

An Experimental Analysis of the Prize–Probability Tradeoff in Stopping Problems*

Yair Antler[†] and Ayala Arad[‡]

Abstract

We experimentally examine how individuals commit to a cutoff stopping rule when facing a sequence of independent lotteries. We identify two main behavior patterns: (1) a small share of participants consistently choose stopping rules whose gain bound (i.e., the accumulated gain at which the sequence stops) is larger than the loss bound, and (2) a larger share of participants consistently choose rules whose loss bound is larger than the gain bound. We introduce a procedural decision-making model that accounts for these patterns and show that the behavior of most of our participants is inconsistent with prominent theories of decision-making under risk.

Keywords: commitment, compound lotteries, decision procedure, experiment, stopping problems, type classification.

JEL Codes: C91, D01, D81, D90

* We thank Aluma Dembo, Dudi Dilenberger, Kfir Eliaz, and Ariel Rubinstein for their comments, and Tomer Baratz, Shira Bukchin, Sean Joffe, Eli Mograbi, and Hadar Weisman for research assistance. We acknowledge support from the Collier Foundation, the European Union’s Horizon 2020 Research and Innovation programme (Grant Agreement No. 952574), the Henry Crown Institute of Business Research in Israel, as well as the International Visiting Fellowship at the University of Essex.

[†] Collier School of Management, Tel Aviv University. yair.an@gmail.com

[‡] Collier School of Management, Tel Aviv University. ayala.arad@gmail.com

1. Introduction

Stopping problems appear in numerous contexts in economics and finance, ranging from option pricing and job search to experimentation, technology adoption, and gambling. In these problems, an individual observes a sequence of realizations of a stochastic process and has to decide when to stop it. According to several prominent theories of decision-making under risk (e.g., expected utility), an optimal stopping plan can be described by a simple *cutoff* rule, namely, stopping the process once an individual's payoff reaches a threshold.

Our main research objective is to experimentally examine how individuals make *binding* stopping plans and what forces shape these plans. Understanding to which plans individuals commit is important not only because such commitment is relevant in practice but also because it reveals individuals' preferences over the induced outcomes of dynamic play in situations where there is no commitment (e.g., casino gambling). To see this, note first that commitment to a cutoff rule turns a dynamic stopping problem into a *static* one, which may simplify it and help individuals better understand certain aspects of the problem. Second, such commitment enables individuals to choose their preferred stopping plan without worrying about their ability to implement it. Finally, from the researcher's perspective, observing an individual's binding stopping plan enables learning about her preferences without the data being contaminated by biases and inconsistencies that may arise during a dynamic play.¹

In order to understand our setting, consider a decision-maker (DM) who faces an infinite sequence of lotteries, where each lottery pays 1 with probability p and -1 with probability $1 - p$. Under various theories of decision-making under risk, the DM's optimal stopping plan can be described by an upper bound $h > 0$ and a lower bound $l \leq 0$ such that the DM stops the process once her payoff hits one of these bounds. The higher h is, the less likely the process is to reach h before it reaches l ; the lower l is, the less likely the process is to reach l before it reaches h . Thus, when choosing these bounds, the DM trades off between two aspects: the probability of winning and the size of the potential gain/loss.² This tradeoff is at the heart of our experimental design.

As an illustration, consider the two cutoff rules given in Figure 1. Under both rules, the sequence stops once the DM accumulates a net loss of 20. Under a (resp., b), the sequence stops once the DM accumulates a gain of 10 (resp., 30). We refer to cutoff rules for which the upper bound is smaller (resp., larger) in absolute value than the lower bound as *left-biased* (resp., *right-*

¹ For example, the negative feelings associated with realizing losses may lead investors to hold on to badly performing stocks (Shefrin and Statman, 1985).

² When $p \neq 0.5$, the probability of stopping the process at a gain is $\frac{1-q^{-l}}{1-q^{h-l}}$, where $q = \frac{1-p}{p}$ (Feller, 1970).

biased). The likelihood that the sequence ends with a loss is smaller under the left-biased rule *a*, while the potential gain is greater under the right-biased rule *b*. Thus, when the DM chooses between the two rules, she trades off between the potential gain and the probability of a gain.

Rule	Lower bound	Upper bound
<i>a</i>	-20	+10
<i>b</i>	-20	+30

Figure 1. Two cutoff rules with the same lower bound.

To gain intuition, consider a risk-neutral expected utility maximizer who has to choose between stopping rules with a fixed lower bound, as in Figure 1. When the baseline lottery is unfavorable (i.e., $p < 0.5$), she will obtain a higher expected utility under the left-biased rule. To see this, note that the left-biased rule induces a smaller expected number of negative expected value lotteries: the two rules induce the same number of lotteries if the process reaches -20 before reaching $+10$ and rule *b* results in a larger number of lotteries otherwise.³ In a symmetric manner, when the baseline lottery is favorable (i.e., $p > 0.5$), she will choose the right-biased rule *b*. Now consider the case of a risk-neutral expected utility maximizer who has to choose between two cutoff rules that share the same upper bound, as in Figure 2. She will prefer the left-biased rule *d* if $p > 0.5$ and the right-biased rule *c* if $p < 0.5$ as she would like to maximize the expected number of baseline lotteries in the former case and minimize it in the latter case. We can conclude that expected value maximization can lead to a choice of left-biased rules or right-biased rules, depending on the context.

Rule	Lower bound	Upper bound
<i>c</i>	-10	+20
<i>d</i>	-30	+20

Figure 2. Two cutoff rules with the same upper bound.

³ Formally, when $p \neq 0.5$ the expected number of lotteries played given a lower bound of l and an upper bound of h is $\frac{-l}{1-2p} - \frac{-l+h}{1-2p} * \frac{1-q^{-l}}{1-q^{h-l}}$, where $q = \frac{1-p}{p}$ (Feller, 1970).

In our experiment, each participant faced 36 choice problems in this spirit. In each problem, the participants had to choose one rule out of five rules: two right-biased ones, two left-biased ones, and a symmetric rule. The problems varied in the probability of winning in the baseline lottery ($p < 0.5$ in the first part of the experiment and $p > 0.5$ in the second part), and whether the upper bound, the lower bound, or neither of them was fixed within a problem. We ran two treatments: in our main treatment, T_0 , the stopping rules' induced winning probabilities were not provided to the participants, whereas in the second treatment, T_p , they were provided. We now focus on T_0 and later discuss the findings in T_p and their implications.

Our main finding is a general tendency to either consistently choose left-biased rules or consistently choose right-biased ones, across qualitatively and quantitatively different choice problems. We find a larger share of participants who consistently choose left-biased rules than of participants who consistently choose right-biased ones. The participants' choices suggest that many of them categorize the rules into right- and left-biased rules: they tend to choose either a left- or a right-biased rule, but not necessarily the most biased rule in the respective direction. This apparent categorization is in the spirit of the *binary bias*, which suggests that people tend to categorize items into two *main* distinct categories, for example, positive and negative reviews (for recent documentation of this bias in the psychology literature see Fisher and Keil, 2018, and Fisher et al., 2018). This binary categorization seems natural in our setting due to the simplicity of splitting the set of rules according to their directional bias (i.e., right-biased or left-biased). Moreover, as expressed in the participants' written explanations of their choices, each category reflects prioritizing one of the two main aspects of the problem: choosing right-biased rules reflects prioritizing high prizes while choosing left-biased rules reflects prioritizing the probability of winning.

The participants' behavior, together with their explanations, suggests that most of them try to solve a simple tradeoff between the likelihood of winning or losing and the size of the prizes. The particular way in which this tradeoff is solved depends on the favorability of the baseline lottery. Indeed, choices of left-biased rules are more common in problems in which the baseline lottery is unfavorable, whereas choices of right-biased rules are more common in problems in which the baseline lottery is favorable. A potential explanation for this pattern is that when the baseline lottery becomes favorable, participants feel that they are more likely to finish with a gain and hence they shift their attention from the probability of not losing to the size of the potential gain.

Although the participants' choices differ between the two parts of the experiment (i.e., favorable vs. unfavorable baseline lotteries), they are strongly correlated. In fact, for most participants in our main treatment, the solution of the prize–probability tradeoff is virtually unaffected by the specific details of each problem (e.g., the exact probabilities and expected value of the stopping rules' induced lotteries). In Section 5, we show that participants have a qualitative understanding of the prize–probability tradeoff that arises in stopping problems and suggest that this understanding makes them reason in qualitative terms, focusing on this tradeoff rather than considering the fine details of the decision problem.

In light of the above findings, we treat left-biased rules and right-biased rules as two distinct categories, and study the participants' behavior within each category. We find that there are three groups of participants of (roughly) equal size: participants who tend to consistently choose the extreme stopping rule within each category (i.e., the most right-biased rule or the most left-biased rule), participants who tend to consistently choose the moderate rule within a category (i.e., the second-most right- or left-biased rule), and participants who diversify with different intensities between extreme and moderate rules across problems.

To account for the experimental results, we suggest a decision procedure according to which individuals operate in two stages. They begin with “the big picture”: resolving the fundamental prize–probability tradeoff between the two categories of left-biased and right-biased rules. They then continue to the “finer details”: resolving the more incremental prize–probability tradeoff within a category between extreme and moderate rules. We formalize the procedure using a simple qualitative model that consists of two key parameters (and a noise term): one parameter that captures a participant's tendency to choose left- or right-biased rules, and one parameter that captures their tendency to choose extreme or moderate rules within a category. We refer to this procedure as the *two-stage qualitative tradeoff resolution (2S-QTR)* model.

In Section 4, we examine the extent to which the 2S-QTR model can explain our participants' behavior. To this end, we performed a *leave-one-out prediction exercise* for each participant separately: we estimated the model using 35 problems and used the estimate to predict her behavior in the remaining problem. A participant's behavior is considered consistent with the 2S-QTR model if the number of correct predictions across the 36 problems is sufficiently large (the specific threshold was set such that the probability of classifying as consistent a participant who chooses at random is less than 1%). In our main treatment, roughly 75% of the participants exhibited behavior consistent with the model. We then estimated the model for these participants using the 36 problems and classified them as types based on their tendency to choose

left-biased rules, right-biased rules, moderate rules, or extreme rules. The most common tendencies are toward left-biased rules and extreme rules.

Our experimental design enabled us to test whether the above findings can be explained by standard “off-the-shelf” theories of decision-making under risk. We examined, *for each participant*, whether her behavior fits the predictions of several prominent decision theories: (constant relative) risk aversion, cumulative prospect theory (Kahneman and Tversky, 1992), disappointment aversion (Gul, 1991), regret aversion (Bell, 1982; Loomes and Sugden, 1982), and salience theory (Bordalo et al., 2012). To do so, we ran a leave-one-out prediction competition between all of these theories and the 2S-QTR model. We classified a participant into a theory if (i) the theory was able to predict a sufficiently large number of the participant’s choices (the threshold was identical to the one chosen for the 2S-QTR model) and (ii) no other theory was able to predict a larger number of choices. In this more conservative exercise, the 2S-QTR model accounts for the behavior of 69% of the participants in the main treatment. Prospect theory accounts for the behavior of 33% of the participants. None of the other theories we examined accounts for the behavior of more than 3% of the participants.

Studying our second treatment, T_p , in which the rules’ induced probabilities are provided, sheds light on the extent to which the choice patterns observed in T_0 are due to the participants’ lack of knowledge of these induced probabilities. While the literature on stopping problems is not insubstantial, to the best of our knowledge the difference between these two conditions is underexplored. In T_p , the 2S-QTR model accounts for the behavior of 74% of the participants (51% in the prediction competition exercise). Here again, the most common tendencies are toward left-biased rules and extreme rules. The tendency toward extreme rules is significantly greater in T_p than in T_0 . We suggest that the knowledge of the rules’ induced probabilities makes participants in T_p more sensitive to the fine details than the T_0 participants, which leads to different choices in some of the problems. In particular, they are better able to recognize situations in which there is a clear-cut way of resolving the tradeoff and choose accordingly. For example, in situations where a minor deduction of a winning probability leads to a major increase in prizes, they tend to opt for the extreme right-biased rule. These findings suggest that knowing the induced probabilities changes individuals’ behavior by making them consider the finer details of the problem. However, it does not change the way they perceive the big picture, namely, their directional bias.

1.1 Related literature

The present paper is related to a recent strand of the literature that investigates planning in dynamic decision-making under risk. Fischbacher et al. (2017) show that stop-loss and take-gain strategies mitigate the disposition effect in dynamic play. Dertwinkel-Kalt et al. (2020) conduct a lab experiment in which they find that plans and dynamic behavior in a stopping problem are consistent with the predictions of Bordalo et al.'s (2012) salience theory. Alaoui and Fons-Rosen (2021) find that grittier individuals have a higher tendency to over-gamble relative to their original plans. Perhaps closest to our paper is Heimer et al. (2021) who document a discrepancy between investors' initial plans and their actual behavior: most investors choose stopping rules that are right-biased (the modal strategy of 46% of the investors is right-biased while the modal strategy of 32% is left-biased) *ex ante*, but their subsequent choices follow the reverse pattern. To pin down the mechanism behind this discrepancy, they perform an online experiment in which individuals stop a finite sequence of fair binary lotteries. In particular, one of their treatments, dubbed "hard plan," examines how individuals commit to a stopping rule in this situation.

While Heimer et al.'s (2021) elegant design allows them to pin down the mechanism underlying the discrepancy between planning and playing, we focus on *individual* decision-making with commitment, which requires a richer dataset at the individual level. To this end, we recorded 36 choices of a stopping rule (in different contexts) per individual rather than the single choice recorded in their hard plan treatment. In addition to the different focus, there are several differences between the hard plan treatment and our setting, which may explain the differences in the tendency to choose right-biased rules. Perhaps the most significant difference is that the baseline lotteries in Heimer et al. have an expected value of zero while ours have either a strictly positive or a strictly negative expected value.⁴ Observe that, *given a stopping rule*, an "almost fair" baseline lottery such as the ones used in our experiment is likely to induce winning probabilities that are very far from fair, which means that the decision problems the participants faced in the two experiments, ours and Heimer et al.'s, are quite different. For example, the rule $(-15, +15)$ induces a fair lottery if the baseline lottery is fair, but it induces a probability of winning of 30.7% when $p = 18/37$, as in the first part of our experiment.

⁴ Additional noteworthy differences are that (1) Heimer et al.'s participants choose a stopping rule freely while in our setting, to better understand the tradeoffs that participants make and to distinguish between the predictions of prominent theories, we let our participants choose from various fixed sets of five rules, and (2) the participants in our experiment were STEM and management students, who are presumably more familiar with basic statistics and may have a better understanding of the implications of different stopping rules compared to the typical online subject pool.

Other papers study stopping decisions without planning. In Strack and Viefers (2021), the participants choose when to stop a multiplicative random walk and exhibit history-dependent behavior, which is consistent with regret aversion and inconsistent with cutoff rules.⁵ Sandri et al. (2010) examine exit decisions and find that most individuals tend to hold on to a badly performing asset longer than is consistent with real option reasoning.

Stopping plans have been studied indirectly in the experimental literature on dynamic inconsistency, which focuses on *deviations* from planning when individuals face a small number of lotteries. Barkan and Busmeyer (1999, 2003) and Ploner (2017) find evidence of dynamically inconsistent behavior in settings where individuals decide whether to participate in an additional lottery after experiencing one outcome. Cubitt and Sugden (2001) do not reject the dynamic consistency hypothesis when participants have to decide how many all-or-nothing additional gambles to participate in after winning in four mandatory rounds.

Our work also relates to the literature on skewness-seeking and prudent behavior. Skewness corresponds to our notion of left/right-biased stopping rules. The more right-biased a rule is, the greater is the skewness of its induced lottery.⁶ Golec and Tamarkin (1998) find evidence of skewness-seeking behavior in horse-race betting. Brunner et al. (2011), Deck and Schlesinger (2010, 2014), Ebert and Wiesen (2011, 2014), Ebert (2015), Grossman and Eckel (2015), Maier and Ruger (2012), and Noussair et al. (2014) provide evidence for skewness-seeking and/or prudent behavior in lab experiments. Bleichrodt and van Bruggen (2018) find prudent behavior in the gain domain and imprudent behavior in the loss domain.

There are several differences between our setting and the typical setting in this strand of the literature. The experiments on skewness-seeking and prudent behavior typically examine choices between lotteries with identical means and variance. By contrast, the stopping rules in our setting induce compound lotteries with different means and variance such that prudence does not imply a tendency to choose right-biased rules (e.g., facing the two rules in Figure 1, a prudent individual may choose the left-biased rule when $p < 0.5$ as it induces a greater expected value and a smaller variance than the right-biased rule). Moreover, the stopping rules' framing is different from the standard lottery framing, even when participants are provided with the rules' induced probabilities (as in our second treatment). The dynamic story underlying stopping

⁵ This type of behavior is also consistent with other theories of decision-making under risk such as cautious stochastic choice (Henderson et al., 2022).

⁶ It should be noted that a left-biased rule can induce a positively skewed lottery when $p < 0.5$ and a right-biased rule can induce a negatively skewed lottery when $p > 0.5$.

problems and the participants' qualitative understanding of the prize–probability tradeoff in this context may encourage reasoning in qualitative terms, which is less likely to be triggered when choosing between standard binary lotteries.

Recent theoretical work by Ebert and Karehnke (2021) characterizes the skewness preferences implied by a large number of theories of decision-making under risk. In particular, they find that prudent expected utility, disappointment aversion, regret aversion, and salience theory imply skewness-seeking (of different orders), and that cumulative prospect theory with the conventional S-shaped value function can imply both skewness-seeking and skewness aversion, depending on the parameters. These findings suggest that, relative to other prominent theories of decision-making under risk, cumulative prospect theory has the greatest potential to explain the behavior of a large share of our participants. The econometric estimation performed in Section 4 confirms this intuition. Yet, it indicates that the behavior of the majority of the participants is inconsistent with cumulative prospect theory.

Finally, our finding that many participants use qualitative decision rules contributes to the behavioral literature that aims to identify and model individuals' decision-making procedures instead of assuming that choices are guided by some utility maximization (see, for example, Güth et al., 2009; Arieli et al., 2011; Salant, 2011; Halevy and Mayraz, 2021). In particular, our participants' category-based behavior is reminiscent of the decision procedure suggested in Manzini and Mariotti (2012).

The paper proceeds as follows. Section 2 presents our experimental design and Section 3 describes the results at both the aggregate and individual levels. In Section 4, we introduce the two-stage qualitative tradeoff resolution model and classify the T_0 participants into theory-based types according to their choices. In Section 5, we examine the extent to which the participants are able to estimate the stopping rules' induced probabilities and investigate the influence of lack of probabilities in T_0 by analyzing behavior in T_p . Section 6 concludes.

2. Experimental Design

The experiment was carried out in the Interactive Decision-Making Lab at Tel Aviv University in April–May 2017. The participants were 114 Tel Aviv University undergraduate students in management and STEM, 44% of whom were women. The average age was 25. Recruitment of participants was done via ORSEE (Greiner, 2004).

Each participant received 55 NIS (roughly \$15) at the beginning of the experiment. In an attempt to make the participants internalize this endowment, one week prior to the session we

notified them that they would receive this amount and could lose part of it (at most 30 NIS) or win an additional amount, depending on their choices in the experiment. A reminder of that was sent on the day before the session as well. The experiment included 57 computerized decision problems (we refer to these decision problems as Questions 1–57 or Q1–Q57), one of which was randomly selected at the end of the experiment to determine the payment for the participants. The amount won (or lost) in that game was added to (or subtracted from) the initial endowment. In practice, each participant could win at most an additional 45 NIS and could lose at most 28 NIS of her initial endowment. All sessions were completed within an hour.

2.1 Detailed description of the experiment

In each session, the participants were randomly assigned to two treatments, denoted by T_0 and T_p , each with four parts, which are described below. Of the 114 participants, 67 participants were assigned to our main treatment, T_0 , and 47 were assigned to⁷ T_p . The complete questionnaire can be found in Appendix B. In short, Part A (respectively, Part B) examines the choice of a stopping rule when the baseline lottery has a negative (respectively, positive) expected value, and Part C explores the participants' ability to estimate the rules' induced probabilities. Part D studies the participants' behavior in a simpler setting to identify whether their choices in Parts A and B are related to a pure taste for skewed lotteries.

Part A. In this part, participants faced a sequence of computerized lotteries, each with an 18/37 probability of winning 1 NIS and a 19/37 probability of losing 1 NIS. These probabilities resemble the win/loss probability in the “Red or Black” European roulette game. In each decision problem, the participants were asked to choose a cutoff stopping rule. The participants faced 18 decision problems, in each of which they chose one out of five alternative cutoff stopping rules. If one of these problems was randomly selected for payment, then the stopping rule was automatically and instantaneously implemented by the computer.

The only difference between the two treatments was that in T_p the participants were informed about the probability of ending the game with a gain given each of the five stopping rules, whereas in T_0 they were not (in both treatments the participants were informed about the winning probability in the baseline lottery). This difference allows us to explore the extent to

⁷ Participants were assigned to the main treatment with probability 0.6.

which the patterns in T_0 result from the participants' lack of knowledge of the rule's induced probability.

We considered three types of decision problems, which are illustrated in Figure 3. In Q1–Q6 (*fixed loss*), the participants were required to choose between five stopping rules that induce the same potential loss and vary in the potential gains that they induce. In Q7–Q12 (*fixed gain*), the participants were required to choose between five stopping rules that induce the same potential gain and vary in the potential losses that they induce. In Q13–Q18 (*not fixed*), the stopping rules vary in both the potential gains and the potential losses that they induce.

In each decision problem, there were two rules in which the potential loss was greater than the potential gain, two rules in which the potential gain was greater than the potential loss, and one rule in which the potential gain and the potential loss were equal. We refer to these rules as left-biased, right-biased, and symmetric rules, respectively. We refer to the most left-biased rule (i.e., with the largest loss and the smallest gain) as *Rule ll*, the second-most left-biased rule as *Rule l*, the symmetric rule as *Rule s*, the most right-biased rule (i.e., with the largest gain and smallest loss) as *Rule rr*, and the second-most right-biased rule as⁸ *Rule r*. The five stopping rules were presented to the participants either in order from the left-biased rule with the largest loss and smallest gain to the right-biased rule with the largest gain and smallest loss (as in Figure 3) or in the reverse order.⁹ Thus, the five stopping rules were always ordered either from the highest probability of a gain to the lowest probability of a gain or the other way around.

Part B. This part consisted of 18 decision problems (Q19–Q36) and was similar in structure to Part A. The main difference between the two parts was that the probabilities of gain and loss in the baseline lottery were reversed in Part B (i.e., the probability of winning in a single lottery was 19/37). In addition, we tried to diversify the problems in Parts A and B to prevent a sense of repetition. Thus, the stopping rules in Part B were similar to the ones in Part A, yet they were not identical.

At the end of Parts A and B, the participants were asked to explain the principles that guided them in their choices. We examined the participants' explanations in order to obtain a better understanding of their reasoning process.

⁸ The notation *ll, l, s, r, rr* is for the reader's convenience and was not presented to the participants.

⁹ The randomly selected order was used consistently throughout Parts A and B. The results suggest that the order did not affect the choices in the experiment and hence we merged the data from the two variations in the analysis.

Type (i): Fixed loss

	Loss	Gain	Probability of gain
<i>Rule ll</i>	-21	+9	52%
<i>Rule l</i>	-21	+15	35%
<i>Rule s</i>	-21	+21	24%
<i>Rule r</i>	-21	+27	17%
<i>Rule rr</i>	-21	+33	12%

Type (ii): Fixed gain

	Loss	Gain	Probability of gain
<i>Rule ll</i>	-20	+12	42%
<i>Rule l</i>	-16	+12	39%
<i>Rule s</i>	-12	+12	34%
<i>Rule r</i>	-8	+12	28%
<i>Rule rr</i>	-4	+12	18%

Type (iii): Not fixed

	Loss	Gain	Probability of gain
<i>Rule ll</i>	-27	+15	38%
<i>Rule l</i>	-24	+18	31%
<i>Rule s</i>	-21	+21	24%
<i>Rule r</i>	-18	+24	19%
<i>Rule rr</i>	-15	+27	14%

Figure 3. The three types of questions in Part A. The probability of a gain for each stopping rule is provided for the reader's convenience. Only participants in T_p received information on the probability of a gain and a loss for each stopping rule, which was presented in a sentence below the description of the rule's upper and lower cutoffs (see Appendix B).

Part C. This part included three problems (Q37–Q39), where each problem presented a different stopping rule. In each of the three problems, the participants were asked to consider a baseline lottery that paid 1 NIS with probability $18/37$ and -1 NIS with probability $19/37$ (as in Part A) and to estimate the probability that the game would end with a gain, given the stopping rule. In particular, in the first problem, they had to gauge the probability of finishing the game with a gain of 25, given that the stopping rule was $(-25, +25)$. The second and third problems were similar except that the stopping rules were $(-25, +50)$ and $(-25, +100)$, respectively. The correct answers to these three questions were roughly 20.5%, 5%, and 0.3%, respectively. The payment for each of the problems in Part C (in case one of these problems was selected for payment) was 40 NIS minus the size (in absolute terms) of the error in the participant’s estimation. There was no difference between the two treatments in this part.

Part D. In this part, the participants faced 18 decision problems (Q40–Q57). In each problem they chose between two binary *lotteries* with known probabilities of loss and gain, as illustrated in Figure 4. In each problem, the two lotteries were a “mirror image” of each other (i.e., $-x$ with probability p and $+y$ with probability $1 - p$ vs. $-y$ with probability $1 - p$ and $+x$ with probability p), and had an expected value of roughly 0. In fact, we chose the prizes and the probabilities of the lotteries to reflect two stopping rules, one right-biased rule and one left-biased rule, with a baseline lottery’s winning probability of¹⁰ 0.5. When the baseline lottery is fair, right-biased rules induce positively skewed lotteries and left-biased rules induce negatively skewed lotteries. In each problem, the order of appearance of the two lotteries was randomly and independently determined. There was no difference between the two treatments in this part.

The participants’ decisions in this part were simpler than those in Parts A and B in two main dimensions: the winning probabilities were given and the lotteries were not presented as stopping rules. The lotteries’ mirror structure together with the simplicity of the setting allowed us to better understand the participants’ preference for skewed prospects and connect it to their choices between stopping rules in the main parts of the experiment (Parts A and B). This analysis can be found in Appendix A.2.

¹⁰ We structured the lotteries as follows: we simulated a stopping problem with a repeated lottery that yielded +1 with probability 0.5 and -1 with probability 0.5. We examined what would be the induced probabilities of a stopping rule with the bounds $-y$ and $+x$, and rounded the probabilities to make the problem seem simpler. Then, we did the same for $-x$ and $+y$.

Part D: Game 2

Choose your preferred lottery from the following two lotteries:

<i>a.</i>	probability	24%	76%
	amount	-25	+8
<i>b.</i>	probability	76%	24%
	amount	-8	+25

Figure 4. An example of a decision problem in Part D.

Discussion: Choosing from restricted sets of rules

In each of the decision problems in Parts A and B, the participants chose one out of five stopping rules. Alternatively, we could have asked them to make a single choice of the stopping rule's upper and lower bounds within a range $[X, Y]$, where $X < 0$ and $Y > 0$. We decided to let the participants face many problems and varied the sets of rules they faced for two reasons. First, this enabled us to focus on the effects of some fundamental properties of the stopping rules (e.g., the effect of the favorability of the baseline lotteries, how the choices differed given a fixed loss/gain, etc.) on the participants' choices while keeping the decision problems relatively simple. Second, observing choices from varied sets of rules reveals more information on the participants' preferences than a single choice when all rules are available. This additional information improves our ability to disentangle different theoretical explanations of the observed behavior.

Our restricted sets of rules resemble risk questionnaires that investment banks often use to elicit investors' preferences over investment strategies. In these questionnaires, individual investors often have to choose pairs of cutoffs that represent the maximal loss that they are willing to bear in a given time period and the gains that they expect to obtain in that period. In practice, investors are often given a fixed set of cutoffs to choose from rather than allowed to choose the cutoffs freely. Fixing the set of cutoffs allows the bank to classify the investors into a manageable number of categories and implement an investment strategy suitable for each category.

3. General Description of the Participants' Choices

We now focus on the main treatment, T_0 , in which the participants were not provided with the rules' induced probabilities. In Section 5.2 we shall present the results obtained in T_p , in which the rules' induced probabilities were provided, and compare them to the results in T_0 . We categorized the participants' choices into (i) left-biased rules (l and ll) and right-biased rules (r and rr), and into (ii) extreme rules (ll and rr) and moderate rules (l and r). An additional potential category is of symmetric rules. However, such rules were not frequently chosen. These categorizations enabled us to present the participants' choice patterns succinctly. We describe the behavior in Part A and Part B side by side, which allows us to observe both differences and similarities in choice patterns.

3.1 Aggregate-level data

At the aggregate level, in each part of the experiment there were 1,206 choices (67×18). We found that 66% of the choices in Part A were of left-biased rules and only 25% were of right-biased ones. Remarkably, only 9% of the choices were of the symmetric rule. In Part B, 46% of the chosen rules were left-biased whereas 35% were right-biased (see Table 1). The symmetric rule was chosen in 19% of the cases. Thus, the choices in Part B reflect a somewhat weaker tendency to choose left-biased rules than those in Part A. Finally, in each part, extreme and moderate rules were chosen with roughly similar proportions.

	Part A ($p < 0.5$)	Part B ($p > 0.5$)
<i>Rule ll</i>	31%	23%
<i>Rule l</i>	35%	23%
<i>Rule s</i>	9%	19%
<i>Rule r</i>	10%	19%
<i>Rule rr</i>	15%	16%

Table 1. The proportions of choices in T_0 , out of 1,206 (67×18) choices that were made in each part.

3.2 Individual-level analysis

Examining the participants' choices at the individual level reveals that many of them were consistent in their tendency to choose either left-biased rules or right-biased rules. To measure

the extent of this tendency, we consider the number of times each participant chose a left-biased rule, which ranges from 0 to 18 in each of the main parts of the experiment, A and B. We refer to this measure as the *number of left-biased choices*. In a similar manner, we consider the number of times each participant chose a right-biased rule and refer to this measure as the *number of right-biased choices*. It turns out that 70% of the participants in Part A and 63% of the participants in Part B chose the same category of rules in more than two-thirds of the decision problems. That is, for these participants, either the number of left-biased choices or the number of right-biased choices was 13 or higher (out of 18). The probability of observing such a pattern when a participant chooses uniformly at random is less than 1%.

The number of left-biased choices is higher on average in Part A than it is in Part B, according to a paired-samples t-test¹¹ (11.9 vs. 8.3, $t(66) = 4.44, p < 0.001$). Figure 5a shows that the cumulative distribution of the number of left-biased choices per individual in Part A stochastically dominates the corresponding distribution in Part B. The number of right-biased choices is higher on average in Part B than it is in Part A (6.31 vs. 4.54, $t(66) = -2.57, p = 0.012$). Figure 5b shows that the cumulative distribution of the number of right-biased choices per individual in Part B first-order stochastically dominates the corresponding distribution in Part A.

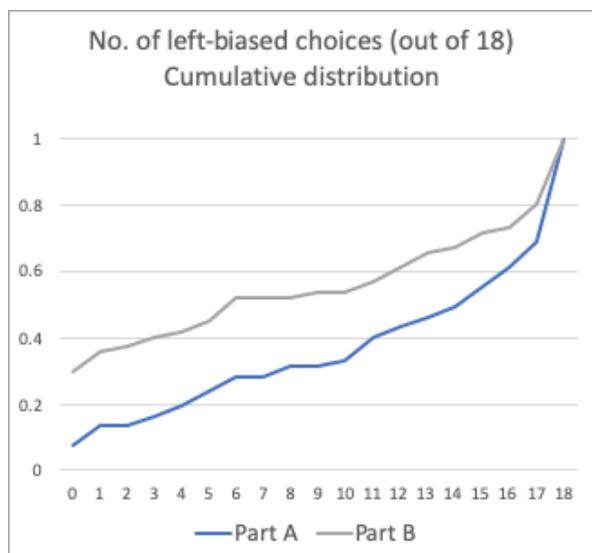


Figure 5a. Cumulative distribution of the number of left-biased choices per participant in Part A vs. Part B.

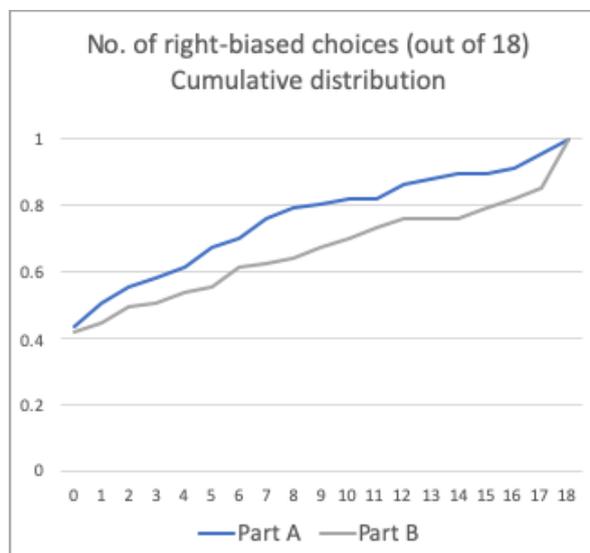


Figure 5b. Cumulative distribution of the number of right-biased choices per participant in Part A vs. Part B.

¹¹ All the statistical results in Section 3 are robust to non-parametric testing.

Despite the differences in the participants' behavior in Parts A and B, their choices in these two parts are highly correlated in terms of the number of left-biased choices (Pearson's $r = 0.56$, $p < 0.001$) and in terms of the number of right-biased choices (Pearson's $r = 0.64$, $p < 0.001$). The combination of these findings suggests that there exists an individual tendency either to choose left-biased rules or to choose right-biased rules, though the favorability of the baseline lottery reduces the tendency to choose left-biased rules.

Comment: Directional bias in different types of problems.

Examining each of the 36 decision problems in Parts A and B separately suggests that left-biased choices are more prevalent than right-biased ones in all but two of them. In Appendix A.1, we thoroughly examine how the type of problem (i.e., whether the loss or gain is fixed) affects the tendency to choose left-biased rules. Although there are some differences in behavior between the problems, in all the types of problems, there are more participants who consistently choose left-biased rules than participants who consistently choose right-biased rules.

Moving on to choices of extreme vs. moderate rules, we consider the number of times each individual chose an extreme rule, which ranges from 0 to 18 in each of the main parts of the experiment, A and B. We refer to this measure as the *number of extreme choices*. In a similar manner, we consider the number of times each individual chose a moderate rule and refer to this measure as the *number of moderate choices*.

Figures 6a and 6b show the cumulative distribution of these measures and suggest that the participants' behavior in Parts A and B is more similar in terms of extreme and moderate choices than in terms of left- and right-biased choices. It turns out that 57% of the participants in Part A and 47% of the participants in Part B chose the same category of rules in more than two-thirds of the decision problems. That is, for these participants, either the number of extreme choices or the number of moderate choices was 13 or higher (out of 18). The rest of the participants diversified between extreme and moderate rules. As before, the participants' behavior in the two parts of the experiment is highly correlated. The number of extreme choices in Part A is correlated with the corresponding number in Part B (Pearson's $r = 0.75$, $p < 0.001$). The number of moderate choices in Part A is correlated with the corresponding number in Part B (Pearson's $r = 0.58$, $p < 0.001$).

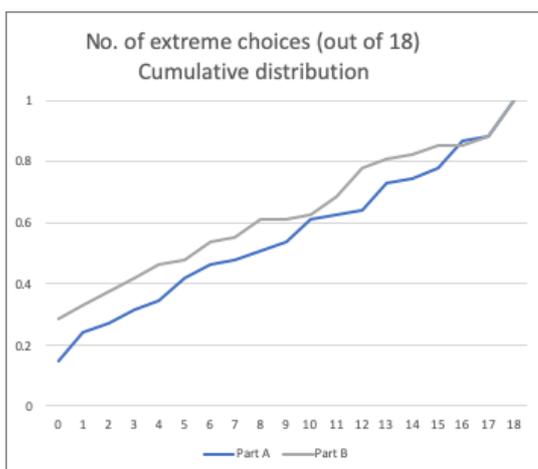


Figure 6a. Cumulative distribution of the number of extreme choices per participant in Part A vs. Part B.

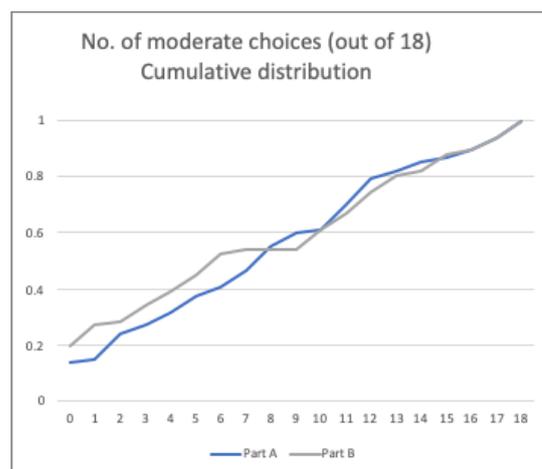


Figure 6b. Cumulative distribution of the number of moderate choices per participant in Part A vs. Part B.

3.3 Joint analysis of Parts A and B

While there are subtle differences between the participants' behavior in the context of favorable and unfavorable baseline lotteries, overall, it seems that their behavior across different parts of the experiment is highly correlated. Thus, it makes sense to examine the behavior in Parts A and B jointly. Figure 7 presents the cumulative distributions of the number of choices of left-biased rules and the number of choices of right-biased rules in the two parts of the experiment together. In addition to these distributions, as a benchmark, the figure presents the distribution that is obtained if individuals choose a stopping rule uniformly at random. This comparison illustrates the participants' strong tendency to either consistently choose left-biased rules or consistently choose right-biased rules across different problems. While the probability of choosing at least 22 times a rule that is biased in a particular direction is less than 1% when choosing uniformly at random, the figure shows that roughly 66% of our participants chose either at least 22 left-biased rules or at least 22 right-biased rules.

A similar pattern is obtained when we consider the participants' tendency to choose an extremely biased rule or a moderately biased rule. This is illustrated in Figure 8, which presents the distribution of the number of choices of extreme and moderate rules against a benchmark of choosing uniformly at random. Here too, roughly 66% of the participants are at the tails of the benchmark distribution.

In the next section, we dig deeper into the individual-level behavior and suggest a model that is based on these categorizations and accounts for the above behavior.

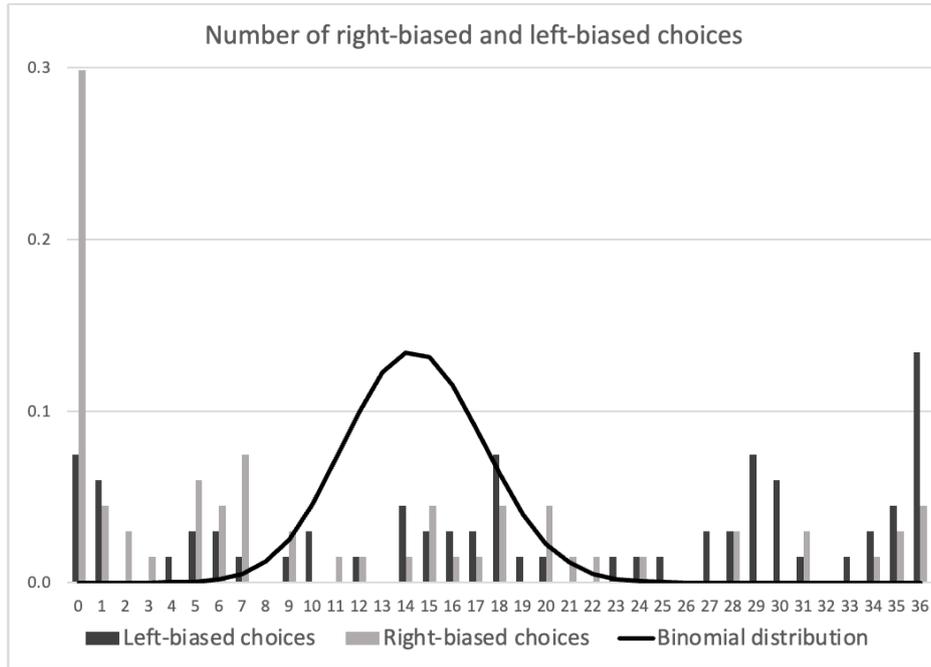


Figure 7. The percentage of participants with each number of right-biased and left-biased choices in Parts A and B together and the probability of observing each number (per participant) given a binomial process with 0.4 probability of a right- (left-) biased choice in each problem.

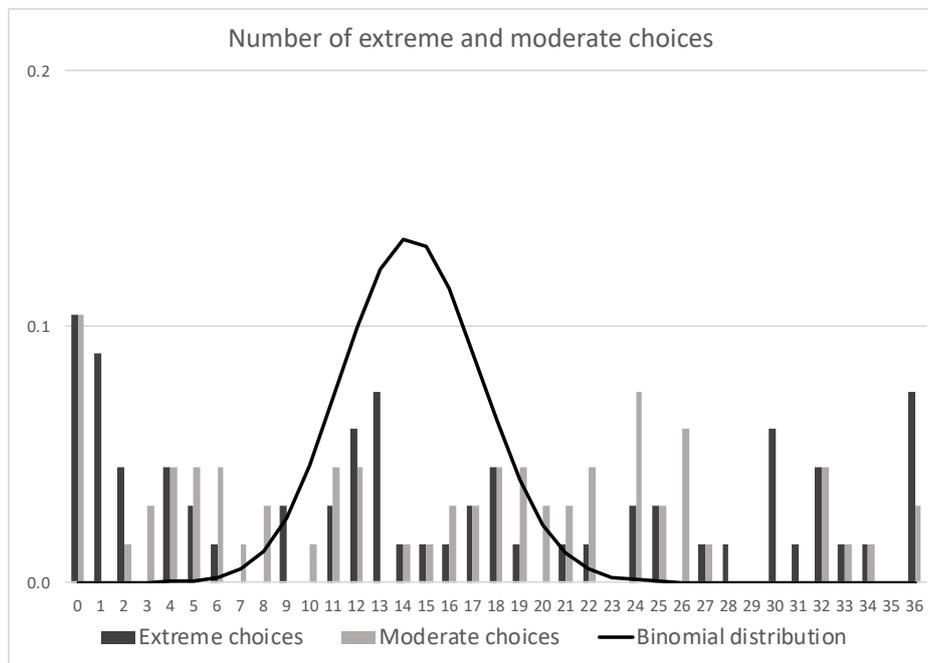


Figure 8. The percentage of participants with each number of extreme and moderate choices in Parts A and B together and the probability of observing each number (per participant) given a binomial process with a 0.4 probability of an extreme (moderate) choice in each problem.

4. A Model: Two-Stage Qualitative Tradeoff Resolution (2S-QTR)

In a stopping problem, participants essentially choose a potential gain and a potential loss. The larger the gain is, the less likely a participant is to finish the game with a gain, and the larger the loss is, the more likely she is to finish the game with a gain. Thus, when facing a stopping problem, individuals trade off between prizes and probabilities. As this qualitative feature is intuitive and easy to grasp (as we shall establish in the discussion of Part C's results in Section 5.1), we suggest that this tradeoff is solved in a qualitative manner. Although the solution may be affected by the context, as explained in the previous section, participants appear to be consistent in the way they solve the tradeoff (i.e., they are virtually unaffected by the fine details of the problems). However, there are different types of individuals who tend to resolve the prize–probability tradeoff in different manners.

To capture the qualitative reasoning described above, we introduce the two-stage qualitative tradeoff resolution model (2S-QTR), which is essentially a qualitative variation of the “categorize then choose” model (Manzini and Mariotti, 2012). In this model, individuals categorize the stopping rules at their disposal into two categories: one that consists of left-biased rules and one that consists of right-biased ones. The former category reflects the resolution of the prize–probability tradeoff in favor of a high probability of winning (or, consistent with the participants' explanations, a low probability of losing), whereas the second category reflects the resolution of the tradeoff in favor of large potential prizes (and smaller losses). After choosing a category, the same tradeoff is then resolved within the category: either in the same direction (i.e., choosing the most right-biased rule rr or the most left-biased rule ll) or in the opposite direction (choosing one of the moderately biased rules: r or l). Thus, individuals implement a (two-stage) sequential procedure, starting from “the big picture” (resolving the prize–probability tradeoff between the two categories of left-biased and right-biased rules), and then deciding about the “fine details” (resolving the prize–probability tradeoff within a category between extreme and moderate rules).¹²

The 2S-QTR model captures the above ideas by assuming that each participant i can be described by three parameters: α_i , β_i , and ϵ_i . The first parameter, α_i , reflects the participant's tendency to choose a left-biased category. The second parameter, β_i , reflects her tendency to choose an extreme rule within a category. The third parameter, ϵ_i , is a random noise term that reflects the probability that she chooses uniformly at random. Formally, with probability ϵ_i

¹² While we describe a natural order of these two stages, our formalism is independent of that order.

participant i chooses a rule uniformly at random. Conditional on not choosing a rule uniformly at random, she chooses a rule in the left-biased category with probability α_i and a rule in the right-biased category with probability $1 - \alpha_i$. Within each category she chooses an extreme rule with probability β_i and a moderate rule with probability $1 - \beta_i$. Thus, Table 2 specifies the probability that participant i chooses rule $j \in \{ll, l, m, r, rr\}$ in a given problem.¹³

Rule	Probability
<i>ll</i>	$0.2\epsilon_i + (1 - \epsilon_i)\alpha_i\beta_i$
<i>l</i>	$0.2\epsilon_i + (1 - \epsilon_i)\alpha_i(1 - \beta_i)$
<i>s</i>	$0.2\epsilon_i$
<i>r</i>	$0.2\epsilon_i + (1 - \epsilon_i)(1 - \alpha_i)(1 - \beta_i)$
<i>rr</i>	$0.2\epsilon_i + (1 - \epsilon_i)(1 - \alpha_i)\beta_i$

Table 2. The probability that participant i chooses rule $j \in \{ll, l, m, r, rr\}$ in a given problem in the 2S-QTR model.

To examine whether the 2S-QTR model explains the behavior of a large share of our participants, we performed a leave-one-out prediction exercise. For each participant and each of the 36 problems in Parts A and B, we employed a maximum likelihood estimation of the model's parameters based on the participant's choices in the other 35 problems. Subsequently, based on these estimated parameters, we tried to predict the answer to the 36th problem. We classified a participant as a 2S-QTR type if the number of correct predictions in this exercise was 14 or higher. The guiding principle in choosing the threshold was that the probability of predicting 14 or more choices when a participant chooses rules uniformly at random is less than 1%. Importantly, our results remain virtually the same if we choose a cutoff of 13 or 15 choices (the probability of predicting at least 13 or 15 choices when a participant chooses uniformly at random is less than 2% or 0.25%, respectively). Overall, we classified 50 (74.6%) of the participants in our main treatment, T_0 , as exhibiting behavior consistent with the 2S-QTR model. The mean number of predicted choices for these participants was 24.24.

Next, to study the behavior of the participants who chose consistently with the 2S-QTR model, we estimated the model's parameters for each of them based on all 36 problems in Parts A and B. We used the estimation to classify the participants into types. Participants whose tendency to choose the left-biased rules category, α , was significantly greater (resp., less) than

¹³ While the 2S-QTR formalization is tailored to our experimental setting, it is possible to modify it to situations in which the set of possible stopping rules has a different structure.

0.5 (at the 5% level) were classified as L (resp., R) types. The remaining seven participants were classified as unbiased. We took a similar approach when classifying participants as extreme and moderate. We classified as extreme (resp., moderate) types participants whose tendency to choose an extreme rule within a category, β , was significantly greater (resp., less) than 0.5. The remaining participants were classified as diversifying. Tables 3 and 4 present this classification, where, for each type, L and R, we report the average number of predicted choices in the leave-one-out exercise, the average number of choices of left-biased, right-biased, moderate, and extreme rules, and the average estimated parameters, α and β . The tables show that estimated parameters for right-biased/left-biased/moderate/extreme types are quite far from 0.5, and that their choices in the experiment match their classification.

<i>L types</i>	Extreme	Moderate	Diversifying	Overall
<i>Proportion (n)</i>	24% (16)	12% (8)	9% (6)	45% (30)
<i>Mean alpha (SD)</i>	0.93 (0.09)	0.98 (0.05)	1 (0)	0.96 (0.08)
<i>Mean beta (SD)</i>	0.84 (0.10)	0.12 (0.14)	0.47 (0.01)	0.57 (0.34)
<i>No. of predicted choices (SD)</i>	26.25 (4.95)	27 (5.37)	18.16 (2.64)	24.83 (5.84)
<i>No. of left choices (SD)</i>	32 (3.66)	31.62 (4.50)	31.66 (4.41)	31.83 (3.86)
<i>No. of right choices (SD)</i>	3.12 (3.58)	1.75 (2.55)	1.16 (2.04)	2.36 (3.06)
<i>No. of moderate choices (SD)</i>	6.31 (3.74)	28.12 (5.36)	17.33 (3.61)	14.33 (10.49)
<i>No. of extreme choices (SD)</i>	28.81 (4.36)	5.25 (5.44)	15.5 (3.21)	19.86 (11.38)

Table 3. A description of the 2S-QTR model's L types: the proportions of extreme, moderate, and diversifying L types; the estimated parameters α and β ; the number of predicted choices; the mean number of choices of left-biased, right-biased, moderate, and extreme rules.

As suggested in Table 3, a large share of the participants in T_0 exhibited a tendency toward left-biased rules. In fact, 30 (45%) participants were classified as L types. While the formal selection rule was 14 correct predictions according to our model, on average, the number of correct predictions for the 2S-QTR model's L types was 24.83. Participants who were classified as L-extreme and L-moderate exhibited behavior that was consistent with their classification as the average number of choices of left-biased rules was 31.83, the number of extreme rules chosen by extreme types was 28.8 and the number of moderate rules chosen by moderate types was 28.1.

<i>R types</i>	Extreme	Moderate	Diversifying	Overall
<i>Proportion (n)</i>	10% (7)	4% (3)	4% (3)	19% (13)
<i>Mean alpha (SD)</i>	0.06 (0.12)	0.12 (0.08)	0.04 (0.04)	0.07 (0.1)
<i>Mean beta (SD)</i>	0.95 (0.07)	0.04 (0.06)	0.36 (0.08)	0.6 (0.41)
<i>No. of predicted choices (SD)</i>	30.14 (5.84)	27 (6.66)	18.67 (0.58)	26.76 (7.05)
<i>No. of left choices (SD)</i>	2.71 (4.46)	4.66 (3.21)	1.66 (2.08)	2.92 (3.68)
<i>No. of right choices (SD)</i>	32.28 (4.79)	25.66 (8.14)	31 (3)	30.46 (5.64)
<i>No. of moderate choices (SD)</i>	3.14 (4.34)	28 (3.46)	20.33 (1.53)	12.85 (11.75)
<i>No. of extreme choices (SD)</i>	31.85 (5.49)	2.33 (1.53)	12.33 (3.06)	20.53 (13.84)

Table 4. A description of the 2S-QTR model's R types: the proportions of extreme, moderate, and diversifying R types; the estimated parameters α and β ; the number of predicted choices; the mean number of choices of left-biased, right-biased, moderate, and extreme rules.

A smaller share of the participants exhibited a tendency to choose right-biased rules. Again, the mean number of predicted choices according to our model was much larger than the cutoff of 14 predictions as, overall, the average number of predictions of participants who were classified as the 2S-QTR model's R types was 26.77. As before, the R types exhibited behavior that was consistent with their classification, which is reflected in the average values of α and β as well as in the number of choices of right-biased, extreme, and moderate rules.

4.1 Alternative theory-based explanations

A natural question that arises is whether there is a different, more standard explanation of the findings described above. In this section, we address this question by considering leading theories of decision-making under risk: expected utility with risk aversion, disappointment aversion (DA; Gul, 1991), regret aversion (RA; Bell, 1982; Loomes and Sugden, 1982), salience theory (ST; Bordalo et al., 2012), and cumulative prospect theory (CPT; Kahneman and Tversky, 1992).

Before examining the different theories, it will be useful to examine a relatively simple explanation of our findings. To this end, we explore the behavior of a risk-neutral expected utility maximizer who faces the decision problems in our experiment. Not only is expected value maximization a special case of all the theories we examine, but it can also provide clear intuition for the quantitative reasoning in our experimental setting. The next observation establishes that

expected value maximization implies a completely different ranking of the rules, $\{ll, l, s, r, rr\}$, depending on which bound is fixed in the problem and whether the baseline lottery is favorable. Expected value maximization is, therefore, inconsistent with our findings.¹⁴

Observation 1. *An expected value maximizer would rank the rules as $ll > l > s > r > rr$ in Questions 1–6 and 25–30, and as $rr > r > s > l > ll$ in Questions 7–12 and 19–24.*

Proof. The expected value from choosing rule i is $e \times n_i$, where e is the expected value of the baseline lottery and n_i is the expected number of baseline lotteries played given rule i . When $e < 0$ (as in Part A) the expected value of rule i is decreasing in n_i and when $e > 0$ (as in Part B) its expected value is increasing in n_i . Denote by $l_i < 0$ and $h_i > 0$ the potential loss and gain given rule i , respectively. Consider two rules, i and j , such that $l_i = l_j$ and $h_i > h_j$. Observe that $n_i > n_j$ as the two rules induce the same number of lotteries if the process reaches l_j before reaching h_j , and otherwise rule i results in more lotteries. This proves the claim for Questions 1–6 and 19–24. Symmetrically, consider two rules, i and j , such that $h_i = h_j$ and $l_j > l_i$. Here too $n_i > n_j$ as the two rules induce the same number of lotteries if the process reaches h_j before reaching l_j , and otherwise rule i results in more lotteries. This proves the claim for Questions 7–12 and 25–30. ■

We conclude that an expected value maximizer would diversify between the most right-biased rule and the most left-biased rule, depending on which bound is fixed and whether the baseline lottery is favorable or not.

After establishing that expected value maximization cannot account for the main patterns in the data, we consider the more nuanced theories and compare their success rate in predicting the data. For each theory, we consider a prominent specification and run a leave-one-out prediction exercise, per participant, similar to the one performed for the 2S-QTR model. We employ a maximum likelihood estimation of each participant’s parameters (assuming a multinomial logit choice model).¹⁵ We begin by exploring the overall performance of each model

¹⁴ In the remaining problems in our experiment, in which no bound is fixed, the ranking of the rules is more nuanced and remains inconsistent with our findings.

¹⁵ The descriptions of the theories’ specifications appear in Appendix A.3. For some of the theories, there is more than one workhorse specification. In such cases, we estimated more than one specification and reported the results for the specification that was consistent with the behavior of the largest share of participants (e.g., for expected utility theory, we estimated both a specification with constant relative risk aversion and a specification with constant absolute risk aversion).

separately, allowing different participants to be characterized by different model parameters. Aggregating over all the participants in T_0 , i.e., considering 2,412 (67×36) choices between 5 rules, the 2S-QTR predicts 1,290 choices, CPT predicts 1,015 choices, ST predicts 645 choices, DA predicts 591 choices, RA predicts 585 choices, and CRRA predicts 443 choices. Thus, 2S-QTR and CPT predict a substantially larger number of choices than the other theories we consider.

Next, we run a prediction competition between these theories, *per participant*, and classify a participant into a theory if (i) the theory predicts at least 14 of the participant’s choices (a criterion that is identical to the one used in the 2S-QTR classification exercise above), and (ii) there is no other theory that predicts a higher number of choices. Table 5 summarizes the prediction competition for participants in T_0 . It illustrates three main findings. First, even when we allow for alternative explanations, the share of individuals whose behavior is best explained by the 2S-QTR model is 69%. This suggests that there is no better explanation of our participants’ behavior among the prominent specifications of decision-making under risk theories. Second, a considerable share of the participants (33%) were classified as CPT types, where about two-thirds of them were also classified as 2S-QTR types (i.e., the two theories tie for those participants). Third, only a small share of the participants were classified into one of the other theories.

Theory	Proportion (n)
<i>2S-QTR</i>	69% (46)
<i>Constant Relative Risk Aversion</i>	1.5% (1)
<i>Disappointment Aversion</i>	1.5% (1)
<i>Regret Aversion</i>	3% (2)
<i>Saliency Theory</i>	4.5% (3)
<i>Cumulative Prospect Theory</i>	33% (22)

Table 5. The proportion and the number (in parentheses) of participants in T_0 , out of the 67 participants, that were classified into each of the decision theories.

A potential limitation of our classification method is that a theory is disqualified as an explanation for a participant’s behavior if it is only slightly outperformed by another theory. In Appendix A.4, we modify our classification method in a manner that relaxes the competition between the different theories. In this robustness exercise we consider a theory to be a plausible

explanation of a participant's behavior if it is the best at predicting that participant's behavior or if it is only slightly less successful than the best predicting theory. The results of this exercise are similar to those presented in Table 5 and can be found in Table S5 in the appendix.

Recent findings by Ebert and Karehnke (2021) provide an intuition for why CPT seems to be the best explanation for our participants' behavior among the quantitative theories we considered. Ebert and Karehnke show that among the leading theories of decision-making under risk, CPT is essentially the only theory that can imply both skewness-seeking and skewness-averse behavior, depending on the parameters. To see the connection, let $skew_i$ be the skewness of the lottery induced by rule $i \in \{rr, r, s, l, ll\}$ and note that, in all of the problems in Parts A and B, it holds that $skew_{rr} > skew_r > skew_s > skew_l > skew_{ll}$. Thus, skewness-seeking is closely related to choosing right-biased rules and skewness aversion is closely related to choosing left-biased rules. As for CPT, Ebert and Karehnke (2021) suggest that skewness-seeking follows from probability weighting that overweights small probabilities and underweights large probabilities, whereas skewness aversion follows from a diminishing sensitivity to gains and losses. While Ebert and Karehnke's definition of skewness-seeking and skewness aversion is "in the small," the next observation considers the canonic representation suggested by Kahneman and Tversky (1992) and shows that sufficiently overweighting small probabilities implies a preference for right-biased rules, whereas sufficiently diminishing sensitivity to gains and losses implies a preference for left-biased rules.

Observation 2. Consider the representation suggested by Kahneman and Tversky (1992):

$$u(x) = \begin{cases} -\lambda(-x)^\alpha & \text{for } x < 0 \\ x^\alpha & \text{for } x \geq 0 \end{cases}$$

with the probability weighting function

$$w(p) = \frac{p^\delta}{(p^\delta + (1-p)^\delta)^{\frac{1}{\delta}}}.$$

Fix a problem from Part A or Part B.

- (i) For any $\lambda \geq 1$ and $\alpha \leq 1$ there exists a δ^* such that for all $\delta \in (0, \delta^*)$ an individual with such preferences would rank the rules as $rr > r > s > l > ll$.
- (ii) For any $\lambda \geq 1$ and $\delta \leq 1$ there exists an α^* such that for all $\alpha \in (0, \alpha^*)$ an individual with such preferences would rank the rules as $ll > l > s > r > rr$.

Proof. Denote by U_i and L_i the absolute values of the upper and lower bounds of rule i , respectively. For part (i), note that $w(p)$ goes to 1 as δ goes to 0. Hence, in this limit rule i induces a value of $U_i^\alpha - \lambda L_i^\alpha$. It follows that a decision-maker prefers rule i to rule j if $U_i^\alpha - U_j^\alpha > \lambda(L_i^\alpha - L_j^\alpha)$. Since $\alpha \in (0,1]$ and $\lambda \geq 1$, this inequality implies the ranking $rr > r > s > l > ll$ in all of the problems in our experiment. The continuity of the weighting function completes the proof of (i). To prove (ii), consider the $\alpha = 0$ limit, in which rule i induces $w(q_i) - \lambda w(1 - q_i)$, where q_i is the probability of finishing the game with a gain. In this limit, the ranking of the rules is $ll > l > s > r > rr$ as the expected value induced by a rule depends only on the probability q_i . The continuity of the value function in α guarantees (ii). ■

We conclude that CPT can accommodate both consistently choosing right-biased rules and consistently choosing left-biased rules. It should be stressed, however, that it can also accommodate other choice patterns (in fact, Observation 1 shows such a pattern for $\alpha = \lambda = \delta = 1$), and that some of the participants who were classified as CPT types (6 out of 22) were not classified as either L or R types.

Let us take an even more conservative approach and classify participants into the 2S-QTR model only if it predicts a strictly greater number of the participant's choices than any of the other theories. Under this approach, 30 participants (45%) are classified as 2S-QTR types, 26 (39%) are classified into one of the quantitative theories considered, and 11 (16%) are unclassified. It is noteworthy that for the 30 participants classified into 2S-QTR, the quantitative theories' predictions are substantially less successful (e.g., for 24 of them, the difference in the number of predictions is at least 4).

We wish to point out that for some of the 22 CPT types, the estimated parameters are inconsistent with the range of parameters usually estimated in the literature (Stott, 2006). If we follow the literature and restrict CPT parameters to a more conventional range ($0.15 < \alpha < 1, 0.15 < \delta < 1, 1 < \lambda < 5$), then the theory explains the behavior of only 11 participants (16%) in T_0 and the 2S-QTR model becomes the *single* best explanation for 39 participants' behavior (58%).

The analysis in this section suggests that the 2S-QTR model provides the best explanation for the behavior of a considerable share of the participants. As for the other participants, the leading explanation is CPT.

5. The Absence of the Rules' Probabilities and Its Implications

In this section, we examine to what extent the choices of the participants in the main treatment, T_0 , were affected by not knowing the rules' induced winning probabilities. Not knowing the induced probabilities should have no effect if the participants can infer these probabilities from the likelihood of winning a single baseline lottery. Thus, the first step of the analysis must examine the participants' ability to make such an inference. Part C of the experiment explores this question and shows that the participants' inferences are very far from the true winning probabilities (consistent with Gneezy, 1996, and Halevy, 2007). In the second part of this section, we present the results of our second treatment, T_p , in which the induced probabilities were explicitly given to the participants. A comparison of the two treatments sheds light on the effects of the unknown probabilities on the participants' behavior.

5.1 Can the participants infer the rules' induced winning probabilities? (Part C)

In each of the three problems in Part C, we presented the participants with a stopping rule. The rules were $(-25, +25)$, $(-25, +50)$, and $(-25, +100)$ in the first, second, and third problems, respectively. The participants were asked to assess the rules' induced winning probabilities given that the probability of winning a single baseline lottery is $18/37$, as in Part A. The correct induced winning probabilities were 20.5%, 5%, and 0.3%, respectively.

The participants' average estimates in T_0 were 39.6%, 24.3%, and 17.4%. The mean errors in absolute terms were 23.2%, 20.6%, and 17.4%. Moreover, only 26.8% of the answers were within a range of 5% from the correct answer (e.g., an estimate of 15.6%–25.6% in the first problem in Part C).¹⁶ While most of the participants failed to estimate the winning probabilities correctly, they did exhibit a qualitative understanding of the prize–probability tradeoff, where 86.8% of them provided monotone estimates (an estimate is monotone if the estimate for $(-25, +25)$ is weakly greater than the estimate for $(-25, +50)$ and the latter is weakly greater than the estimate for $(-25, +100)$). The fact that the vast majority of the participants failed to estimate the induced winning probabilities provides additional motivation for our investigation of the T_p treatment in which the participants were provided with the rules' induced winning probabilities.

¹⁶ In T_p , where the participants observed the probabilities of the stopping rules in Parts A and B, the average estimates in Part C were 37%, 22.7%, and 12.3%, and the average error size slightly decreased in all three problems. Only 29.8% of the answers in T_p were within a range of 5% from the correct answer.

Our findings in Part C complement Gneezy’s (1996) findings, which relate to positive expected value lotteries. He finds that individuals use the stage-by-stage probability as an anchor and adjust insufficiently: estimations are biased toward the direction of the single-lottery probability, resulting in an underestimation of the overall probability of winning. The combination of these findings and our results can have significant implications for situations in which processes are perceived to be “almost fair.” It could lead to over-optimism and over-participation in situations where the baseline drift is slightly negative (e.g., casino gambling) and over-pessimism and under-participation in situations where the baseline drift is slightly positive (e.g., stock market trading).

5.2 Known vs. missing induced winning probabilities (T_p vs. T_0)

We briefly describe the behavior in T_p , in which the participants were provided with the stopping rules’ induced probabilities of winning and losing. We show both similar and different patterns from those observed in T_0 and compare the behavior statistically.

At the aggregate level, the behavior patterns that are exhibited by the T_p participants are mostly similar to the ones observed in T_0 . First, when the baseline lottery is unfavorable, there is a tendency to prefer left-biased stopping rules to right-biased ones. Second, this tendency is weaker when $p > 0.5$. We found that in Part A, 62% of the 846 choices (47×18) are of left-biased rules and 28% are of right-biased ones, whereas in Part B, 49% of the choices are of left-biased rules and 37% are of right-biased ones (see Table 6). A prominent difference from T_0 is the higher ratio of extreme vs. moderate rules in T_p : in both parts, extreme rules are more frequent than moderate rules within both the category of left-biased rules and the category of right-biased rules.

	T_p		T_0	
	Part A ($p < 0.5$)	Part B ($p > 0.5$)	Part A ($p < 0.5$)	Part B ($p > 0.5$)
<i>Rule ll</i>	44%	32%	31%	23%
<i>Rule l</i>	18%	17%	35%	23%
<i>Rule s</i>	11%	15%	9%	19%
<i>Rule r</i>	10%	13%	10%	19%
<i>Rule rr</i>	18%	24%	15%	16%

Table 6. The proportions of choices in T_p out of the 846 choices (47×18) that were made in each part, presented next to the proportions of choices in T_0 out of the 1,206 choices (67×18) in each part.

At the individual level, the mean number of choices of left-biased rules in Part A of T_p is higher than that in Part B of T_p , according to a paired-samples t-test¹⁷ (11.09 vs. 8.74, $t(46) = 2.96$, $p = 0.005$). The mean number of choices of right-biased rules in Part A of T_p is lower than that in Part B of T_p , according to a paired-samples t-test (4.98 vs. 6.6, $t(46) = -2.26$, $p = 0.028$). Nonetheless, the participants' choices in Part A and Part B are correlated in terms of the number of left-biased choices (Pearson's $r = 0.62$, $p < 0.001$) and the number of right-biased choices (Pearson's $r = 0.57$, $p < 0.001$). A comparison of the number of left-biased choices across the two treatments, T_0 and T_p , reveals that there are no significant differences in either part or overall (when the two parts are analyzed jointly). Similarly, there are no significant differences between the treatments in the number of right-biased choices.

By contrast, there are significant differences between the two treatments in the number of extreme and moderate choices. In particular, participants in T_p tended to choose the extreme stopping rules more often than those in T_0 in both parts and overall (the mean number out of 36 choices was 21.13 vs. 15.45, $t(112) = -2.88$, $p = 0.005$). Accordingly, the mean number of moderate rules was lower in T_p than in T_0 (10.28 vs. 15.6, $t(112) = 3.43$, $p < 0.001$). Thus, the uncertainty over the induced lotteries in T_0 mitigated the individuals' extreme choices.

The above observations suggest that the directional (left-biased vs. right-biased) tendencies in T_p are similar to those found in T_0 . A comparison of the two treatments indicates that in both parts, the distribution of our measures of the number of the participants' left-biased or right-biased choices is not significantly different between the treatments. Furthermore, there are no significant differences between the treatments in the number of left-biased or right-biased choices for any of the six types of questions, as described in Appendix A.1. The only significant difference in behavior between the treatments is the tendency mentioned above of choosing more extreme stopping rules in T_p (i.e., rules ll and rr are more common than rules l and r).

Consequently, roughly 74% of the participants in T_p exhibited behavior consistent with the 2S-QTR model, as in T_0 . Table 7 presents the classification into theory-based types in T_p , based on a leave-one-out prediction competition, as in Section 4.1. It appears that the main difference between the two treatments is that in T_p a larger share of the participants were classified as CPT types ($p < 0.017$, $\chi^2 = 5.73$), while a somewhat smaller share of the participants were classified as 2S-QTR types ($p = 0.058$, $\chi^2 = 3.61$).¹⁸ Among the 24 participants (51%) who were classified

¹⁷ All the results in Section 5.2 are robust to non-parametric testing.

¹⁸ We modify our classification method in a robustness exercise reported in Appendix A.4.

as 2S-QTR types, 17 were L types, 5 were R types, and 2 were unbiased. A possible interpretation of these differences is that when the probabilities of gains and losses are provided, the participants better recognize situations in which the tradeoff between prizes and probabilities is clear-cut and adjust their choices accordingly. For example, in situations where a minor deduction of a winning probability leads to a major increase in prizes, they tend to opt for the most right-biased rules. Symmetrically, in situations where a major increase in the probability of winning is accompanied by a small change in prizes, they tend to opt for the most left-biased rules. This may explain why the quantitative theories we considered, such as CPT, account for the behavior of a larger share of the participants in T_p .

Theory	T_p (N=47)	T_0 (N=67)
	Proportion (n)	Proportion (n)
<i>2S-QTR</i>	51% (24)	69% (46)
<i>Constant Relative Risk Aversion</i>	15% (7)	1.5% (1)
<i>Disappointment Aversion</i>	11% (5)	1.5% (1)
<i>Regret Aversion</i>	13% (6)	3% (2)
<i>Saliency Theory</i>	11% (5)	4% (3)
<i>Cumulative Prospect Theory</i>	55% (26)	33% (22)

Table 7. The proportion and the number (in parentheses) of participants in T_p who were classified into each of the decision theories, next to the corresponding results in T_0 .

CPT accounts for the behavior of 55% of the participants in T_p . As in T_0 , for some of the CPT types, the estimated parameters are inconsistent with the range of parameters usually estimated in the literature. If we follow the literature and restrict CPT's parameters to a more conventional range ($0.15 < \alpha < 1, 0.15 < \delta < 1, 1 < \lambda < 5$), then only 18 participants (38%) in T_p are classified as CPT types.

We conclude that there are both similarities and differences between the patterns of behavior observed in T_p and T_0 . First, the participants' tendency to consistently choose left- or right-biased rules is quite similar between the two treatments. Recall that the T_p and T_0 participants received the same experimental instructions. Moreover, the stopping rules' framing in the two treatments was identical: each alternative was presented as a lower and an upper bound rather than as a standard lottery. The only difference between the treatments was an

additional sentence in T_p that provided, for every alternative, the induced probabilities of reaching the lower and the upper bound given p . We suggest that the stopping rules' framing and the qualitative understanding of the prize-probability tradeoff in this context (established in Section 5.1) makes many participants reason in qualitative terms when trading off prizes and probabilities *even when the probabilities are known*.

Second, the participants' resolution of the tradeoff between extreme and moderate rules is different between the two treatments: the T_p participants tend to opt for extreme rules more than the T_0 participants. The tendency to choose extreme rules results in a larger share of participants who behave in a manner that is consistent with CPT, which also results in fewer 2S-QTR types in this more standard exercise (nonetheless, more than half of the participants were classified as such types). However, the results of T_0 suggest that in situations where participants lack information about the induced probabilities of stopping with a gain or a loss, as often occurs in reality, a larger share of them reason qualitatively.

6. Conclusion

We examined individuals' preferences over stopping rules when they have commitment power. We suggest a simple qualitative model whereby individuals tend to trade off between the size of the prize and the probability of winning in a consistent manner, either in favor of right-biased stopping rules or in favor of left-biased stopping rules. Then, they resolve this tradeoff again in favor of either the extreme or the moderate rule within the category of left- or right-biased rules. Our model accounts for behavior patterns in the data, which cannot be explained by prominent theories of decision-making under risk.

Our analysis suggests that many individuals use qualitative decision procedures even when the stopping rules' induced winning probabilities are known. These individuals consistently focus either on the winning probability or on the size of the potential gains and losses. More generally, our results provide indications of qualitative reasoning: individuals think in relative terms and are not responsive to a decision problem's fine numerical details. An interesting direction for future research would be to examine whether this type of reasoning arises in stopping problems in other contexts, such as job search and experimentation in R&D, as well as when choosing between other kinds of prospects.

References

- Alaoui, L. and Fons-Rosen, C. (2021): "Know When to Fold 'em: The Flip Side of Grit," *European Economic Review*, 136, 103736.
- Arieli, A., Ben-Ami, Y., and Rubinstein, A. (2011): "Tracking Decision Makers under Uncertainty," *American Economic Journal: Microeconomics*, 3, 68–76.
- Barberis, N. (2012): "A Model of Casino Gambling," *Management Science*, 58, 35–51.
- Barkan, R. and Busemeyer, J. (1999): "Changing Plans: Dynamic Inconsistency and the Effect of Experience on the Reference Point," *Psychonomic Bulletin & Review*, 6, 547–554.
- Barkan, R. and Busemeyer, J. (2003): "Modeling Dynamic Inconsistency with a Changing Reference Point," *Journal of Behavioral Decision Making*, 16, 235–255.
- Bell, D. E. (1982): "Regret in Decision Making under Uncertainty," *Operations Research*, 30, 961–981.
- Bleichrodt, H. and van Bruggen, P. (2018): "Reflection for Higher Order Risk Preferences," Working Paper.
- Bordalo, P., Gennaioli, N., and Shleifer, A. (2012): "Salience Theory of Choice under Risk," *Quarterly Journal of Economics*, 127(3), 1243–1285.
- Brunner, T., Levínský, R., and Qiu, J. (2011): "Preferences for Skewness: Evidence from a Binary Choice Experiment," *European Journal of Finance*, 17, 525–538.
- Cubitt, R. and Sugden, R. (2001): "Dynamic Decision-Making under Uncertainty: An Experimental Investigation of Choices between Accumulator Gambles," *Journal of Risk and Uncertainty*, 22, 103–128.
- Dertwinkel-Kalt, M., Frey, J., and Koster, M. (2020): "Optimal Stopping in a Dynamic Salience Model," Working Paper.
- Deck, C. and Schlesinger, H. (2010): "Exploring Higher-Order Risk Preferences," *Review of Economic Studies*, 77, 1403–1420.
- Deck, C. and Schlesinger, H. (2014): "Consistency of Higher-Order Risk Preferences," *Econometrica*, 82, 1913–1943.
- Ebert, S. (2015): "On Skewed Risks in Economic Models and Experiments," *Journal of Economic Behavior and Organization*, 112, 85–97.

- Ebert, S. and Karehnke, P. (2021): "Skewness Preferences in Choice under Risk," Working Paper.
- Ebert, S. and Strack, P. (2015): "Until the Bitter End: On Prospect Theory in a Dynamic Context," *American Economic Review*, 105, 1618–1633.
- Ebert, S. and Wiesen, D. (2011): "Testing for Prudence and Skewness Seeking," *Management Science*, 57, 1334–1349.
- Ebert, S. and Wiesen, D. (2014): "Joint Measurement of Risk Aversion, Prudence, and Temperance," *Journal of Risk and Uncertainty*, 48, 231–252.
- Feller, W. (1970): *An Introduction to Probability Theory and Its Applications*, 3rd Ed., Wiley.
- Fischbacher, U., Hoffmann, G., and Schudy, S. (2017): "The Causal Effect of Stop-Loss and Take-Gain Orders on the Disposition Effect," *Review of Financial Studies*, 30, 2110–2129.
- Fisher, M. and Keil, F. C. (2018): "The Binary Bias: A Systematic Distortion in the Integration of Information," *Psychological Science*, 29 (11), 1846–1858.
- Fisher, M., Newman, G. E., and Dhar, R. (2018): "Seeing Stars: How the Binary Bias Distorts the Interpretation of Customer Ratings," *Journal of Consumer Research*, 45(3), 471–489.
- Gneezy, U. (1996): "Probability Judgments in Multi-stage Problems: Experimental Evidence of Systematic Biases," *Acta Psychologica*, 93, 59–68.
- Golec, J. and Tamarkin, M. (1998): "Bettors Love Skewness, Not Risk, at the Horse Track," *Journal of Political Economy*, 106, 205–225.
- Greiner, B. (2004): "An Online Recruitment System for Economic Experiments," Working Paper.
- Grossman, P. and Eckel, C. (2015): "Loving the Long Shot: Risk-Taking with Skewed Lotteries," *Journal of Risk and Uncertainty*, 51, 195–217.
- Gul, F. (1991): "A Theory of Disappointment Aversion," *Econometrica*, 59, 667–686.
- Güth, W., Levati, V., and Ploner, M. (2009): "An Experimental Analysis of Satisficing in Saving Decisions," *Journal of Mathematical Psychology*, 53(4), 265–272.
- Halevy, Y. (2007): "Ellsberg Revisited: An Experimental Study," *Econometrica*, 75, 503–536.
- Halevy, Y. and Mayraz, G. (2021): "Identifying Rule-Based Rationality," Working Paper.
- Heimer, R., Iliewa, Z., Imas, A., and Weber, M. (2021): "Dynamic Inconsistency in Risky Choice: Evidence from the Lab and Field," Working Paper.

- Henderson, V., Hobson, D., and Zeng, M. (2022): "Cautious Stochastic Choice, Optimal Stopping and Deliberate Randomization," *Economic Theory*, 1–36.
- Loomes, G. and Sugden, R. (1982): "An Alternative Theory of Rational Choice under Uncertainty," *Economic Journal*, 92, 805–824.
- Maier, J. and Rüger, M. (2012): "Experimental Evidence on Higher-Order Risk Preferences with Real Monetary Losses," Working Paper.
- Manzini, P. and Mariotti, M. (2012): "Categorize Then Choose: Boundedly Rational Choice and Welfare," *Journal of the European Economic Association*, 10(5), 1141–1165.
- Noussair, C. N., Trautmann, S. T., and van de Kuilen, G. (2014): "Higher Order Risk Attitudes, Demographics, and Financial Decisions," *Review of Economic Studies*, 81(1), 325–355.
- Ploner, M. (2017): "Hold on to It? An Experimental Analysis of the Disposition Effect," *Judgment and Decision Making*, 12, 118–127.
- Salant, Y. (2011): "Procedural Analysis of Choice Rules with Applications to Bounded Rationality," *American Economic Review*, 101(2), 724–748.
- Sandri, S., Schade, C., Musshoff, O., and Odening, M. (2010): "Holding on for Too Long? An Experimental Study on Inertia in Entrepreneurs' and Non-entrepreneurs' Disinvestment Choices," *Journal of Economic Behavior & Organization*, 76, 30–44.
- Shefrin, S. and Statman, M. (1985): "The Disposition to Sell Winners Too Early and Ride Losers Too Long: Theory and Evidence," *Journal of Finance*, 40, 777–790.
- Stott, H. (2006): "Cumulative Prospect Theory's Functional Managerie," *Journal of Risk and Uncertainty*, 32, 101–130.
- Strack, P. and Viefers, P. (2021): "Too Proud to Stop: Regret in Dynamic Decisions," *Journal of the European Economic Association*, 19, 165–199.
- Tversky, A. and Kahneman, D. (1992): "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty*, 5, 297–323.