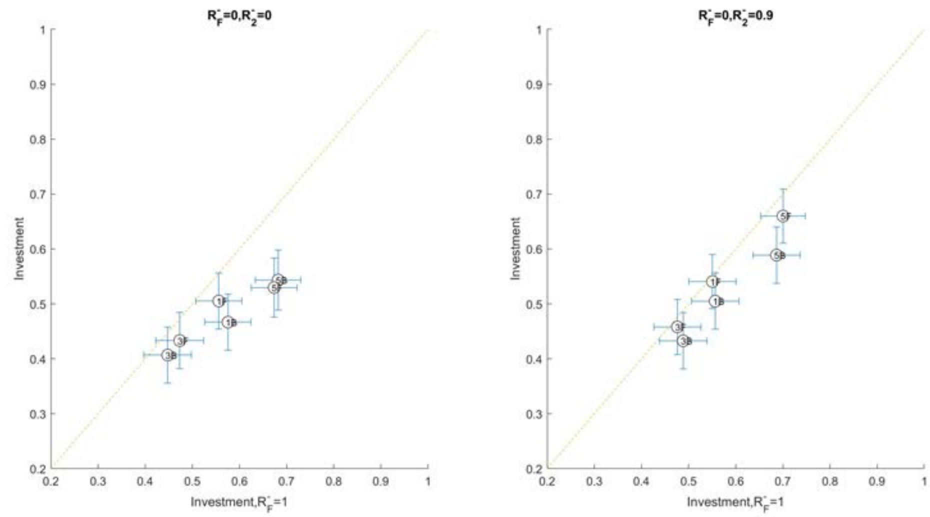

Notes: Each dot plots the investment in decision problems from groups 1, 3, and 5, averaged across all subjects, under $R_F^- = 1$ on the x-axis, and under $R_F^- = 0$ on the y-axis. For each decision problem, we plot the front-loaded (F) and back-loaded (B) versions of the problem separately. The left plot corresponds to $R_2^- = 0$ and the right to $R_2^- = 0.9$. The error bars indicate 95% confidence intervals in both dimensions.

of individual differences in investment, respectively.[15] Recall from Table 1 that for the three-period decision problems, we have only two R_F^- values, 0 and 1, but we also have variation across two R_2^- values, 0 and 0.9. Also, we only have problems 1, 3, and 5 with two versions of each, one of which the probability of the highest return is front-loaded (F) in the second period and one in which it is back-loaded (B) in the third period.

The three-period treatment replicates the major patterns from the two-period treatment. Figure 4.5 shows that, as in the two-period treatment, subjects react to the inconsequential return: they invest less when $R_F^- = 0$ than when $R_F^- = 1$, although not always significantly. They also rationally invest more in decision problems from group 5 than group 3 because R_1 is lower in group 5, a finding we confirm is significant at the 5% level. We again find little evidence of learning - Figure A.5 of Appendix A shows similar patterns to Figure 4.5 when we restrict to the second half of the data after 13 decision problems have been completed. Parametrically, the coefficient on task order in an aggregate regression is -0.004 ($p = 0.781$) and with interaction effects, only 4.4% of subjects significantly increase investment over

[15]We include the comparison between lotteries and decision problems in Figure A.4 of Appendix A.

Figure 4.6. Three-period Investments by Individual

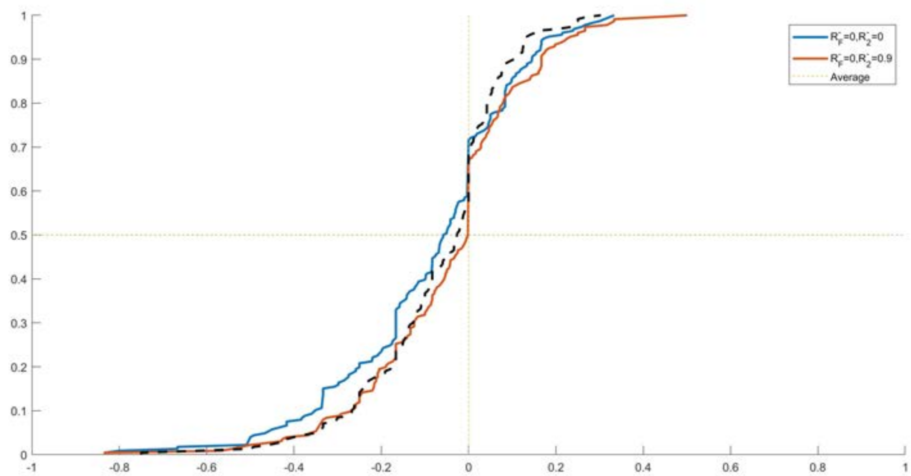

Notes: Empirical cdf of the individual differences between investment with $R_F^- = 1$ and $R_F^- = 0$ separately for $R_2^- = 0$ and $R_2^- = 0.9$ (solid lines). Each difference is averaged across decision problems. The dashed line represents the individual differences averaged across $R_2^- = 0$ and $R_2^- = 0.9$.

time.[16] Finally, Figure 4.6 shows that the differences across decision problems with different R_F^- values are again concentrated among a subset of subjects.

Result 5. *Three-period problems replicate the main results from the two-period problems. Investment significantly increases when R_1 decreases. Investment decreases, although not always significantly, when the irrelevant final period return is less than the outside option. Subjects fail to learn with experience.*

As before, we can also filter out subjects that violate FOSD in the lottery tasks.[17] Here, subjects make decisions in two lotteries, corresponding to decision problems 5a and 5d, for which the investment returns first order stochastically dominate the safe return. We filter out subjects that violate FOSD in either lottery task but note that 80% of subjects that violate FOSD in either lottery violate it in both. Figures 4.7 and 4.8 provide the plots corresponding to Figures 4.5 and 4.6 for the 42% of subjects that do not violate FOSD in either lottery. We find similar levels of

[16]In individual regressions, only 4.4% have significantly positive coefficients on task order indicating learning. In three-period decision problems, subjects may be learning about rational decisions both in the second and third periods, but we expect learning to increase investment as in the two-period decision problems.

[17]Subjects can also violate FOSD in their second period decisions. 21% of decisions are to continue when R_2^+ is realized and 21% of decisions are to withdraw when $R_2^- = 0$ is realized (26% when $R_2^- = 0.9$ is realized). To facilitate comparisons with the two-period problems, we do not filter out these violations of FOSD.

Figure 4.7. Three-Period Investments (Subjects satisfying FOSD)

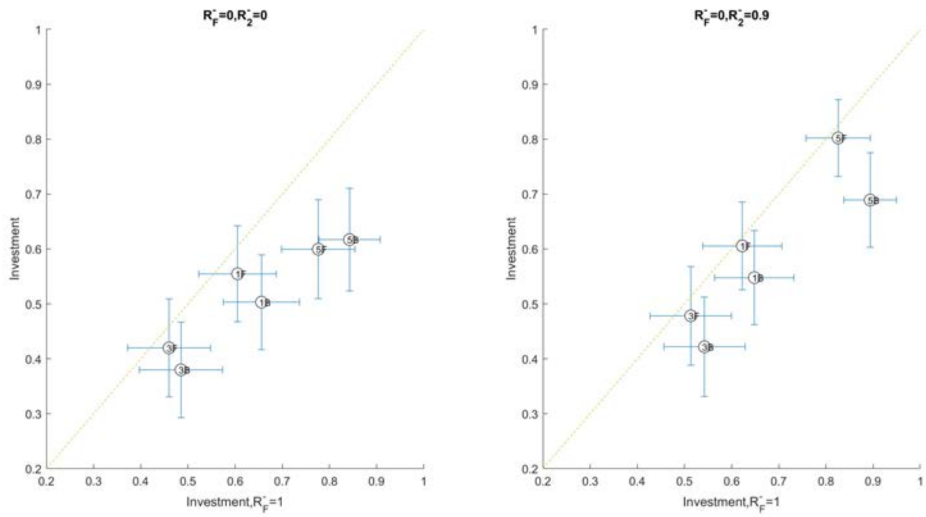

Notes: Each dot plots the investment in decision problems from groups 1, 3, and 5, averaged across subjects that satisfy FOSD in the lottery tasks, under $R_F^- = 1$ on the x-axis, and under $R_F^- = 0$ on the y-axis. For each decision problem, we plot the front-loaded (F) and back-loaded (B) versions of the problem separately. The left plot corresponds to $R_2^- = 0$ and the right to $R_2^- = 0.9$. The error bars indicate 95% confidence intervals in both dimensions.

irrationality among those that satisfy FOSD. Overall, the results of the three-period decision problems provide additional support for Results 1-4 that were based on the two-period data.

The three-period problems generate additional comparative statics of interest. Here, we focus on the subset of subjects that satisfy FOSD (the results are similar in the overall sample, but noisier). Rational subjects should not respond to the dominated second period return R_2^- or to the probability of the best return being front-loaded versus back-loaded.[18] However, we find more investment with $R_2^- = 0.9$ (0.63) than with $R_2^- = 0$ (0.58), a difference that is significant ($p < 0.001$ with paired t-test of average individual investments).

We find no difference between the front-loaded and back-loaded decision problems (0.61 and 0.60, respectively; $p = 0.82$ with paired t-test) when averaging across all R_F^- and R_2^- values, but we can observe a clear pattern in Figure 4.7: the front-loaded versions generally lie up and to the left of the back-loaded versions. The average investment in front-loaded problems is 0.58 versus 0.53 in back-loaded problems when $R_F^- = 0$, but 0.63 versus 0.68 when $R_F^- = 1$, both of which are

[18]When comparing the front-loaded versus back-loaded versions of the lotteries, we find no difference, and also no difference with the corresponding reduced lotteries from the two-period treatment, suggesting that subjects seem to reduce compound lotteries correctly ($p > 0.44$ for all three comparisons via t-tests).

FIGURE 4.8. Three-period Investments by Individual (Subjects satisfying FOSD)

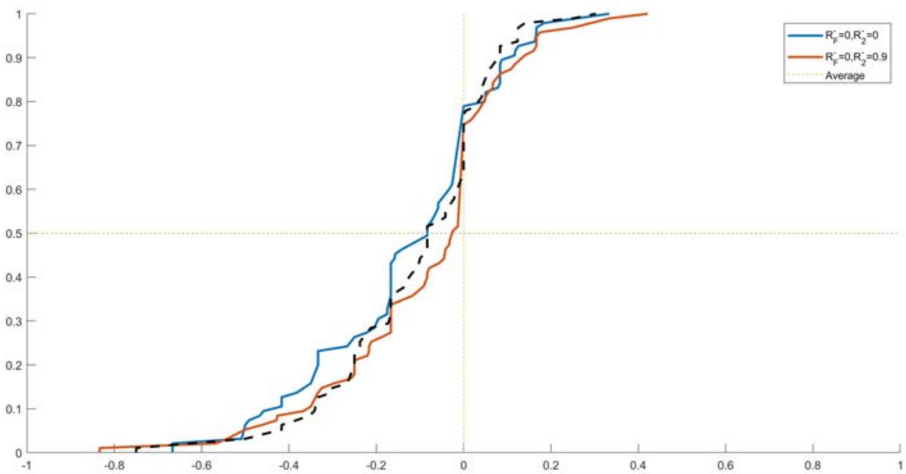

Notes: Empirical cdf of the individual differences between investment with $R_F^- = 1$ and $R_F^- = 0$ separately for $R_2^- = 0$ and $R_2^- = 0.9$ (solid lines) for the subset of subjects that satisfy FOSD in the lottery tasks. Each difference is averaged across decision problems. The dashed line represents the individual differences averaged across $R_2^- = 0$ and $R_2^- = 0.9$.

significant differences ($p = 0.01$ for each with two-sample t-tests). Thus, contrary to the rational prediction, subjects do respond to front versus back-loading the probability of getting the high return, but in a subtle way that depends upon the lowest possible return in the final period.

Result 6. *Subjects respond to the second-period returns and to how the probability of the highest possible return is distributed across the second and third periods in ways a rational subject would not.*

To make comparisons between the two-period and three-period decision problems, we use only the comparable decision problems across the two treatments, those from groups 1, 3, and 5 and with $R_F^- \in \{0, 1\}$, and we continue to focus on the subsets of subjects that do not violate FOSD in the lottery tasks. We find no overall difference between the two and three-period problems: average investment is 0.62 and 0.60, respectively ($p = 0.62$ via two-sample t-test), contrary to the hypothesis the complexity is driving the difference across treatments.

However, when we break the average investment down by R_F^- as in Table 2, we see that the lack of a difference on average masks substantial heterogeneity: when $R_F^- =$ 0, subjects invest less in the two-period decision problems (although not significantly

Table 2. Comparison of Two and Three-Period Problems

	Two-period	Three-period
$R_F^-= 0$	0.49	0.55
$R_F^-= 1$	0.75	0.66

Notes: Average investment across comparable decision problems in the two and three-period treatments.

so, $p = 0.17$ in two-sample t-test), but when $R_F^- = 1$, subjects significantly invest more ($p < 0.001$ in two-sample t-test). Thus, subjects respond less to changes in the final period return in the three-period problems. The differences in average investments between problems with $R_F^- = 0$ and $R_F^- = 1$ are -0.26 in two-period decision problems and -0.10 in three-period problems. The difference in differences is highly significant in a two-sample t-test ($p < 0.001$). Combining these results with the fact that investment is higher in the three-period problems with $R_2^- = 0.9$ versus $R_2^- = 0$ provides some suggestive evidence for a form of discounting, something which we consider further in the following section.

5. Model Estimation

The reduced-form results of the previous sections paint a picture of bounded rationality. On one hand, subjects do not forecast their future actions correctly when making initial investment decisions. They respond to the irrelevant worst-case return in the final period and respond to changes in probabilities and intermediate returns in the three-period problems in ways not predicted by a rational model. On the other hand, they respond to changes in other returns as a rational actor would. Over the next five subsections, we describe several candidate models of boundedly rational behavior (5.1), and estimate and compare them. We first estimate CRRA risk-preference parameters from the lottery tasks (5.2) and then use those to estimate the models parametrically by individual (5.3). We then compare the models (5.4) and show that the fit is better when we allow for cognitive discounting (5.5). We conclude with a discussion of the model that fits best (5.6).

5.1. **Models.** We start by specifying a somewhat general behavioral model for a two-period decision problem, in which, in the first period, the subject believes that she chooses R_F^+ and R_F^- with probabilities q_1 and q_2 conditional on each return realizing.[19] This model nests the fully rational model with $q_1 = 1$ and $q_2 = 0$, but allows the subject to imperfectly perceive second-period choices. The first-period

[19]Thus the outside return R_F^O is chosen with probabilities $1 - q_1$ and $1 - q_2$ when the R_F^+ and R_F^- returns realize, respectively.

maximization problem is

$$\max_x pq_1 u\left(xR_F^+ + (1-x)R_1\right) + (1-p)q_2 u\left(xR_F^- + (1-x)R_1\right)$$
$$+ \left[p(1-q_1) + (1-p)(1-q_2)\right] u\left(xR_F^O + (1-x)R_1\right) \tag{5.1}$$

In it's full generality, it is difficult to behaviorally interpret q_1 and q_2 in this model, but additional assumptions generate popular classes of behavioral models. We consider five such models here, and discuss how each can be extended to three periods.

1) The *rational* model: $q_1 = 1, q_2 = 0$. This has an obvious extension to the three-period case and has zero free parameters.

2) The *naïve* model in which the decision-maker ignores the outside option, R_F^O completely: $q_1 = 1, q_2 = 1$. This model also has zero free parameters. The extension to three periods is straightforward - we assume the decision-maker applies backwards induction, subject to the constraint that the outside option in the final period is ignored.

3) The *cursed* model, a model based upon Eyster and Rabin [2005]'s concept of cursedness. A fully cursed subject has correct beliefs 'on average' but ignores the correlation between her future action and realized return. Thus, she thinks that both of the high and low returns are chosen at the true unconditional probability,

$$q_1 = q_2 = p \times 1 + (1-p) \times 0 = p$$

As in Eyster and Rabin (2005), we allow a decision-maker to be partially cursed: fully cursed with probability χ and fully rational with probability $1-\chi$, so that beliefs become

$$q_1 = \chi p + 1 - \chi, q_2 = \chi p$$

χ is the single free parameter.

In three-period games, we assume that the decision-maker is either fully cursed and understands none of the correlations between actions and payoffs (in neither the second nor third periods) or is fully rational and understands them perfectly. A fully cursed decision-maker's beliefs about her final period action are exactly as described above for the two-period fully cursed agent with p_2 replacing p. Her belief about her second period actions is that she continues with probability $1-p_1$, the probability with which a rational decision-maker continues, irrespective of the realized return. For a partially-cursed decision-maker, the probability of continuing in the second period is the χ-weighted average of the fully cursed and fully rational cases: that is, $\chi(1-p_1)$ when R_2^+ realizes and $\chi(1-p_1)+(1-\chi)$ when R_2^- realizes.

4) The ϵ-tremble or *tremble* model is one in which a subject thinks that at every decision node, she takes the rational action with probability $1-\epsilon$ and takes the

other action with probability ϵ. Thus, $q_1 = 1-\epsilon$ and $q_2 = \epsilon$. This model has one free parameter, ϵ.

This model has an additional interpretation in terms of ambiguity. Suppose that subjects perceive the true future return $R_2^*(+)$ from the hypothetical event when R_2^+ realizes, as being ambiguous and distributed over the range $[R_2^O, R_2^+]$ with full support. Similarly, they perceive the true future return $R_2^*(-)$ from the hypothetical event when R_2^- realizes, as being ambiguous and distributed over the range $[R_2^-, R_2^O]$ with full support. Then, an α-maxmin [Ghirardato et al., 2004] subject who assigns probability $\alpha = \epsilon$ to the best case scenario (best returns from the ranges realizing) and $1-\alpha$ to the worst case scenario (worst returns from the ranges realizing), would behave exactly as in the ϵ-tremble model.

In three-period games, the extension is straightforward - trembles in the second period occur with probability ϵ, but the ambiguity interpretation in this case is not as straightforward.

5) The *noisy self* mode is similar to quantal response equilibrium: a decision-maker believes that she will make correct decisions in future periods in proportion to utility differences, but relaxes the assumption of her future actions matching those beliefs.[20] Thus, in this model, q_1, q_2 depend on the expected utilities from the two choices in the following way:

$$\begin{aligned} q_1 &= \frac{\exp\left(\nu u\left(xR_F^+ + (1-x)R_1\right)\right)}{\exp\left(\nu u\left(xR_F^+ + (1-x)R_1\right)\right) + \exp\left(\nu u\left(xR_F^O + (1-x)R_1\right)\right)} \\ q_2 &= \frac{\exp\left(u\left(\lambda xR_F^- + \lambda(1-x)R_1\right)\right)}{\exp\left(\nu u\left(xR_F^- + (1-x)R_1\right)\right) + \exp\left(\nu u\left(xR_F^O + (1-x)R_1\right)\right)} \end{aligned} \tag{5.2}$$

Contrary to quantal response equilibrium, our model allows subjects to make correct decisions with probability 1 in the future, despite holding probabilistic beliefs. This model has one free parameter, ν, which measures the sensitivity to payoffs. $\nu \to \infty$ corresponds to rational choice.

In the three-period game, the beliefs about third-period actions are the same as above. We derive beliefs about second-period actions recursively again using the logistic function. The expected utility of withdrawing is simply a function of the realized withdrawal rate, R_F^+ or R_2^-. The expected utility of continuing depends upon the beliefs over final period returns, q_1 and q_2, and the objective probability, p_2. The probabilities of continuing are then formed as in (5.2).

In addition to these belief-based models, we consider a sixth model inspired by the work of Ke [2019]. The *generalized mean* model differs from the above models

[20] We use utility differences rather than payoff differences, because in the three-period model, the extension to using payoffs would be to use expected value for second-period choices. But, using expected utility would mean that subjects ignore their own risk preferences when forming beliefs. Qualitatively, none of our conclusions change if we use payoffs instead of utilities.

in that the decision-maker's misperceptions enter through aggregations in payoffs (utilities) in a way which cannot be modeled through expected utility with mistaken beliefs. Specifically, the decision-maker aggregates utilities using a generalized mean. For the two-period decision problem, when R_F^+ realizes, the decision-maker perceives the utility

$$u^+ = \left(\frac{1}{2}\left(u(xR_F^+ + (1-x)R_1)\right)^{\gamma} + \frac{1}{2}\left(u(xR_F^O + (1-x)R_1)\right)^{\gamma}\right)^{1/\gamma}$$

Here, γ captures the degree of rationality, encompassing several possibilities. When $\gamma = 1$, we have the average of the two possible utilities, when $\gamma \to \infty$ we have the maximum utility (rational), and when $\gamma \to -\infty$ we get the minimum utility (deciding with the worst case scenario in mind), assuming utilities are positive.[21] We can similarly define u^- when R_F^- realizes by replacing R_F^+ by R_F^- in the above expression. The decision-maker then maximizes the expected utility, $pu^+ + (1-p)u^-$. For the three-period problems, we define the overall utility recursively.[22]

5.2. **Estimating risk-preferences.** We begin by estimating the risk preferences of each subject from their investment choices in the lottery tasks only.[23] We follow Friedman et al. [2022] who show how to transform the lottery problem in a way that allows preferences to be estimated via OLS, assuming a CRRA utility function, $u(x) = \frac{x^{1-\alpha}}{1-\alpha}$. In this case, the first-order condition for the investment, x, that maximizes expected utility in a reduced lottery task is[24]

$$p\left(R_2^+ - R_1\right)\left(xR_2^+ + (1-x)R_1\right)^{-\alpha} + (1-p)\left(R_F - R_1\right)\left(xR_F + (1-x)R_1\right)^{-\alpha} = 0$$

After rearranging and taking the logarithm, this first-order condition can be written

$$\log\frac{(xR_F + (1-x)R_1)}{\left(xR_2^+ + (1-x)R_1\right)} = -\frac{1}{\alpha}\log\frac{p\left(R_2^+ - R_1\right)}{(1-p)\left(R_1 - R_F\right)} \tag{5.3}$$

[21]When utilities are negative, which for the CRRA utility function we specify below will occur when $\alpha > 1$, the generalized mean becomes undefined. Because when $\alpha > 1$, all utilities are negative, we handle this case in estimation by defining the generalized mean as the negative of the generalized mean of the negative of the utilities. In this case, $\gamma \to -\infty$ corresponds to rational behavior.

[22]For example, in the second period, when R_2^- realizes, utility is $u^- = \left(\frac{1}{2}\left(u(xR_2^- + (1-x)R_1)\right)^{\gamma} + \frac{1}{2}\left(u^O\right)^{\gamma}\right)^{1/\gamma}$ where u^O is the utility from continuing to the third period which is constructed as in the two-period model, $u^O = p_2u^+ + (1-p_2)u^-$.

[23]We experimented with estimating risk aversion along with the other parameters via maximum likelihood, but ran into convergence problems and thus adopted the two-step procedure.

[24]Given evidence that subjects reduced the compound lotteries in the three-period treatment, we simply estimate the reduced-form lottery.

FIGURE 5.1. Risk Preferences

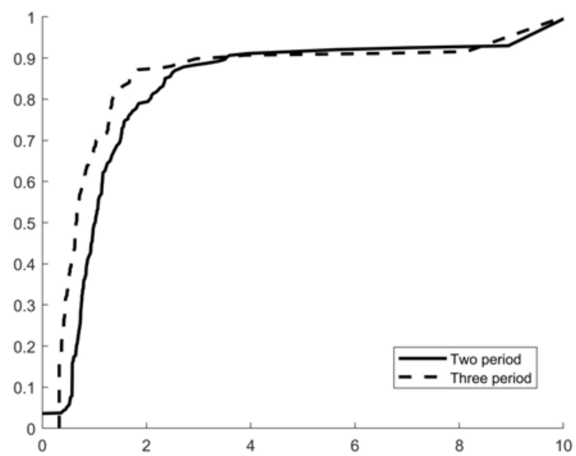

Notes: Empirical cdfs of the estimates of the CRRA coefficient, α, for each subject estimated from the lottery choices in each treatment.

The left-hand side of (5.3) can be constructed from the investment choices observed in the data. Positing a normally-distributed error term, we can then regress the left-hand side on the logarithmic term on the right-hand side to obtain an estimate of α. Note that we must exclude the lotteries in which the investment returns first-order stochastically dominate the safe return ($R_1 = R_F$), but these are uninformative about risk preferences in any case. We therefore have four data points for each subject in each treatment. We continue to exclude subjects which we excluded from the reduced-form analysis - those who chose dominated payoffs more than twice in the final period. Figure 5.1 plots empirical cdfs of the estimated risk preferences in each treatment.

With the exception of four subjects in the two-period treatment, all subjects are estimated to be risk-averse to varying degrees. The median estimate of α is 1.02 in the two-period treatment and 0.66 in the three-period treatment.[25]

5.3. **Estimating Model Parameters.** Taking the estimated risk preferences as fixed, we then use the data from the decision problems to estimate each of the twelve candidate models.[26] We estimate each model separately by individual.[27]

[25]The fact that subjects are estimated to be more risk-averse in the two-period treatment is likely due to the inclusion of a lottery in which the probability of the high return is only 0.2 - this lottery does not appear in the three-period treatment.

[26]To avoid numerical issues with the estimation procedure, we bound α between -1 and 5. These limits bind for about 10% of subjects in each treatment.

[27]We estimated finite mixture models as well, obtaining qualitatively similar results and reaching the same conclusions in terms of which models do best. But, given that we have sufficient data at the subject level, we prefer the individual estimates.

For the two-period treatment, we use only the twenty first-period investment decisions (i.e. we ignore the second-period decisions because they are almost always rational given we have removed subjects that make dominated choices). For each candidate model, we calculate the expected utility for each possible investment, $x \in [0, 1]$, on a grid-size of 0.01, as in the data. We then use a logit error structure to calculate the probability of observing the actual investment choice, $\hat{x}$:

$$P(\hat{x}) = \frac{\exp\left(\lambda EU(\hat{x})\right)}{\sum_{x=0,0.01,\ldots,1} \exp\left(\lambda EU(x)\right)} \tag{5.4}$$

λ parameterizes the level of noise in the investment data - choices are random when $\lambda = 0$ and the model fits perfectly as $\gamma \to \infty$. This approach has been used in estimating quantal response equilibrium models [McKelvey and Palfrey, 1995, 1998] and has been suggested by Harrison and Rutström [2008] to be favorable to some from of non-linear least squares when corner choices are present because it treats corner choices identically to others. It also has the distinct advantage of not having to calculate optimal investment levels, which cannot be analytically calculated for some of the models. Given the individual choice probabilities, the log-likelihood in expression (5.4) is constructed simply by taking the logarithm of each choice probability and summing them (thus assuming independent errors in each problem, as is standard). We then estimate the parameters via maximum likelihood.

For the three-period model, we use both first-period investment choices and the second-period continuation decisions from the twenty-five decision problems. For the second-period continuation decisions, we construct the choice probability as in expression (5.4) using the amount invested in the first period, but for the binary choice of continuing versus withdrawing. We allow for a different noise parameter, λ, than that for first-period choices because we don't expect the distributions of noise to be similar for first and second-period choices. We then construct the likelihood assuming independent errors across choices in the first and second periods and across tasks so that the log-likelihood is again the summation of the logarithm of the probabilities of each choice.[28]

Figure 5.2 plots empirical cdfs of the estimates of the rationality parameters for each of the cursed, noisy self, tremble, and generalized mean models. In these figures, we restrict to subjects with estimated noise parameters, $\gamma > 0.001$, because for very noisy subjects, the other parameters of the model are not identified.[29] Overall, the figures indicate significant heterogeneity across subjects. About 40 to

[28]If nothing is invested in the first period, the second-period choice is irrelevant so contributes nothing to the likelihood.

[29]7.5 to 48 percent of subjects are estimated to have $\gamma < 0.001$, depending on the model.

FIGURE 5.2. Rationality Parameters Estimates

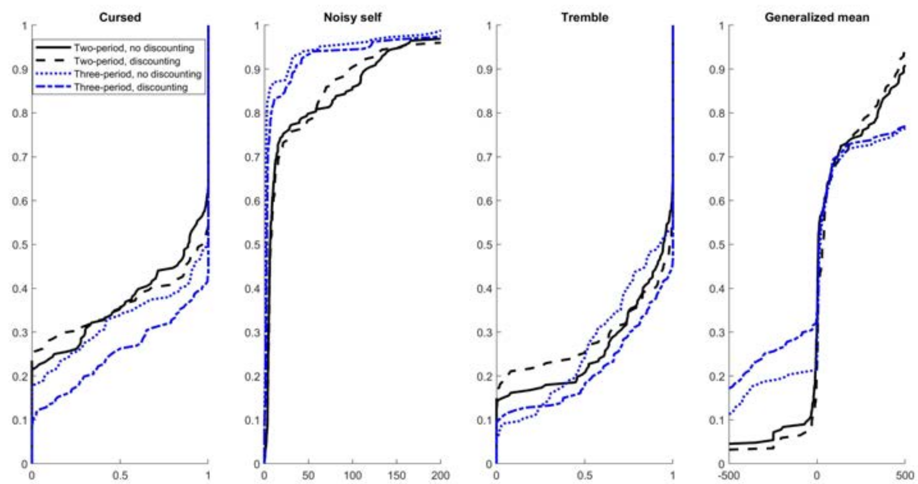

Notes: Empirical cdfs of the rationality parameters for each subject (χ, ν, ϵ, γ from left to right), estimated from the decision problems in the two-period and three-period treatments. Only estimates for subjects with noise parameter $\lambda > 0.001$ are included because the rationality parameters are not identified for small values of λ.

60 percent of subjects are estimated to be rational, depending on the model.[30] At the other extreme, approximately 10 to 30 percent of subjects are estimated to be completely irrational (but consistently so given that we are excluding very noisy subjects).

5.4. **Model Comparisons.** We use likelihood ratio tests to compare nested models (i.e. each model with and without cognitive discounting, as well as the rational models to the cursed, noisy self, and tremble models). We use Vuong tests to compare non-nested models (i.e. naive to the other four models and pairs of the cursed, noisy self and tremble models). We do comparisons both at the individual level and by aggregating the likelihoods. By comparing each model to the rational model, we avoid the need to calculate standard errors for the estimated rationality parameters, which is problematic in our setting for two reasons. First, due to noise in the estimate of risk-aversion, α, which would have to be accounted for and, second, due to the fact that standard errors cannot be calculated using standard (outer product of gradient) methods when an estimate is at a corner.

[30]The levels of the rationality parameters required for a subject to be deemed rational in the noisy self and generalized mean models is somewhat arbitrary. In Section 5.6, we look at the probabilities with which subjects believe that they will make mistakes in the noisy self model, which are more easily interpreted.

Table 3. Model Comparisons : Two-period

Model	Rational	Naïve	Cursed	Noisy Self	Tremble	Generalized mean
Rational	-	0.07	-	-	-	-
Naïve	**0.16**	-	**0.53**	0.23	**0.53**	**0.50**
Cursed	**0.22**	0.08	-	0.02	0.01	0.07
Noisy self	**0.55**	**0.22**	**0.46**	-	**0.38**	**0.37**
Tremble	**0.24**	0.08	**0.15**	0.02	-	0.08
Generalized mean	**0.35**	0.12	**0.16**	0.05	**0.08**	-

Notes: Fraction of subjects for which the row model beats the column model by a model comparison test at the five percent level. For nested models, only a one-way comparison is done via a likelihood ratio test. For non-nested models, two-way comparisons are done using Vuong tests. Boldface entries indicate the row model beats the column model in the aggregate at the five percent level.

Table 4. Model Comparisons : Three-period

Model	Rational	Naïve	Cursed	Noisy Self	Tremble	Generalized mean
Rational	-	**0.22**	-	-	-	-
Naïve	0.04	-	0.31	0.18	0.24	0.11
Cursed	0.01	**0.14**	-	0	0.08	0.14
Noisy self	**0.52**	**0.39**	**0.39**	-	**0.42**	**0.35**
Tremble	**0.15**	**0.18**	**0.10**	0.05	-	**0.15**
Generalized mean	**0.34**	**0.17**	0.15	0.10	0.13	-

Notes: Fraction of subjects for which the row model beats the column model by a model comparison test at the five percent level. For nested models, only a one-way comparison is done via a likelihood ratio test. For non-nested models, two-way comparisons are done using Vuong tests. Boldface entries indicate the row model beats the column model in the aggregate at the five percent level.

Table 3 reports the results of the model comparison tests for the two-period decision problems. Each entry indicates the fraction of subjects for which the row model is better than the column model, and boldface entries indicate that the model is better in the aggregate (at five percent levels of significance). The clear winner is the noisy self model, which outperforms all of the other models in the aggregate. At the other extreme, the cursed model is outperformed by all of the other behavioral models. The tremble, naïve and generalized mean models are in the middle, with the naïve model outperforming the other two: the extra parameter in each of the tremble and generalized mean models does not significantly improve their fit.

Table 4 reports the results of the model comparisons for the three-period decision problems. We again see that the noisy self model is the clear front-runner. It beats every other model in the aggregate and explains at least 35 percent of individuals better than any other model. The tremble model, followed by the generalized mean model, are the only other models which fit better than the rational model in the aggregate.

One possibility is that some of the behavioral models fit better than others because of inattentive subjects. To test for this possibility, we redo the model comparison tests excluding subjects that violated FOSD in the lottery choices and those with risk-aversion estimates greater than one (very high risk aversion may be an artifact of choosing 'middle' values in the lottery tasks). After excluding these subjects, we have 57 and 66 subjects in the two-period and three-period treatments, respectively. Within these selected subsets of subjects, our results continue to hold: the noisy self model dominates all of the other models in both the two-period and three-period decision problems as before. In fact, the results are actually stronger than in the full sample: in the three-period decision problems, noisy self beats each of the other five models for at least 58 percent of these subjects.

One last way to see the dominance of the noisy self model is to look at the fraction of subjects each model cannot explain at all (we use $\lambda < 0.001$ to identify such subjects). In the two-period decision problems, this fraction varies from 7.5 percent with noisy self to 48 percent with naïve. But, more interestingly the 7.5 percent of subjects whose behavior the noisy self model cannot explain, can also not be explained by any of the other models. In the three-period decision problems, the noisy self model cannot explain behavior of 19 percent of subjects, but the fraction of subjects that can't be explained by any model is 5.8 percent. In almost all cases, it is the generalized mean model that can explain the behavior of a subject that the noisy self model cannot explain, which likely reflects the fact that the generalized mean model allows subjects to choose the dominated payoff more than 50 percent of the time, while the noisy self model does not. Overall, the model estimates strongly suggest that the noisy self model best explains the data.

Result 7. *A model in which subjects believe they will make the correct choice more often when the payoff differences are larger (noisy self) best explains the data.*

An obvious question is why is it that the noisy self model does so much better than the other models? Given that the tremble model performs second best in the three-period decision problems, it seems important that a model allow for (perceived) probabilistic behavior in the second period. The other three models only allow mistakes in the second period when the perceived value of continuing to the third period is less than the realized second period payoff, which can only happen when R_2^- is higher (0.9). However, given that the noisy self model outperforms all of the other models in the two-period decision problems, this advantage does not appear to be the only reason for its success. Instead, note that the only difference between the tremble and noisy self models is that in the noisy self model, the decision-maker believes that the probability of taking the correct action in the future is proportional to the difference in utilities across actions rather than being constant. Given that the noisy self model fits much better than the tremble model, this proportionality seems to be crucial.

5.5. **Cognitive Discounting.** Intuitively, if subjects solve the decision problems forwards instead of backwards, they may place more weight on immediate payoffs than future payoffs. The fact that subjects invest more in the three-period problems when R_2^- increases, and respond less to R_F^- in these problems than the two-period problems is suggestive of some form of cognitive discounting or myopia.

To explore this possibility further, we consider variants of the previous models in which payoffs are exponentially discounted by some discount factor, $\delta < 1$: payoffs one period ahead are multiplied by δ and those two periods ahead by δ^2.[31] We emphasize that multiplying payoffs by δ is a simple, reduced-form way to capture a lowering of perceived future payoffs due to working forwards or due to uncertainty about which payoff will be received - we do not believe δ reflects literal discounting over time given that all payoffs are received at the same time.

Figure 5.3 plots the estimated discount parameters, δ, for all fix models in the two-period problems (left panel) and three-period problems (right-panel). The fraction of subjects with $\delta < 1$ varies from around 30% to 80% depending on the model and type of problem, but the estimates are highly correlated across models. Across the 30 possible pairs of correlations (15 for two-period and 15 for three-period), the average Spearman correlation is 0.60 and all of the correlation coefficients are significantly positive ($p < 0.001$ with one-sided test). These results suggest that the estimates of the exponential discounting parameter are driven by some inherent characteristic of data, and not by some regularity in the data that one particular model is not capturing.

Further evidence for cognitive discounting comes from comparing the models with and without discounting via likelihood ratio tests. For 4 out of 6 of the models in two-period problems (the exceptions being the tremble and generalized mean models) and all 6 models in three-period problems, the tests reject the model without discounting at the 5 percent level.

Result 8. *Estimates of structural models provide evidence for cognitive discounting: every model in the three-period decision problems is improved by incorporating cognitive discounting.*

[31] In the noisy self model, we also use discounted payoffs when calculating the utilities that enter the logistic function that determines beliefs. Our conclusions do not change if we instead use payoffs that are not discounted.

FIGURE 5.3.

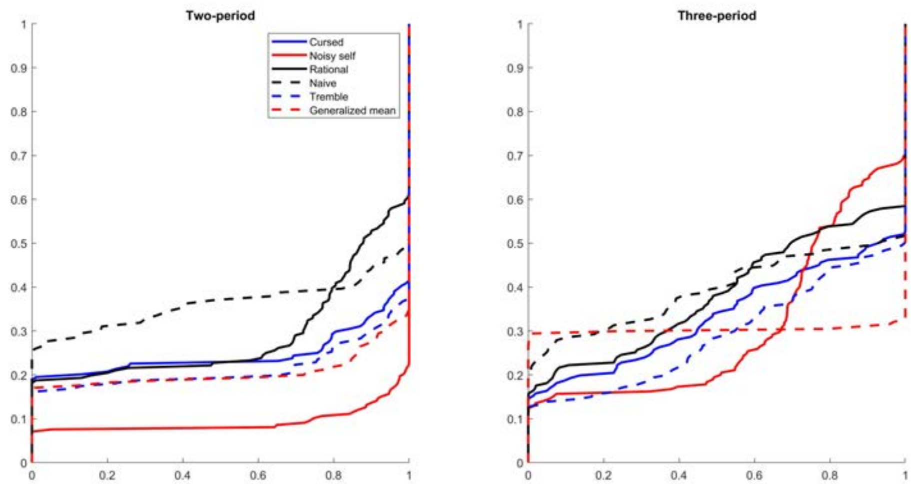

Notes: Empirical cdfs of the amount by which final period payoffs are cognitively discounted (δ in two-period problems and δ^2 in three-period problems) for each subject, estimated from the decision problems in the two-period treatment (left panel) and three-period treatment (right panel). Only estimates for subjects with $\lambda > 0.001$ are included, because δ is not identified for small values of λ.

5.6. **Noisy Self Model Results.** Given that the noisy self model with cognitive discounting best explains results, we investigate it in greater depth. Because the rationality parameters of the noisy self model are difficult to interpret, we calculate the probabilities with which each subject believes they will make a rational choice in the final period given the decision problem parameters and estimated model parameters. Figure 5.4 plots the empirical cdfs of the average (across decision problems) beliefs of rational choice by subject (restricting to subjects with $\lambda > 0.001$).

We see that beliefs about making a rational choice when R_F^- is realized are systematically smaller than when R_F^+ is realized, consistent with the difference in utilities being smaller in this case.[32] Furthermore, the beliefs are smaller in the three period decision problems than two period decision problems which is consistent with the lower rationality parameter estimates in these decision problems (see Figure 5.2. On average, subjects believe they will make the correct choice 75 percent of the time in two-period decision problems and 70 percent of the time in three-period decision problems. Importantly, except for a small fraction of subjects that the model cannot explain, these results imply that, on average, subjects believe they will make mistakes, but not that their future actions will be completely noisy.

[32]In fact, the beliefs about rational choice when R_F^- is realized are lowered by the fact that, when $R_F^- = R_F^O$, subjects are indifferent independent of their rationality parameter.

FIGURE 5.4. Noisy Self Beliefs

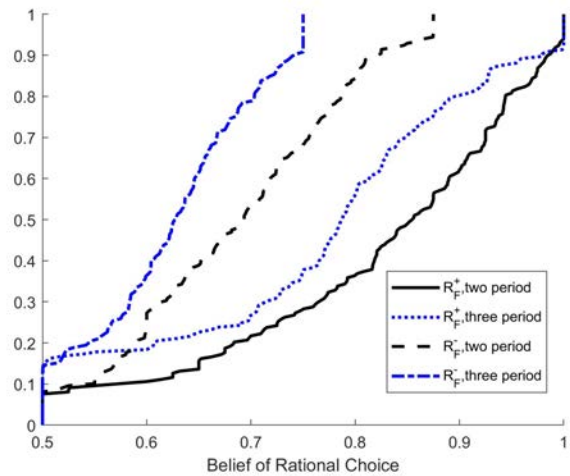

Notes: Empirical cdfs of the beliefs about the probability of rational choice in the final period. For each subject, beliefs are averaged across decision problems. Only estimates for subjects with $\lambda > 0.001$ are included.

Finally, we ask whether the explanatory power of the noisy self model with cognitive discounting is driven by heterogeneity, with some subjects being explained by cognitive discounting and others being explained by incorrect beliefs about future mistakes. Figure 5.5 plots the cognitive discounting parameter estimate versus the rationality parameter estimate for both the two-period (left panel) and three-period (right panel) decision problems. To ensure differences in the rationality parameter at low values are visible, we have capped the rationality parameters at 50.

From Figure 5.5, particularly in the three-period decision problems, we can clearly see a mass of subjects for which both mistaken beliefs are important (rationality parameter is small) and cognitive discounting is also important (parameter less than one). In fact, in these problems, 64 percent of subjects believe that they will make mistakes when R_F^+ is realized more than 5 percent of the time and also discount payoffs by more than 5 percent. Thus, even at the individual level, both mistaken beliefs and cognitive discounting appear to be important.[33]

6. DISCUSSION

We designed an experiment to capture rich behavioral data in which failures of backwards induction can be attributed to failures to correctly predict one's own future actions, isolating failures of backwards induction itself from other causes of departures from standard theory found in multi-player games (strategic ambiguity,

[33]We also looked for correlations between risk-aversion and rationality, and risk-aversion and cognitive discounting, but found no obvious relationships.

FIGURE 5.5. Noisy Self Parameters

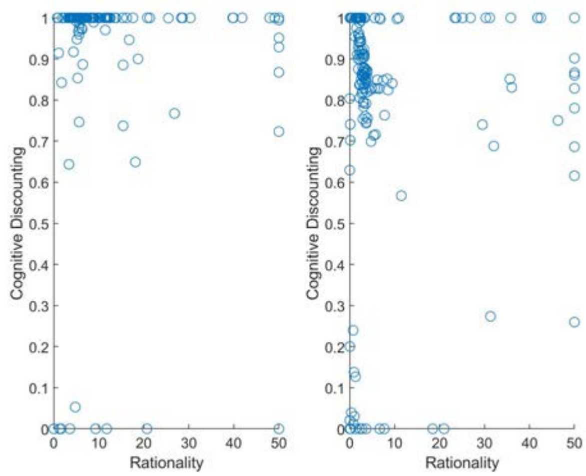

Notes: Parameter estimates from noisy self model with cognitive discounting. The rationality parameters have been capped at 50. Only estimates for subjects with $\lambda > 0.001$ are included.

other-regarding preferences, etc.). In perhaps the simplest possible setting requiring backwards induction (our two-period decision problems), we find extensive failures of backwards induction reasoning - subjects fail to predict their own future actions even after extensive experience taking those actions. In more complex, three-period problems, failures are also common, but, contrary to the hypothesis that subjects simply 'play it safe' in more complicated problems, we find more subtle patterns in the data.

From both reduced-form results and structural estimates, we find that subject behavior can best be described by two intertwined reasoning failures: (i) subjects imperfectly forecast their future actions, and (ii) they treat payoffs received in later period as if they are discounted (cognitive discounting). Exploring several models of incorrect foresight, we find that a model in which subjects believe that they will make mistakes less often when the difference in utilities is larger best fits behavior (the noisy self model).

We do not claim that the noisy self model with cognitive discounting is the best model possible model of behavior. But, we do believe that its two features are somewhat natural. The noisy self part of the model says that big differences between the payoffs of future actions makes it more likely subjects will recognize the correct action, forming correct beliefs. The cognitive discounting part of the model says that when future actions, and thus payoffs are more certain, subjects place more weight on payoffs that they can obtain immediately. It suggests that subjects search forward through the decision problem, rather than backwards as standard theory would have them do (something which Johnson et al. [2002] provide direct evidence for through MouseLab).

We hope theorists will take up the task of creating other models of behavior that can explain the data we have collected. One plausible means of unifying mistaken beliefs and cognitive discounting is through noisy perception. Gabaix and Laibson [2017] have shown that noisy perception induces a form of myopia. And, Frydman and Jin [2022] have shown that noisy perception can lead to dominated payoffs being chosen. If subjects, when thinking about their future plan perceive future payoffs noisily, it would then seem to lead to both a form of discounting and mistaken beliefs. Producing a tractable model of noisy foresight that can be applied in decision-making problems and games alike is a challenging, but worthy goal.

References

Ken Binmore, John McCarthy, Giovanni Ponti, Larry Samuelson, and Avner Shaked. A backward induction experiment. *Journal of Economic theory*, 104 (1):48–88, 2002.

Ciril Bosch-Rosa and Thomas Meissner. The one player guessing game: a diagnosis on the relationship between equilibrium play, beliefs, and best responses. *Experimental Economics*, 23(4):1129–1147, 2020.

Gary Charness and Dan Levin. The origin of the winner's curse: a laboratory study. *American Economic Journal: Microeconomics*, 1(1):207–36, 2009.

Benjamin Enke and Thomas Graeber. Cognitive uncertainty in intertemporal choice. Technical report, National Bureau of Economic Research, 2021.

Ignacio Esponda and Emanuel Vespa. Hypothetical thinking and information extraction in the laboratory. *American Economic Journal: Microeconomics*, 6(4): 180–202, 2014.

Ignacio Esponda and Emanuel Vespa. Contingent thinking and the sure-thing principle: Revisiting classic anomalies in the laboratory. *working paper*, 2019.

Erik Eyster and Matthew Rabin. Cursed equilibrium. *Econometrica*, 73(5):1623–1672, 2005.

Daniel E Fragiadakis, Daniel T Knoepfle, and Muriel Niederle. Who is strategic? *working paper*, 2016.

Daniel Friedman, Sameh Habib, Duncan James, and Brett Williams. Varieties of risk preference elicitation. *Games and Economic Behavior*, 133:58–76, 2022.

Cary Frydman and Lawrence Jin. Efficient coding and risky choice. *Quarterly Journal of Economics*, 137:161–213, 2022.

Xavier Gabaix and David Laibson. Myopia and discounting. *working paper*, 2017.

Xavier Gabaix and David I Laibson. A boundedly rational decision algorithm. *American Economic Review*, 90(2):433–438, 2000.

Xavier Gabaix, David Laibson, Guillermo Moloche, and Stephen Weinberg. Costly information acquisition: Experimental analysis of a boundedly rational model. *American Economic Review*, 96(4):1043–1068, 2006.

Paolo Ghirardato, Fabio Maccheroni, and Massimo Marinacci. Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory*, 118(2):133–173, 2004.

Uri Gneezy and Jan Potters. An experiment on risk taking and evaluation periods. *The Quarterly Journal of Economics*, 112(2):631–645, 1997.

Jacob K Goeree, Charles A Holt, and Thomas R Palfrey. Stochastic game theory for social science: a primer on quantal response equilibrium. In *Handbook of Experimental Game Theory*, pages 8–47. Edward Elgar Publishing, 2020.

Werner Güth and Reinhard Tietz. Ultimatum bargaining behavior: A survey and comparison of experimental results. *Journal of Economic Psychology*, 11(3):417–449, 1990.

Glenn W Harrison and E Elisabet Rutström. Risk aversion in the laboratory. In *Risk aversion in experiments*. Emerald Group Publishing Limited, 2008.

Yuval Heller. Three steps ahead. *Theoretical Economics*, 10(1):203–241, 2015.

Phillippe Jéheil. Limited horizon forecast in repeated alternate games. *Journal of Economic Theory*, 67(2):497–519, 1995.

Philippe Jehiel. Limited foresight may force cooperation. *The Review of Economic Studies*, 68(2):369–391, 2001.

Eric J Johnson, Colin Camerer, Sankar Sen, and Talia Rymon. Detecting failures of backward induction: Monitoring information search in sequential bargaining. *Journal of Economic Theory*, 104(1):16–47, 2002.

John H Kagel and Dan Levin. The winner's curse and public information in common value auctions. *The American Economic Review*, 76(5):894–920, 1986.

Shaowei Ke. Boundedly rational backward induction. *Theoretical Economics*, 14 (1):103–134, 2019.

Chad Kendall and Anujit Chakraborty. Future self-proof elicitation mechanisms. *working paper*, 2022.

Marco Mantovani et al. Limited backward induction: Foresight and behavior in sequential games. *Unpublished paper, Working Papers*, 289, 2014.

Alejandro Martínez-Marquina, Muriel Niederle, and Emanuel Vespa. Failures in contingent reasoning: The role of uncertainty. *American Economic Review*, 109 (10):3437–74, 2019.

Richard D McKelvey and Thomas R Palfrey. An experimental study of the centipede game. *Econometrica*, 60(4):803–836, 1992.

Richard D McKelvey and Thomas R Palfrey. Quantal response equilibria for normal form games. *Games and economic behavior*, 10(1):6–38, 1995.

Richard D McKelvey and Thomas R Palfrey. Quantal response equilibria for extensive form games. *Experimental economics*, 1(1):9–41, 1998.

Friederike Mengel. Learning by (limited) forward looking players. *Journal of Economic Behavior & Organization*, 108:59–77, 2014.

Jeevant Rampal. Opponent's foresight and optimal choices. Technical report, 2018.

Jeevant Rampal. Limited foresight equilibrium. *Games and Economic Behavior*, 132:166–188, 2022.

Appendix A. Additional Figures

Figure A.1. Three-Period Decision Problems Which Interchange p_1 and p_2

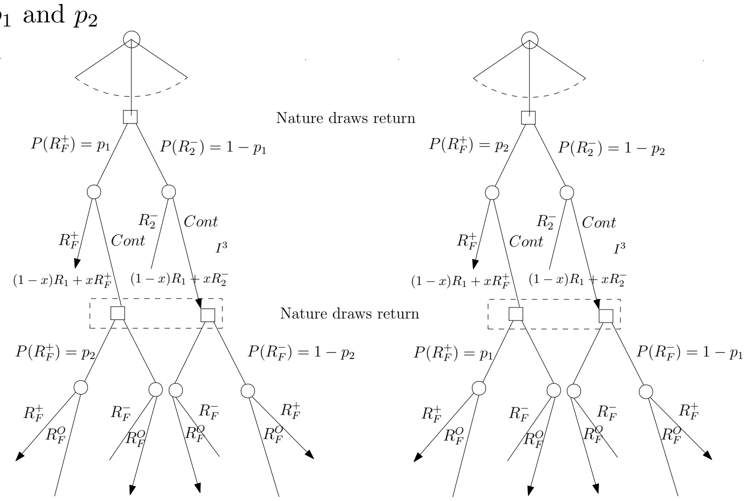

Notes: With $p_1 > p_2$,the decision problem on the left is back-loaded while that on the right is front-loaded. The parameters are from Table 1.

FIGURE A.2. Two-period Investments Compared to Lottery Investments

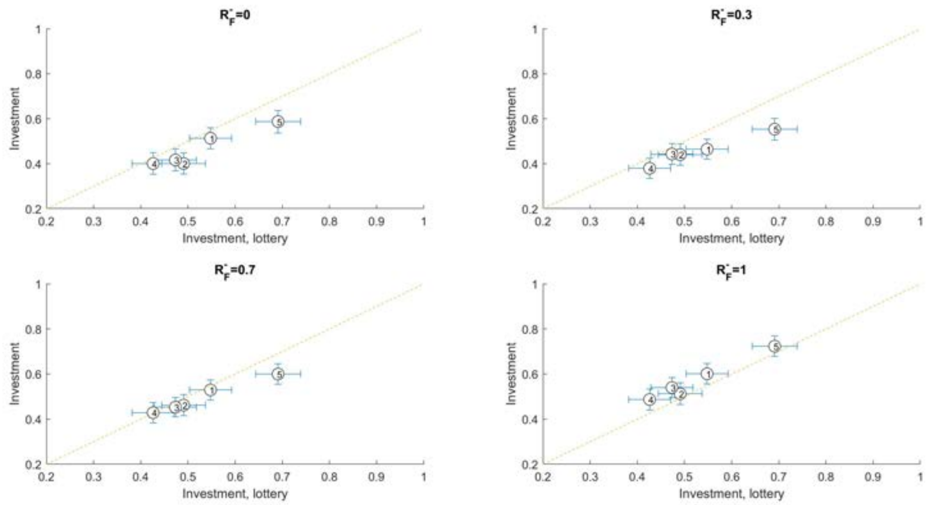

Notes: Each dot plots the investment in lotteries from groups 1-5, averaged across all subjects, on the x-axis, and for decision problems on the y-axis. The error bars indicate 95% confidence intervals in both dimensions.

FIGURE A.3. Two-period Investments (Second Half of Data)

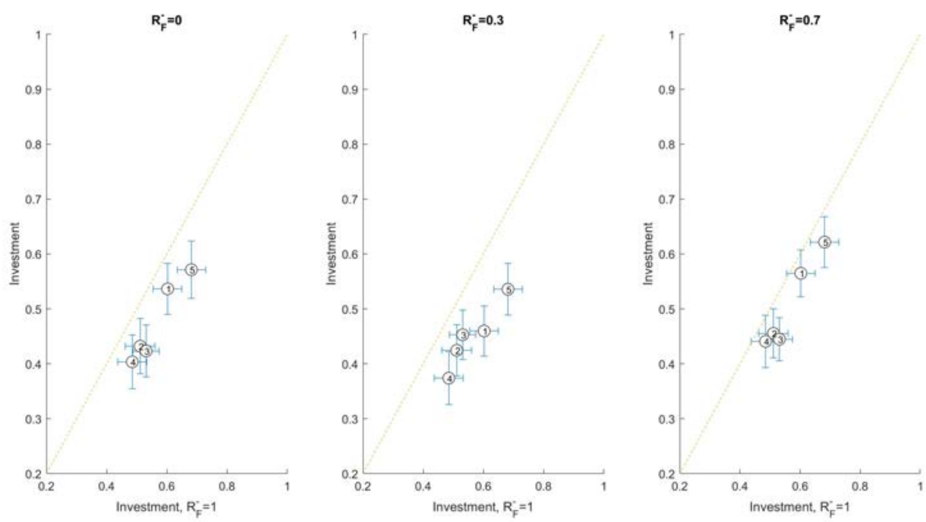

Notes: Each dot plots the investment in lotteries from groups 1-5, averaged across all subjects, on the x-axis, and for decision problems on the y-axis. Only the second half of the data is used - data for the first 10 problems is dropped. The error bars indicate 95% confidence intervals in both dimensions.

Figure A.4. Three-period Investments Compared to Lottery Investments

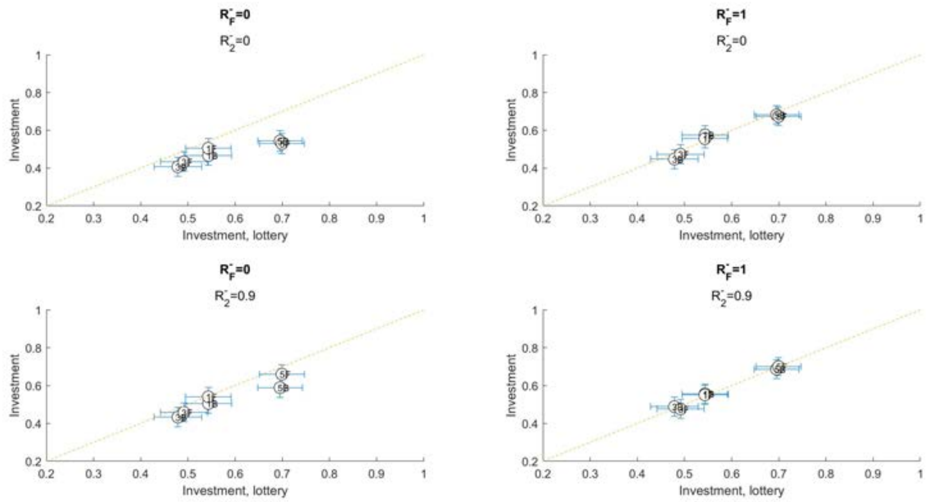

Notes: Each dot plots the investment in lotteries from groups 1, 3, and 5, averaged across subjects, on the x-axis, and for decision problems on the y-axis. For each decision problem, we plot the front-loaded (F) and back-loaded (B) versions of the problem separately. The left plot corresponds to $R_2^- = 0$ and the right to $R_2^- = 0.9$. The error bars indicate 95% confidence intervals in both dimensions.

Figure A.5. Three-period Investments (Second Half of Data)

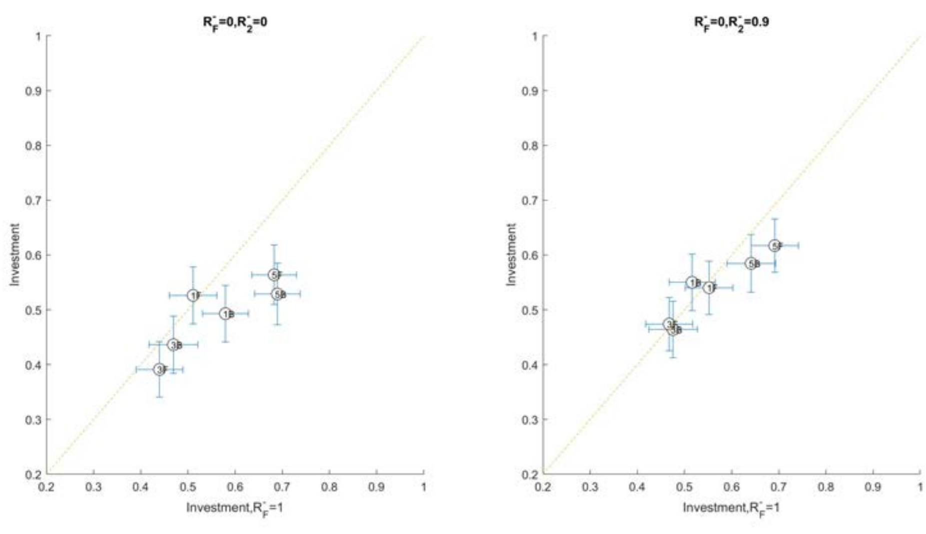

Notes: Each dot plots the investment in lotteries from groups 1-5, averaged across all subjects, on the x-axis, and for decision problems on the y-axis. Only the second half of the data is used - data for the first 13 problems is dropped. The error bars indicate 95% confidence intervals in both dimensions.

Appendix B. Theoretical Results

Proposition 1. *Assume a unique interior optimal solution, x^*, exists to the general model in (5.1) with $u' > 0, u'' < 0, q_1 > 0, q_1' = 0$, and $q_2' = 0$. The following statements hold:*
i) x^* is strictly decreasing in R_1. x^* is strictly increasing in R_F^+ and p.
ii) x^* is strictly increasing in R_F^O for any q_1, q_2 such that $\max\{q_1, q_2\} < 1$, including for the rational case of $q_1 = 1, q_2 = 0$.
iii) x^* is strictly decreasing in R_F^- if and only if $q_2 > 0$.

Proof. Proof of Proposition 1: Assuming an interior solution, the first order condition is

$$pq_1(R_F^+ - R_1)u'\left(x^*R_F^+ + (1-x^*)R_1\right) + (1-p)q_2(R_F^- - R_1)u'\left(x^*R_F^- + (1-x^*)R_1\right)$$
$$+[p(1-q_1) + (1-p)(1-q_2)](R_F^O - R_1)u'\left(x^*R_F^O + (1-x^*)R_1\right) = 0$$

Next, we derive the comparative static with respect to R_F^+. Differentiating implicitly with respect to R_F^+, we get

$$\underbrace{pq_1u'\left(x^*R_F^+ + (1-x^*)R_1\right)} + pq_1(R_F^+ - R_1)^2u''\left(x^*R_F^+ + (1-x^*)R_1\right)\frac{dx^*}{dR_F^+}$$
$$+(1-p)q_2(R_F^- - R_1)^2u''\left(x^*R_F^- + (1-x^*)R_1\right)\frac{dx^*}{dR_F^+}$$
$$+[p(1-q_1) + (1-p)(1-q_2)](R_F^O - R_1)^2u''\left(x^*R_F^O + (1-x^*)R_1\right)\frac{dx^*}{dR_F^+} = 0$$

As $u'' < 0$, we have that $\frac{dx^*}{dR_F^+} > 0$ if and only if the underbraced term is strictly positive which happens if and only if $q_1 > 0$ as assumed. Similarly, one can show that $\frac{dx^*}{dR_F^-}, \frac{dx^*}{dR_F^O} \geq 0$ and $\frac{dx^*}{dR_1} \leq 0$ under the relevant conditions.
We also derive the comparative static with respect to p, as it requires one extra step. Differentiating the first order condition implicitly with respect to p, we get

$$q_1(R_F^+ - R_1)u'\left(x^*R_F^+ + (1-x^*)R_1\right) + pq_1(R_F^+ - R_1)^2u''\left(x^*R_F^+ + (1-x^*)R_1\right)\frac{dx^*}{dp}$$
$$\underbrace{-q_2(R_F^- - R_1)u'\left(x^*R_F^- + (1-x^*)R_1\right)} + (1-p)q_2(R_F^- - R_1)^2u''\left(x^*R_F^- + (1-x^*)R_1\right)\frac{dx^*}{dp}$$
$$+\underbrace{(q_2 - q_1)(R_F^O - R_1)u'\left(x^*R_F^O + (1-x^*)R_1\right)}$$
$$+[p(1-q_1) + (1-p)(1-q_2)](R_F^O - R_1)^2u''\left(x^*R_F^O + (1-x^*)R_1\right)\frac{dx^*}{dp} = 0$$

The sum of the underbraced terms are positive, as $q_2(R_1 - R_F^-) \geq (q_2 - q_1)(R_1 - R_F^O)$, and u' is a decreasing function. Thus, $\frac{dx^*}{dp}$ must be positive. □

Appendix C. Instructions

The instructions for the two-period treatment are reproduced on the following pages.

Please read the following instructions on the following page carefully. You will be asked several questions about them before the experiment begins.

Instructions

Instructions

In this experiment, you will complete investment tasks, divided into two parts. We will explain the tasks for the first part here, and give you new instructions for the second part after the first part is done. One of the tasks you complete will be randomly chosen to determine your bonus (all are equally likely to be chosen).

Instructions for Part 1

In each task, you will start with $1.00 in an Investment. You will have two opportunities to withdraw your money from the Investment. Your bonus will be the **sum** of your earnings from the Investment and the amounts you withdraw.

For each task, you will see a table like the one below.

Withdrawal Multiplier (first step)	1
Possible Withdrawal Multipliers (second step)	3 ; 3 ; 3 ; 3 ; 3 0 ; 0 ; 0 ; 0 ; 0

You will make your withdrawal decisions in two steps:

- In the first step, you can decide to withdraw any amount between $0 and $1.00 from the Investment. As part of your bonus, you will earn the amount you withdraw multiplied by the Withdrawal Multiplier for the first step shown in the table. You will choose the amount to withdraw using a slider that will appear below the table.

- In the second step:

  1. The computer will roll a (virtual) 10-sided die like the one shown above. Each of the 10 sides of the die corresponds to one of the Withdrawal Multipliers in the table. So, in the

example above, numbers 1-5 would result in a Multiplier of 3, and numbers 6-10 would result in a Multiplier of 0.

2. **You will first be shown the chosen Multiplier**, and then decide whether to withdraw whatever money is left in the investment, or not. If you choose to withdraw, you will earn the amount you have remaining in the Investment multiplied by the Multiplier as part of your bonus. If you choose to continue investing, you will instead **earn the amount you have remaining in the investment (a Multiplier of 1)** as part of your bonus.

For example, if you withdraw $0.50 in the first step, you will earn 1 x $0.50=$0.50 in the first step. In the second step, you will have $0.50 remaining. If you choose not to withdraw, you will earn another $0.50. If you learn the Multiplier is 3 and withdraw, you will earn 3 x $0.50=$1.50. If you learn the Multiplier is 0 and withdraw, you will earn 0 x $0.50=$0. Your total earnings will be the sum of the amount you earn in each step.

When you have understood these instructions, please proceed to the next page to answer several comprehension questions.

Comprehension Questions

Please answer the following comprehension questions. Click 'Show/hide instructions' if you need to review the instructions again.

Show/hide instructions

1. Please indicate which of the following statements is true:

- your total bonus comes from the sum of your earnings in each step
- your bonus comes from your earnings in the first step only
- your bonus comes from your earnings in the second step only

2. True or False: In the second step, you learn the Multiplier for withdrawals **before** deciding whether or not to withdraw the remainder of your money from the Investment.

- True
- False

3. Suppose you have $0.75 remaining in the investment in the second step. If you learn that the Multiplier for withdrawals is 2, how much will you earn from the second step if you **withdraw** your money?

- $0
- $0.75
- $1.50
- it is random

4. Suppose you have $0.75 remaining in the investment in the second step. If you learn that the Multiplier for withdrawals is 2, how much will you earn from the second step if you **do not withdraw** your money?

- $0
- $0.75
- $1.50
- it is random

You have answered all of the comprehension questions correctly. Please continue to begin the tasks for Part 1.

Instructions part 2

Instructions for Part 2

The investment tasks in Part 2 are very similar to those in Part 1 except that **you will only make ONE withdrawal decision**. In each task, you will start with $1.00 in an Investment. You will have one opportunity to withdraw your money from the Investment. Your bonus will be the **sum** of your earnings from the Investment and the amount you withdraw.

For each task, you will see a table like the one below.

Withdrawal Multiplier	1
Possible Investment Multipliers	3 ; 3 ; 3 ; 3 ; 3 0 ; 0 ; 0 ; 0 ; 0

In each task, you will choose how much to withdraw. You will earn the amount you withdraw multiplied by the Withdrawal Multiplier shown in the table. You will choose the amount to withdraw using a slider that will appear below the table.

The amount you do not withdraw remains invested in the Investment. The computer will roll a (virtual) 10-sided die to determine the Investment Multiplier from one of the 10 possibilities shown in the table. You will earn the amount you have left in the Investment multiplier by this Multiplier.

For example, if you withdraw $0.50, you will earn 1 x $0.50=$0.50. If the Investment Multiplier is 3, you will earn an additional 3 x $0.50=$1.50. If it is 0, you will earn an additional 0 x $0.50=$0.

When you have understood these instructions, please proceed to the next page to answer a comprehension question.

1. True or False: You can decide whether to continue investing or not after you learn the Investment Multiplier.

Show/hide instructions

- True
- False

You have answered all of the comprehension questions correctly. Please continue to begin the tasks for Part 2.

The instructions for the three-period treatment are reproduced on the following pages.

Please read the following instructions on the following page carefully. You will be asked several questions about them before the experiment begins.

Instructions

Instructions

In this experiment, you will complete investment tasks, divided into two parts. We will explain the tasks for the first part here, and give you new instructions for the second part after the first part is done. One of the tasks you complete will be randomly chosen to determine your bonus (all are equally likely to be chosen).

Instructions for Part 1

In each task, you will start with $1.00 in an Investment. You will have three opportunities to withdraw your money from the Investment. Your bonus will be the **sum** of your earnings from the Investment and the amounts you withdraw.

For each task, you will see a table like the one below.

Withdrawal Multiplier (first step)	1
Possible Withdrawal Multipliers (second step)	50 balls with 1 50 balls with 0
Possible Withdrawal Multipliers (third step)	50 balls with 3 50 balls with 0

You will make your withdrawal decisions in three steps:

- In the first step, you can decide to withdraw any amount between $0 and $1.00 from the Investment. You will earn the amount you withdraw multiplied by the Withdrawal multiplier for the first step as part of your bonus. You will choose the amount to withdraw using a slider that will appear below the table.

- In the second step:
  1. The computer will draw a ball from a (virtual) bag containing 100 balls like the one shown above. As shown in the table, the balls correspond to the possible Withdrawal Multipliers. So, in the example above, 50 of the balls result in a Multiplier of 1, and the other 50 result in a Multiplier of 0.
  2. **You will first be shown the chosen Multiplier**, and then decide whether to withdraw whatever money is left in the investment, or not. If you choose to withdraw, you will earn the amount you have remaining in the Investment multiplied by the Multiplier as part of your bonus. If you choose to continue investing, you will go to the third step.
- The third step is identical to the second step except that if you do not withdraw (after seeing the Multiplier that the computer randomly chose for the third step), you will **earn the amount you have remaining in the investment (a Multiplier of 1)** as part of your bonus. The chance that each Multiplier is chosen in the third step is unrelated to the Multiplier chosen in the second step.

For example:

- If you withdraw $0.50 in the first step, you will earn 1 x $0.50=$0.50 in the first step. In the second step, you will have $0.50 remaining.
- In the second step, if you learn the Multiplier is 1 and withdraw, you will earn 1 x $0.50=$0.50 and the task ends. If you learn the Multiplier is 0 and withdraw, you will earn 0 x $0.50=$0 and the task ends. If you do not withdraw, we go to the third step.
- In the third step, if you learn the Multiplier is 3 and withdraw, you will earn 3 x $0.50=$1.50. If you learn the Multiplier is 0 and withdraw, you will earn 0 x $0.50=$0. If you do not withdraw, you will earn $0.50.

When you have understood these instructions, please proceed to the next page to answer several comprehension questions.

Comprehension Questions

Please answer the following comprehension questions. Click 'Show/hide instructions' if you need to review the instructions again.

Show/hide instructions

1. Please indicate which of the following statements is true:

- your total bonus comes from the sum of your earnings in each step
- your bonus comes from your earnings in the first step only
- your bonus comes from your earnings in the second step only
- your bonus comes from your earnings in the third step only

2. True or False: In the second and third steps, you learn the Multiplier for withdrawals **before** deciding whether or not to withdraw the remainder of your money from the Investment.

- True
- False

3. Suppose you have $0.75 remaining in the investment in the third step. If you learn that the Multiplier for withdrawals is 2, how much will you earn from the third step if you **withdraw** your money?

- $0
- $0.75
- $1.50
- it is random

4. Suppose you have $0.75 remaining in the investment in the third step. If you learn that the Multiplier for withdrawals is 2, how much will you earn from the third step if you **do not withdraw** your money?

- $0
- $0.75
- $1.50
- it is random

Instructions2

Instructions for Part 2

The investment tasks in Part 2 are very similar to those in Part 1 except that **you will only make ONE withdrawal decision in the first period, before you learn the actual Multipliers**. In each task, you will start with $1.00 in an Investment. You will have one opportunity to withdraw your money from the Investment. Your bonus will be the **sum** of your earnings from the Investment and the amount you withdraw.

For each task, you will see a table like the one below.

Withdrawal Multiplier	1
Possible Investment Multipliers (first draw)	50 balls with 1 50 balls with DRAW AGAIN
Possible Investment Multipliers (second draw)	50 balls with 3 50 balls with 0

In each task, you will first choose how much to withdraw. You will earn the amount you withdraw multiplied by the Withdrawal Multiplier shown in the table. You will choose the amount to withdraw using a slider that will appear below the table.

The amount you do not withdraw remains invested in the Investment. The computer will draw a ball from a (virtual) bag with 100 balls. Some of the balls are labeled with an Investment Multiplier and some are labeled with DRAW AGAIN, as shown in the table (first draw). If a ball with an Investment Multiplier is drawn, you will earn the amount you have left in the Investment multiplier by this Multiplier. If the ball that is drawn says DRAW AGAIN, the computer will draw a ball from a second bag with 100 balls. Each of these balls is labeled with the Investment Multiplier as shown in the table (second draw). You will earn the amount you have left in the Investment multiplier by this Multiplier.

For example, if you withdraw $0.50, you will earn 1 x $0.50=$0.50. If the computer first draws a ball with an Investment Multiplier of 1, you will earn an additional 1 x $0.50=$0.50. If it draws a ball labeled DRAW AGAIN, it draws a ball from the second bag. If the ball is 3, you will earn an additional 3 x $0.50=$1.50. If it is 0, you will earn an additional 0 x $0.50=$0.

When you have understood these instructions, please proceed to the next page to answer a comprehension question.

1. True or False: You can decide whether to continue investing or not after you learn the Investment Multiplier.

Show/hide instructions

- True
- False

You have answered all of the comprehension questions correctly. Please continue to begin the tasks for Part 2.