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

Yair Antler† and Ayala Arad‡

March 11, 2019

Abstract

We experimentally examine how individuals stop risky processes (e.g., the evolution of prices) when they have commitment power. Although the participants know the process' drift, they fail to assess the extent to which changes in their stopping decisions alter the probability that the process ends at a gain. Thus, they must trade off between the potential prizes and the winning probabilities in a qualitative manner. We find a type who consistently chooses stopping rules with large potential losses and small potential gains in order to induce a high winning probability, although such choices entail a large downside risk. We also find a less common type who chooses stopping rules with the opposite characteristics. We show that both of these patterns are inconsistent with leading theories of choice under risk.

Keywords: commitment, compound lotteries, downside risk, experiment, risky processes, stopping problems, types classification.

---

* We thank Dudi Dillenberger, Kfir Eliaz and Ariel Rubinstein for their comments and Tomer Baratz, Shira Bukchin, Eli Mograbi and Hadar Weisman for research assistance. We acknowledge a research grant from the Henry Crown Institute of Business Research in Israel, as well as the International Visiting Fellowship at the University of Essex.

† Coller School of Management, Tel Aviv University and CEPR. yair.an@gmail.com

‡ Coller 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 must decide when to stop the process. Well-known results from the theoretical literature on stopping problems have established that optimal stopping can be described by a simple *cutoff* rule, namely, stopping the process once the individual's payoff reaches a threshold.

Remarkably, and despite the optimality of cutoff rules and the fact that they are quite simple and easy to describe, individuals often exhibit behavior that is inconsistent with a cutoff rule. The dynamic nature of stopping problems triggers emotions such as regret, disappointment and elation, which may lead to dynamically inconsistent behavior and make it difficult for individuals to implement their preferred stopping plan. For example, in a recent lab experiment, Strack and Viefers (2018) documented history-dependent behavior that is consistent with regret aversion. Biases and departures from optimal stopping can lead to detrimental outcomes. For example, the negative feelings that are associated with realizing losses may lead investors to hold on to badly-performing stocks (Shefrin and Statman, 1985). Moreover, these biases make it difficult to infer individuals' preferences from their observed behavior and in turn the effect of these biases on their behavior.

A natural way to mitigate biases that arise during dynamic play is to commit to a stopping plan in advance or to delegate the execution to a third party (e.g., stop-loss and take-gain orders in stock trading). The ability to commit enables individuals to choose an optimal stopping plan without worrying about their ability to implement it. Thus, not only does understanding how individuals choose a stopping plan when they have commitment power help us to understand behavior in real-world stopping problems, it can also provide a true picture of the biases that arise in dynamic play, which is the first step in mitigating them. Our main research objective is to experimentally examine how individuals make stopping plans when they have commitment power.

In order to understand the setting and the tradeoffs that emerge in stopping problems, consider a decision maker (DM) who faces an infinite sequence of lotteries, where each lottery pays 1 with probability p and -1 with probability $1-p$. This binomial process can approximate the movement of real prices or the value of a stock. 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 his net payoff hits one of these bounds. The higher h is, the less likely the process is to reach h before it reaches l; correspondingly, 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 the probability of winning and the size of the potential gain/loss. How do individuals settle this tradeoff when the baseline lottery is fair (i.e., $p>0.5$), as is typically the case in the stock market? To what extent, if at all, does their behavior differ when the baseline lottery is unfair (i.e., $p < 0.5$), as in a casino setting?

In order to better understand the problem, consider the two cutoff rules presented in Figure 1. Under both rules, the sequence stops once the DM accumulates a net loss of 20. Under *a* (*b*) the sequence stops once the DM accumulates a gain of 10 (30). We refer to cutoff rules for which the upper bound is smaller (larger) in absolute terms than the lower bound as *left-biased* (*right-biased*). The likelihood that the sequence ends at 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 he is trading off potential gain against the probability of a gain.

Rule	Lower bound	Upper bound
a	-20	+10
b	-20	+30

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

In expectation, under the left-biased rule *a*, the DM participates in a smaller number of lotteries. Thus, when the baseline lottery is unfair (i.e., $p < 0.5$), a risk-neutral (or risk averse) expected utility maximizer would choose the left-biased rule. Our participants' choices match this prediction quite well. In problems in which we fixed the rules' potential loss and varied their potential gain, 63% of the participants consistently chose left-biased rules, whereas only 12% of them consistently chose right-biased rules.[1]

In a symmetric manner, when the baseline lottery has a positive expected value, a risk-neutral individual would choose the right-biased rule *b*. Indeed, 36% of the participants consistently chose right-biased rules in this case. Remarkably, even in this case, many of the participants (37%) consistently put a larger weight on the winning probabilities (or, according to their explanations, the probability of not losing) and chose left-biased rules (see Table 2).

Consider the analogous case of an individual who chooses between two cutoff rules that share the same *upper bound*, as in Figure 2. As a benchmark, note that a risk-neutral expected utility maximizer would prefer the left-biased rule *d* if the baseline lottery is fair ($p>0.5$) and the

[1] In each decision problem, the participants chose one rule out of five: two right-biased, two left-biased, and one symmetric.

right-biased rule *c* if it is not. Despite the larger potential loss, roughly half of the participants chose a left-biased rule in at least 5 out of 6 problems in which the upper cutoff was fixed: 52% when the baseline lotteries were unfair and 48% when the baseline lotteries were fair.

Rule	Lower bound	Upper bound
c	-10	+20
d	-30	+20

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

Our main finding is a general tendency to consistently choose either left-biased rules or right-biased rules. In all the types of problems (no fixed bounds, fixed lower bound, and fixed upper bound), the majority of the participants tended to choose left-biased rules and this tendency was stronger when the baseline lottery was unfair. The participants' explanations suggest that when $p<0.5$ they focus on minimizing the probability of finishing the game with a loss, whereas in the case of $p>0.5$ they put a larger weight on the size of the potential gain. The focus on the stopping rules' induced probabilities, even at the expense of greater potential losses (e.g., choosing rule *d* over rule *c*) is somewhat surprising since *the rules' induced odds were not explicitly given* in the experiment (rather only the baseline lottery's odds *p* were given), which should have made the gains and losses more salient than their respective probabilities.

At the individual level, the participants were consistent in their tendency to choose either left-biased or right-biased rules throughout the experiment and we were able to classify them into types. Participants who chose left-biased rules in a large proportion of the decision problems were classified as *L types* and those who chose a right-biased rule as *R types*.[2] Over 64% of the participants were classified as types, where the vast majority of them were *L* types.

The experimental design allows us to test whether or not the above findings can be explained by various theories of decision-making under risk, such as risk aversion and loss aversion. We classified the participants into types according to these theories using the same method mentioned above and found that solving the tradeoff between the size of the prizes and the winning probability in a consistent manner (i.e., consistently choosing either left-biased rules or right-biased rules) explains the data better than any of the other theories we examined. The participants' behavior,

[2] The proportion was defined such that the probability of being classified as a particular type under random choice was below 0.01.

together with their explanations, suggest that they try to solve a simple tradeoff between the likelihood of winning or losing and the size of the prizes and that the way in which the participants solve this tradeoff depends on the favorability of the baseline lottery and on whether or not the potential loss is fixed. In fact, some of the participants' choices were consistent with employing a lexicographic heuristic, with the criteria being the probability of winning and the size of the prizes.

Related literature

The present paper is related to a recent strand of the experimental literature that investigates dynamic stopping decisions. Oprea et al. (2009) examine irreversible investment decisions in the lab and find that, after some learning, individuals can approximate the optimal exercise of wait options. In Magnani (2017), participants are rewarded according to the higher between the current value of a state variable and a constant value, such that their optimal behavior is to follow a cutoff rule. At the aggregate level, the observed pattern of stopping behavior is consistent with the disposition effect (Shefrin and Statman, 1985) since individuals liquidate too early when they are ahead and too late when they are behind.

Fischbacher et al. (2017) show that stop-loss and take-gain strategies mitigate the disposition effect. The following aspects of their design make it considerably different from ours: (i) the lack of commitment power, (ii) the fact that each participant effectively chooses a stopping rule once, and (iii) the lack of knowledge of the baseline lottery's drift (i.e., p). Thus, while their design allows them to investigate the effect of stop-loss and take-gain strategies on an individual's tendency to hold on to losing assets, it cannot be used to determine the type of rules an individual chooses, the consistency of these choices, and how they depend on the favorability of the underlying process.

Strack and Viefers (2018) investigate dynamic stopping behaviour both theoretically and experimentally. Their participants dynamically choose when to stop a multiplicative random walk and exhibit history-dependent behavior that is consistent with regret aversion and inconsistent with cutoff rules. Sandri et al. (2010) found that even entrepreneurs tend to hold on to a badly-performing asset longer than is consistent with real options reasoning. Aloui and Fons-Rosen (2017) find that grittier individuals have a higher tendency to over-gamble relative to their original plans. In their experiment, the lotteries are unfair and most of the individuals choose to play even though they are not obliged to do so.

In none of the aforementioned papers, however, do the participants have commitment power. Moreover, the bulk of this strand of the literature focuses on dynamic play rather than planning (Aloui and Fons-Rosen, 2017 being the exception).

Stopping plans have been studied indirectly in the experimental literature on dynamic inconsistency, which focuses on *deviations* from planning (without commitment) 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 a single outcome. Cubitt and Sugden (2001) could not reject dynamic consistency when participants had to decide how many all-or-nothing additional gambles to participate in after winning in four mandatory rounds.

The present paper is also related 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 more positively skewed its induced lottery. 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 Rüger (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. A prominent explanation for skewness-seeking is probability distortion (Barberis, 2012; Ebert and Strack, 2015).

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 examine choices between lotteries with identical means and variance. In 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., given the two rules in Figure 1, a prudent individual might choose the left-biased rule when $p < 0.5$ since it induces a greater expected value and a smaller variance than the right-biased rule). Moreover, reducing a stopping rule to its induced lottery is a daunting task since the participants know only the probability of winning a single baseline lottery. (For example, Halevy, 2007, establishes that even in simpler settings individuals often fail to reduce compound lotteries.) Thus, our setting encourages reasoning in qualitative terms, which is less likely to be triggered in the typical setting in the literature on skewness-seeking and prudence.

The paper proceeds as follows. Section 2 presents the experimental design and Section 3 describes the results on both the aggregate and individual levels. In Section 4, we classify the participants into R and L types according to their choices. In Sections 5 and 6, we investigate the mechanisms that underlie our key findings and Section 7 concludes.

2. Experimental Design

The experiment was carried out in the Interactive Decision Making Lab at Tel Aviv University. The participants were 114 Tel Aviv University undergraduate students in various fields of study, 44% of whom were women. The average age was 25. Recruitment of participants was done via ORSEE (Greiner, 2004).

Each participant received 55 shekels (roughly $15) at the beginning of the experiment. In order for the participants to internalize this, we notified them one week prior to the session that they will receive this endowment (and that they may lose part of it or gain an additional amount, depending on their choices in the experiment). This message was sent again the day before the session. The experiment included 57 computerized decision problems (we often refer to these problems as questions), one of which was randomly selected at the end of the experiment to determine the payment for each participant. 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 45 additional shekels or could lose at most 28 shekels out of his initial endowment. All sessions were completed within an hour.

2.1 Detailed description of the experiment

The experiment consisted of two random treatments, denoted T_0 and T_p, and the questions were divided into four parts, as described below. Of the 114 participants, 67 were assigned to the main treatment, T_0, and 47 were assigned to T_p. The complete questionnaire can be found in Appendix B. Part A (Part B) examined the choice of a stopping rule when the baseline lottery has a negative (positive) expected value. Part C explored the participants' ability to estimate the rules' induced probabilities. Part D studied a slightly different setting in which the participants chose between pairs of binary lotteries that were characterized by a known probability of gain, an expected value of zero, and identical variance and kurtosis. Its main objective was to determine whether the participants exhibit a "pure" preference for skewness (and how that preference relates to the choices made in parts A and B).

Part A. In this part of the experiment, the participants faced a sequence of computerized lotteries, each with a probability of 18/37 to win 1 shekel and 19/37 to lose 1 shekel (henceforth, the baseline lottery). These probabilities resemble the win/loss probability in the "Red or Black" roulette game. In each decision problem, the participants were asked to choose a cutoff stopping rule (i.e., to choose two thresholds, one positive and one negative, such that once a participant's accumulated gain or loss hits one of the two thresholds the game would stop). The participants were presented with 18 decision problems and chose one out of five alternative stopping rules in each. In the case of the problem randomly selected for payment, 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 of the probability of ending the game with a gain for 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 examine the extent to which the choice patterns observed in T_0 were affected by the participants' knowledge of the rule's induced probability of winning in T_p.

We considered three types of decision problems, as illustrated in Figure 3. In Questions 1-6 (*fixed loss*) the participants chose from among five stopping rules with the same loss at which the process is stopped but that vary in the size of the gain. In Questions 7-12 (*fixed gain*) the participants chose from among five stopping rules with the same gain at which the process is stopped but that vary in the size of the loss. In Questions 13-18 (*not fixed*) both the gain and loss vary.

In each decision problem, there are two stopping rules in which the potential loss is greater than the potential gain, two in which the potential gain is greater than the potential loss, and one in which they are equal. We refer to them as left-biased, right-biased, and symmetric, respectively. The five stopping rules were randomly ordered either 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 vice versa.[3] Thus, the five stopping rules were either always ordered from the highest to the lowest probability of a gain or the other way around.

[3] 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 merge the data from the two variations in the analysis.

Type (i)

lottery	loss	gain	probability of gain
a	-21	+9	52%
b	-21	+15	35%
c	-21	+21	24%
d	-21	+27	17%
e	-21	+33	12%

Type (ii)

lottery	loss	gain	probability of gain
a	-20	+12	42%
b	-16	+12	39%
c	-12	+12	34%
d	-8	+12	28%
e	-4	+12	18%

Type (iii)

lottery	loss	gain	probability of gain
a	-27	+15	38%
b	-24	+18	31%
c	-21	+21	24%
d	-18	+24	19%
e	-15	+27	14%

Figure 3. The three types of questions in Part A. The probability of a gain given each stopping rule is provided for the readers' convenience; only participants in T_p received information on the probability of a gain or a loss given the stopping rule. In particular, the following sentence: "The probability that the process ends at a gain is *q*% and the probability that it ends at a loss is (100-*q*)%" was presented below the description of the rule (see Appendix B).

Part B. This part is similar in structure to Part A, in that there are 18 problems of the same three types described above. The main difference is that the probabilities of gain and loss in the baseline lottery are reversed in Part B (i.e., the probability of winning in a single lottery is 19/37). In addition, we tried to diversify the problems in Part A and B in order to prevent a feeling of repetition. Thus, the stopping rules in Part B are similar to those in Part A, though not identical.

Part C. Part C consists of three problems, each of which presents a different stopping rule. In all three problems, the participants were asked to consider a baseline lottery that pays 1 shekel with probability 18/37 and -1 shekel with probability 19/37 (as in Part A) and to estimate the probability that the game will end at a gain given the stopping rule. In particular, in Problem 1, they had to gauge the probability of finishing the game with a gain of 25 given that the stopping rule is (-25, +25). Problems 2 and 3 were similar except that the stopping rules were (-25, +50) and (-25, +100), respectively. The correct answers to these three questions are 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 Part C.

Part D. In Part D, the participants faced 18 decision problems. In each, they chose between two binary lotteries with known probabilities of loss and gain (as illustrated in Figure 4). The two lotteries were "mirror images" 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 they had an expected value of 0, identical variance, and identical kurtosis. The departure from the context of a stopping problem, as well as the fact that the lotteries had the same expected value and variance, allowed us to test whether some individuals have a "pure" taste for skewness. Such preferences may be relevant to the choice of a stopping rule in Parts A and B. In each question, the order of appearance of the two lotteries was randomly and independently determined. There was no difference between the two treatments in Part D.

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

Part D – Game 2

Which of the following two lotteries do you prefer?

a.

chance	24%	76%
amount	-25	+8

b.

chance	76%	24%
amount	-8	+25

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

Discussion: Choosing from a fixed set 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 allowed them to select the stopping rule's upper and lower bounds from an unconstrained set. We decided to constrain their choice so as to focus on the effect of the qualitative properties of the stopping rules (e.g., the effect of the fairness of the baseline lotteries, how the choices differ given fixed loss/gain, etc.) while keeping the participants' decision problems relatively simple. Moreover, the "constrained" problems differ from one another and thus allow us to examine whether participants use rules that share similar properties in a large number of different problems. We believe that it would be significantly more difficult to have the participants perceive numerous unconstrained problems as differing from one another, especially since the stopping rules' induced probabilities are not specified in our main treatment.

2.2 Theoretical predictions

Our analysis suggests that individuals solve the tradeoff between prizes and probabilities in a consistent manner when they choose a stopping rule. That is, they consistently choose either left-biased rules, which minimize the loss probability, or right-biased rules, which minimize the size of the potential loss and maximize the size of the potential gain.

We considered several decision procedures to see whether they can explain the findings. To obtain some intuition for what they would predict, consider a risk-neutral expected utility maximizer. In Part A, where the baseline lotteries are unfair, such a decision maker would minimize the induced number of baseline lotteries in which he participates. Thus, in type (i) (type (ii)) problems, where the potential loss (gain) is fixed, such a decision maker would choose the most left-biased (right-biased) rule. In Part B, where $p>0.5$, he would exhibit the opposite behavior, namely he would maximize the expected number of lotteries in which he participates. In type (iii) problems, the stopping rules' induced expected value is U-shaped in their right-biasedness in Part A and hump-shaped in Part B.

Modest levels of risk aversion do not change the predictions. For example, a risk averse individual with CRRA utility would exhibit the same pattern of choice for any wealth level $w \in \{500,600,\ldots,10000\}$ and risk aversion coefficient $\sigma \in \{0,0.1,0.2,\ldots,2\}$. A lower wealth level (e.g., $w = 55$, which reflects a mental accounting procedure, according to which the participants code the experiment's earnings in a separate mental account) induces greater absolute risk aversion in the case of CRRA utility and leads to different predictions (mainly) in type (iii) problems.[4]

In Part C of the experiment, we found that many of the participants estimate the rules' induced winning probabilities as if the baseline lottery is almost fair (i.e., as if *p* is closer to 0.5 than it actually is). To examine whether this type of bias can explain our findings, we compared the participants' choices to those of a risk-averse expected utility maximizer when the baseline lottery's winning probability is 0.5. Due to risk aversion, such an individual would try to minimize the number of lotteries in which he participates. Thus, he would choose left-biased rules in type (i) problems and right-biased rules in type (ii) problems. The choices of such an individual in type (iii) problems will depend on the specific parameters.

In Section 4, we examine which of the following decision procedures better explains the data:

1. Consistently solving the prize-probability tradeoff in favor of rules that induce larger/smaller winning probabilities.
2. Expected utility maximization (with modest levels of risk aversion).
3. Loss aversion with diminishing sensitivity to gains/losses.
4. Risk aversion with mental accounting.
5. Risk aversion with distorted beliefs (i.e., estimating the rules' induced probabilities as if *p*=0.5).

Procedures 2, 3 and 4 lead to the same prediction in type (i) and type (ii) problems. Type (i) problems in Part A and type (ii) problems in Part B distinguish between participants who behave according to these procedures and those who consistently solve the prize-probability tradeoff in favor of a large prize or a small loss (i.e., right-biasedness). Type (ii) problems in Part A and type (i) problems in Part B distinguish between participants who behave according to these procedures and those who consistently solve the prize-probability tradeoff in favor of a large winning probability (i.e.,

[4] In a similar manner, modest levels of probability distortion (in the spirit of cumulative prospect theory, see Kahneman and Tversky, 1992) do not change the prediction by much.

left-biasedness). Type (i) and type (ii) problems in Part A distinguish between Procedure 5 and right-biasedness, and type (i) and type (ii) problems in Part B distinguish between Procedure 5 and left-biasedness. Type (iii) problems distinguish between the five decision procedures; however, the differences are not as sharp as they are in type (i) and type (ii) problems.

3. Choice of stopping rules (Parts A and B)

We shall 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 present the results obtained in T_p, in which the rules' induced probabilities were provided, and compare them to the results obtained in T_0.

Recall that in each problem, the participants chose between five stopping rules: two left-biased rules, one symmetric rule, and two right-biased rules. We refer to the choice of the rule with the largest loss and the smallest gain as *Answer 1*; the second-most left-biased rule as *Answer* 2; the symmetric rule as *Answer* 3; the rule with the largest gain and smallest loss as *Answer 5*, and the second-most right-biased rule as *Answer 4* (though some participants were presented with the reverse order). The behavior in T_0 provides indications of two general patterns. First, in Part A, where the baseline lottery is unfair, there is a strong tendency to choose *left-biased* stopping rules (i.e., Answers 1 and 2). Second, in Part B, this tendency is weaker relative to Part A. These patterns are observed at the aggregate level, at the individual level, and in all three types of problems.

Aggregate-level analysis

At the aggregate level, 1206 (67x18) choices were made in each part of the experiment. 66% of the chosen rules in Part A were left-biased while only 25% were right-biased. In Part B, 46% of the chosen rules were left-biased while only 35% were right-biased (see Table 1).

	Part A	**Part B**
Answer 1	31%	23%
Answer 2	35%	23%
Answer 3	9%	19%
Answer 4	10%	19%
Answer 5	15%	16%

Table 1. The distribution of choices in T_0 (total of 1206 choices).

Individual-level analysis

We use two measures to examine the participants' choices at the individual level. The first is the number of times each individual chose a left-biased rule, which ranges from 0 to 18 for each part of the experiment, and similarly for a right-biased rule. We refer to these measures as the *number of left-biased choices* and the *number of right-biased choices*, respectively. The second type of measure attaches a weight of i to Answer i (in every problem) and sums up these weights such that the measure ranges from 18 to 90 for each individual in each part of the experiment. The closer the measure is to 18 the greater the participant's tendency to solve the prize-probability tradeoff in favor of left-biased rules and vice versa when the measure is closer to 90. We refer to this measure as the *total score*.

The number of left-biased choices is higher on average in Part A than in Part B (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 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 stochastically dominates the corresponding distribution in Part A. Accordingly, the average total score per individual is higher in Part B than in Part A (50.76 vs. 43.76, $t(66)=-4.1$, $p<0.001$).

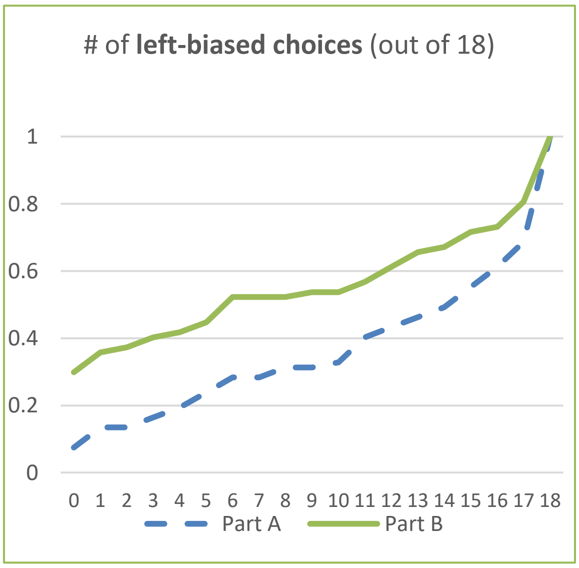

Figure 5a. Cumulative distribution of the number of left-biased choices per participant: Part A vs. Part B.

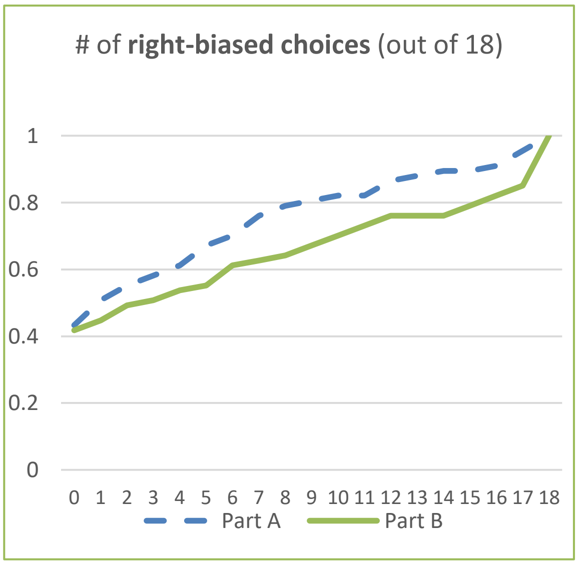

Figure 5b. Cumulative distribution of the number of right-biased choices per participant: Part A vs. Part B.

Note that, despite the significant differences in behavior in Part A and B, the individuals' choices in these two parts are correlated according to the number of left-biased choices measure (Pearson's r=0.56, p<0.001), according to the number of right-biased choices measure (Pearson's r=0.64, p<0.001) and according to the total score measure (Pearson's r=0.77, p<0.001). Viewed together, these findings suggest that an individual has a tendency towards a particular direction (i.e., to choose either left-biased rules or right-biased rules). The favorability of the baseline lottery's odds reduces the common tendency towards left-biased rules, presumably because it moderates the focus on minimizing the probability of losing.[5]

Extreme vs. moderate choices. In both Part A and Part B, the difference between the number of times the individuals chose Answer 1 and the number of times they chose Answer 2, as well as between the number of times they chose Answer 4 and the number of times they chose Answer 5, are small and insignificant.

Analysis of the three types of problems

Examining each of the 36 problems in Parts A and B separately suggests that left-biased choices are more prevalent than right-biased ones in all but two of the problems. In Part A, the median choice was 2 in all of the 18 problems and the average choice was in the range of 2.06-2.87. In Part B, the median choice in most problems was 3 (and 2 in the rest) and the average choice was in the range 2.54-3.13. The range of the average choice suggests that there are some differences between the problems in the extent that left-biased rules were chosen. Therefore, we now examine how the type of problem affects the tendency to choose left-biased rules. Table 2 presents this tendency in each of the three types of problems in Parts A and B (1-6: fixed loss, 7-12: fixed gain and 13-18: not fixed) and compares it to the tendency of choosing right-biased rules.

The results show that left-biased choices are more common in Part A than in Part B, regardless of the type of problem. Furthermore, we find that the tendency toward left-biased choices in problems 1-6 of Part A is greater than that in questions 7-12 of Part A (the average number of left-

[5] The participants' explanations provide some indication that in Part B they shift their attention to the potential gains rather than the potential losses. For example, keywords were classified into the following categories: probability of winning, probability of losing, gain size and loss size. Counting the use of these keyword categories in the participants' explanations suggests that the ratio of the probability of winning to that of losing increases in Part B (54:29 in A and 64:14 in B) as does the size of the gain vs. the size of the loss (66:65 in A and 54:39 in B).

biased choices is 4.4 vs. 3.67, $t(66)=2.78$, $p=0.007$), while the opposite pattern is observed in Part B (2.58 vs. 3.16, $t(66)=-1.95$, $p=0.055$).

	Part A			**Part B**		
	fixed loss	*fixed gain*	*not fixed*	*fixed loss*	*fixed gain*	*not fixed*
Proportion of participants choosing left-biased rules (range in the 6 questions)	66%-79%	55%-64%	55%-69%	34%-51%	48%-54%	42%-49%
Proportion of participants choosing at least 5 left-biased rules (out of 6)	63%	52%	60%	37%	48%	37%
Proportion of participants choosing at least 5 right-biased rules (out of 6)	12%	24%	19%	36%	28%	27%

Table 2. Each type of problem consists of 6 questions. The first row presents, for each type of problem, the maximal and minimal proportion of left-biased choices per question. The second (third) row presents the proportion of participants who tend to choose left-biased (right-biased) rules for each type of problem in Part A and B of T_0.

In Part A, left-biased rules are commonly chosen in the fixed loss problems. Intuitively, since the participants cannot control the size of the loss, they focus on minimizing its likelihood. In the fixed gain problems, participants can also control the size of the potential loss and there is a tradeoff between a higher probability of winning and a smaller potential loss. Thus, a smaller proportion of participants chose the left-biased rules. Nonetheless, *the majority of participants chose left-biased rules in all types of problems.*

In contrast, in Part B, where the odds are on the participants' side, it seems that the participants' focus shifts towards the potential gain. When participants face the fixed gain problems, they focus on maximizing the probability of winning (since they can't control the size of the gain) by choosing left-biased rules. In the fixed loss problems, they can also control the size of the potential gain and hence opt for right-biased rules more often.

4. Classification into Types

We can now classify the participants into types according to their choices. A participant is type L if he consistently chooses left-biased rules and type R if he consistently chooses right-biased rules. We now present two definitions according to which we will classify the participants (although the results under both are very similar). In order to account for the possibility that individuals solve the prize-probability tradeoff in a different manner when the favorability of the baseline lottery changes, we classify the participants into types according to their choices in Part A (18 choices), their choices in Part B (18 choices), and their choices in Parts A and B together (36 choices).[6]

4.1 Two definitions of type

Definition 1: The number of left-/right-biased choices

For Part A and Part B separately, a participant is defined as type L (R) if he chooses a left-biased (right-biased) rule in 13 or more out of 18 problems.[7] Participants who are neither L nor R are referred to as "other". If a participant chooses rules randomly, the probability of being classified as one of the two types is roughly 0.012 (i.e., 0.006 for each type). Consistent with these probabilities, when examining the choices in Part A and B together, we classify a participant as an L type (R type) if he chooses left-biased rules (right-biased rules) in at least 22 out of 36 problems.

Figure 6 and 7 show the experimental distribution of the number of left-biased and right-biased choices (per participant) in Parts A and B, compared to the probability of such realizations given a random choice. Figure 8 displays these distributions for Parts A and B combined.

Definition 2: The total score

In each of Part A and Part B, a participant is classified as an L type if the sum of his answers is in the range [18,40], an R-type if it is in the range [68,90], and "other" if it is in the range [41,67]. If a participant chooses randomly, the probability of the sum being in the low (high) range is roughly 0.01. Combining Part A and B, a participant is classified as an R-type if his total score is in the range [36,88] and as an L type if his total score is in the range [128,180]. The probability of being classified as each of these types for a participant that chooses uniformly at random is roughly 0.01.

[6] In Appendix A, we provide an alternative definition of types which identifies participants that are *relatively extreme* in their choices and report the proportions of such types in our experiment.

[7] This definition is in the spirit of Ebert and Wiesen (2011)'s concept of skewness-seeking.

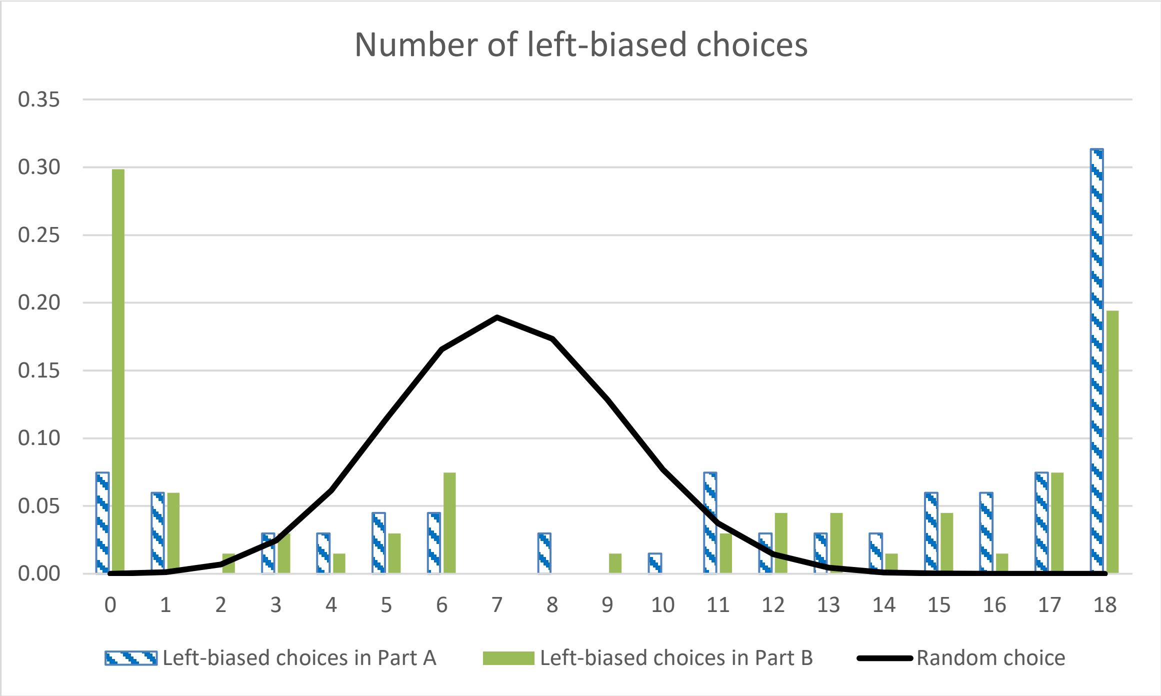

Figure 6. The distribution of participants according to their number of left-biased choices in Part A and B and the probability of each number of choices given random uniform choice in each problem.

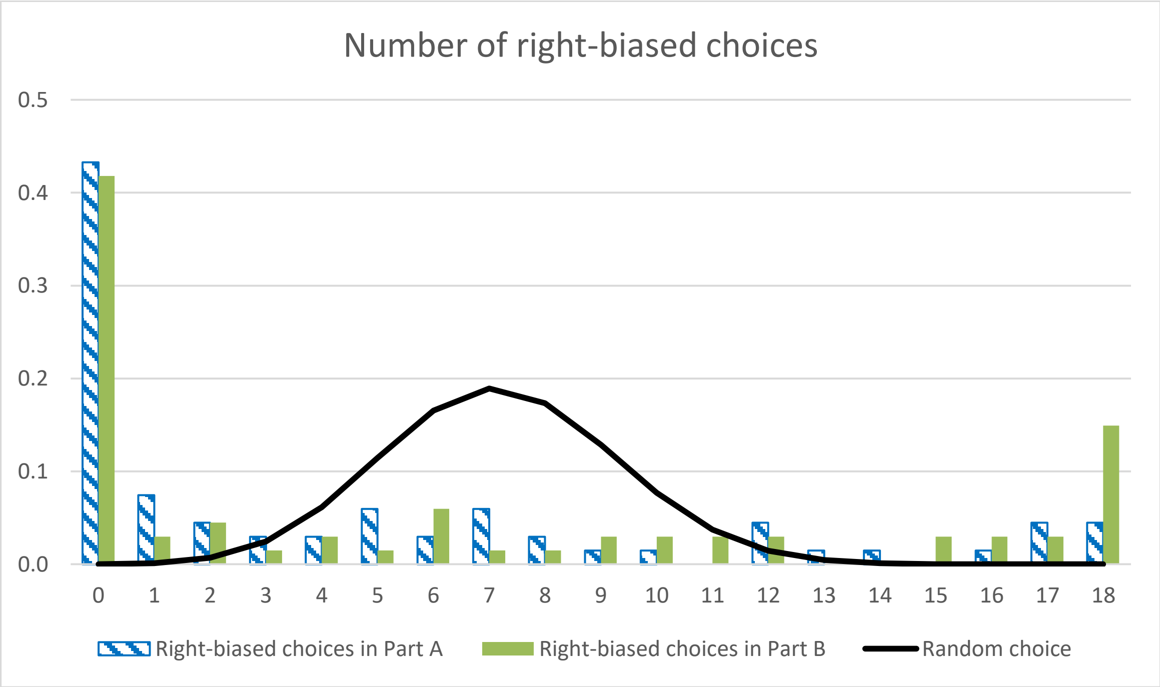

Figure 7. The distribution of participants according to their number of right-biased choices in Part A and B and the probability of each number of choices given random uniform choice in each problem.

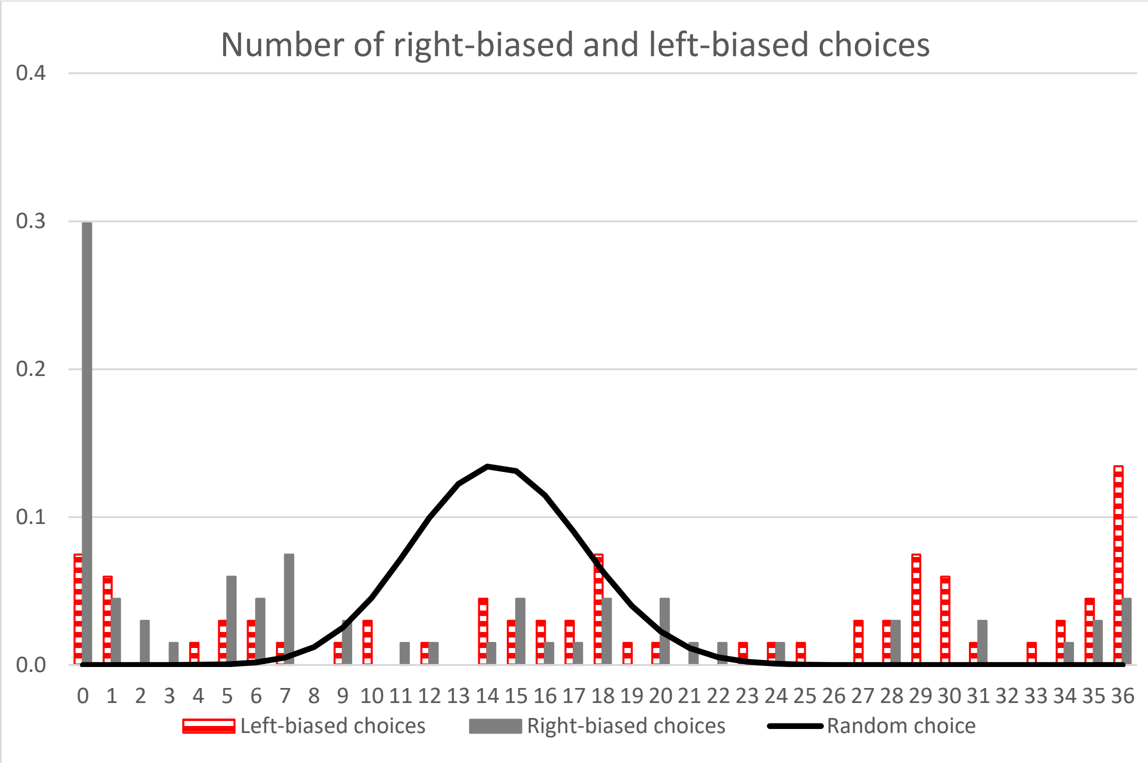

Figure 8. The distribution of participants according to their number of right-biased and left-biased choices in Part A and B combined and the probability of each number of choices given random uniform choice in each problem.

Comment: Internal consistency

The answers to the 18 problems in each part are highly correlated. Cronbach's alpha is 0.95 in Part A and 0.96 in Part B. We further examined the in-sample predictive power of our classification (in each part separately) according to each of the above definitions and concluded that it is possible to use the choices in each subset of 17 problems to classify the participants into the three groups: L, R and Other, and to predict their choices in the 18th problem based on that classification (in the spirit of the typology test suggested in Rubinstein (2016); see Figures S1-S4 in Appendix A).

4.2 Proportion of types

Table 3 presents the proportion of types in T_0, according to Definitions 1 and 2. In both Part A and Part B, the types account for a majority of the participants (62%-72%). The proportion of L types in Part A is higher than in Part B whereas the proportion of R types is lower than in Part B (according to a McNemar test, p=0.023 and p=0.039, respectively). Nonetheless, a participant's type in Part A is correlated with his type in Part B (Spearman's rho =0.5, p<0.001). In particular, a type L in Part A is unlikely (only 13%) to be classified as type R in Part B and vice versa for type R in Part A (none of the R types were classified as type L in Part B); some types switch to "other" in Part B. Considering all of the 36 choices together, we classified 47.7% of the participants as L types and 19.5% as R types according to Definition 1. According to Definition 2, 49.2% of the participants were classified as L types and 17.9% as R-types. Thus, according to each of the six methods of classification, about 2/3 of the participants are classified as types, with the majority of them classified as L types.

		Types in Part A		**Types in Part B**	
		L	*R*	*L*	*R*
T_0 *N=67*	*Number of choices*	57% (38)	13% (9)	39% (26)	24% (16)
	Total score	57% (38)	15% (10)	37% (25)	25% (17)

Table 3: The proportions of types in Part A and B in T_0, according to Definitions 1 and 2.

Consistency with the participants' explanations

At the end of each part, the participants were asked (i) how they would instruct someone else to play on their behalf and (ii) to describe their main considerations in making their choices. To support our interpretation of the type L and type R decision procedure, two research assistants examined whether the explanations fit our identification of types (by Definition 1) using the following protocol: Each assistant separately examined the explanations and classified them into four categories: (1) Explicit – the explanation included an explicit description of choosing (consistently) right- or left-biased stopping rules; (2) Partial – the explanation included the description of a general rule of right or left bias but mentioned exceptions in which the tradeoff was resolved differently; (3) Inconsistent – the explanation contradicted the type identification; and (4) NA – the text did not provide enough information to determine whether the participant's identification is consistent with the reasoning he describes. Following that, the assistants discussed with each other the cases in which their classifications disagreed (it turned out that their classifications were almost always in agreement).

In only a small number of cases did the assistants' final classification indicate that participants described a decision procedure that was different from the type that had been assigned to them. In fact, the vast majority of participants included an explicit description of right or left tendency. In particular, in Part A, the explanation of only *one* of the 69 participants who were classified as types (in the two treatments combined) was inconsistent with his type classification. In addition, the explanations of 6 of those 69 participants were classified as NA. Similarly, in Part B, the explanations of *two* of the 63 participants who were classified as types were inconsistent with their type classification. The explanations of 8 of those 63 participants were classified as NA.

The following are examples of participants' explanations in T_0 in each of the four categories:

Participant 57, Part A (type L) – Explicit: *In each problem, I chose the option in which the loss is larger than the gain so that the chances of achieving a gain would be greater. Nevertheless, I chose the option in which the gap between the loss and the gain is not relatively large. For example, a gain of 6 and a loss of 15.*

Participant 45, Part A (type L) – Partial: *The probability that you win increases if the negative stopping rule is high and the positive stopping rule is low… you need to play safe but not too safe… because there are cases in which the conditions do not pay off.*

Participant 89, Part A (type L) – Inconsistent: *Trying to obtain maximal gain given that a loss is most likely, even if it means risking a loss.*

Participant 4, Part A (type L) – NA: *Choose according to the probability of winning relative to the probability of losing.*

Participant 54, Part B, (type L) – Explicit: *The instructions are almost identical (to Part A). The chance of winning is a bit higher than in the previous part and hence I took a chance, such that the stopping rules of the gain and of the loss are more similar but still the loss is further away from 0 in order to increase the chance of winning.*

Participant 65, Part B, (type R) – Explicit: *The instructions are similar to those of Part A, but just the opposite since the probability of winning is higher. Thus, I looked for the smallest bias towards the "plus".*

In conclusion, many of the participants' explanations focused on a qualitative resolution of the tradeoff between the probability of a gain/loss and the size of the potential gain/loss. Most of the participants who were classified as types (either L or R) explicitly explained their decision as consistently choosing what we refer to as left/right-biased stopping rules.

4.3 Alternative (theory-based) explanations

4.3.1 Explaining the behavior in T_0: A horse race

It is worthwhile considering whether the participants' choices can be explained by any of the theories described in Section 2.2. We start by classifying the participants into types according to each of these theories. For each theory and each problem, we sort the five rules according to the value that they induce according to the theory and define a participant's choice as consistent with the theory if it matches one of the top two rules. We classify a participant as a type in Part A (and similarly in Part B) according to a particular theory if his choices are consistent with that theory in at least 12 of the 18 problems (if a participant chooses uniformly at random, the probability of observing such a pattern is about 2%, which is slightly larger than the combined probability of being classified as an R type or an L type according to Definition 1). As a benchmark, recall that 47 of the 67 participants were classified as an L or R type in Part A and 42 participants were classified as an L or R type in Part B.

We found that the choices of only a small number of participants matched those predicted by expected utility theory under risk neutrality or modest levels of risk aversion (as described in Section 2.2). Only 14 (13) participants made choices that are consistent with this theory in Part A (B). Loss aversion with a piece-wise linear value function and a loss aversion coefficient of $\lambda<3$ leads to a similar outcome. Loss aversion with diminishing sensitivity leads to a slightly better match. For example, for $\lambda=2.2$ and a diminishing sensitivity parameter of $\alpha = 0.83$, the theory produces a match for 14 (23) participants in Part A (B).[8]

We also examined a mental accounting procedure, in which the participants code their earnings in the experiment into a separate account. To capture this, we considered log-utility maximization with an initial endowment of 55 NIS (the number of participants who are classified is not very sensitive to changes in the level of risk aversion). We found that the behavior of 15 (13) participants in Part A (Part B) matches this explanation. We also considered a theory according to which the participants treat the baseline lottery as if its winning probability is 0.5. With logarithmic utility and an initial endowment of 55 NIS this theory explains the behavior of 17 (9) participants in Part A (Part B).[9]

[8] The number of matches is not very sensitive to the parameters, which were chosen based on the calibration exercise in Kahneman and Tversky (1992).

[9] As explained in Section 2.2, the predictions in problems of types (i) and (ii) do not depend on the level of risk aversion and therefore the classification is not very sensitive to the specific utility function/endowment.

If we consider all 36 problems together, the cutoff for classification is 21 matches such that the probability of being randomly classified as a type is roughly 2%. In this case, we were able to classify 17 of the participants as loss averse types and no more than 12 participants for each of the other theories. In comparison, 33 of the participants were classified as L types and 12 as R types.

In conclusion, the alternative theories cannot explain the behavior of most of the participants. It is still nonetheless possible that some of the theories explain the behavior of some of the participants that were classified as R or L types. Therefore, we now modify our initial classification by eliminating participants whose behavior is better explained by one of the alternative theories.

4.3.2 A more conservative classification

In this subsection, we consider the participants who were classified as R types or L types according to the number of times they chose a left- (right-) biased rule, and examine whether *their* behavior can be explained by one of the alternative theories. We change a participant's classification (from L or R) to *other* if his behavior is better explained by one of the alternative theories or if left-/right-biasedness is only slightly better in explaining his behavior than an alternative theory. In particular, if a participant is classified according to an alternative theory and the number of matches according to this alternative theory is lower by 2 or less than the number of matches to his initial type classification (L or R), then we modified his initial classification to *other*.[10]

The conservative classification results are as follows: in Part A, 37 (out of 38) participants remained L and 7 (out of 9) remained R; in Part B, 24 participants (out of 26) remained L and 16 participants (out of 16) remained R. Classifying types based on their combined behavior in Parts A and B, 30 participants (out of 33) remained L and 10 participants (out of 12) remained R. Thus, even a more conservative approach does not change the classification all that much.

5. The effects of probability (mis-)estimation

In this section, we examine to what extent the choices of the participants in the main treatment were affected by the fact that they were not provided with the rules' induced winning probabilities. Not knowing the induced probabilities should not have any effect if the participants can infer them from the likelihood of winning a single baseline lottery. Therefore, the first step of the analysis must be to examine the participants' ability to make such an inference. Part C of the experiment explores this

[10] Changing the criterion to a gap of 4 would not make much difference; only one participant's classification would change.

question and establishes that the participants' inferences are in fact quite different from the true winning probabilities (which is consistent with the finding that individuals typically fail to reduce compound lotteries; see, e.g., Halevy, 2007). In the second part of this section, we present the results of a second treatment in which the induced probabilities were explicitly given to the participants in Parts A and B. Analyzing the behavior of these participants sheds light on the effects of the unknown probabilities on the participants' behavior.

5.1 Are the participants able to infer the rules' induced probabilities? (Part C)

In each of the three problems in Part C, the participants faced a single stopping rule. The rules were (-25,+25), (-25,+50), and (-25,+100) in the first, second, and third problem, 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 are 20.6%, 5.1%, and 0.3%, respectively.

The participants' average estimates in T_0 were 39.7%, 24.2%, and 17.2%. The expected errors in absolute terms were 22.9%, 14.1%, and 17.9%. Moreover, only 26.8% of the answers were within a range of 5% from the correct answer.[11] While most of the participants failed to estimate the winning probabilities correctly, they did exhibit a qualitative understanding of the prize-probability tradeoff, such that 86.8% of them provided monotonic estimates (an estimate is monotonic 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)).

These results indicate that the vast majority of the participants overestimate the likelihood of finishing the game at a gain when the baseline lottery is only slightly unfair.[12] Such overestimation can lead to overconfidence and excessive gambling, since casinos and online gambling websites typically offer unfair baseline gambles which resemble the baseline lotteries in our experiment (e.g., the winning probability in single Blackjack and a Red or Black roulette game is close to 0.5).

After establishing that the vast majority of participants failed to estimate the induced winning probability, we turn to investigate a treatment in which the participants were provided with the rules'

[11] In T_p, where the participants knew the probabilities of the stopping rules in Part A and B, the average estimates in Part C were: 36.8%, 22.8%, and 12.4%, and the average error size was reduced somewhat in all three questions, though the reduction was marginally significant only in the first problem (error difference=3.83, $t(112)=1.89$, $p=0.062$). Only 29.8% of the answers in T_p were within a range of 5% from the correct answer.

[12] While it is unclear whether this overestimation bias persists when the baseline lotteries' expected value is far from zero, this question is out of the present paper's scope and we leave it for future research.

induced probabilities in Part A and B and compare their behavior to that observed in the main treatment.

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

We begin with a description of the behavior in T_p, in which the participants were provided with the stopping rules' induced probabilities of winning and losing. Following that, we will compare the two treatments.

At the aggregate level, the behavior patterns exhibited by the T_p participants are similar to those observed in T_0. Thus, when the baseline lottery is unfair, 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, left-biased rules were chosen in 62% of the 846 (47x18) choices and right-biased rules in 28%, whereas in Part B the figures were 49% and 37%, respectively. It should be noted that the number of left-biased choices in Part A is greater than in Part B and the number of right-biased choices in Part A is smaller than in Part B (see Table 4).

	Part A	**Part B**
Answer 1	44%	32%
Answer 2	18%	17%
Answer 3	11%	15%
Answer 4	10%	13%
Answer 5	18%	24%

Table 4. The distribution of choices in T_p (total of 846 choices).

At the individual level, the mean number of choices of left-biased rules in Part A was higher than in Part B (11.09 vs. 8.74, $t(46)=2.96$, $p=0.005$). Similarly, the total score, which could potentially range from 18 to 90, is higher in Part B than in Part A (on average, 50.40 vs. 43.23, $t(46)=-2.86$, $p=0.005$). Nonetheless, the participants' choices in Part A and Part B are correlated according to the number of left-biased choices measure (Pearson's $r=0.62$, $p<0.001$) and according to the total score measure (Pearson's $r=0.61$, $p<0.001$).

Extreme vs. moderate choices. In contrast to T_0, the differences between the number of choices in the same direction (i.e., Answer 1 vs. Answer 2 and Answer 5 vs. Answer 4) are significant. In particular, participants tended to choose the extreme stopping rules more often. Thus, Answer 1 was chosen more than Answer 2 (in Part A the difference is 4.7, $t(46)=4.14$, $p<0.001$, whereas in Part B the difference is 2.7, $t(46)=2.97$, $p=0.005$) and Answer 5 was chosen more than Answer 4 (in Part A the difference is 1.49, $t(46)=-2.13$, $p=0.039$, whereas in Part B the difference is 1.96, $t(46)=-2.99$, $p=0.005$). Thus, the uncertainty regarding the induced lotteries in T_0 mitigated the individuals' extreme choices.

Table 5 presents two measures of the participants' tendency to choose left-biased rules in each of the three types of problems and compares it to the corresponding tendency to choose right-biased rules. The results suggest that left-biased choices are common in T_p as well. Furthermore, we observed a "mirror" pattern similar to that observed in T_0, such that the tendency of left-biased choices in fixed-loss problems is greater than in fixed-gain problems in Part A (the average number of left-biased choices is 4.87 vs. 3.02, $t(46)=4.8$, $p<0.001$), while the reverse pattern is observed in Part B (1.96 vs. 3.94, $t(46)=-5.09$, $p<0.001$).

	Part A			**Part B**		
	fixed loss	*fixed gain*	*not fixed*	*fixed loss*	*fixed gain*	*not fixed*
Proportion of participants choosing left-biased rules (range in the 6 questions)	68%-87%	43%-57%	47%-55%	23%-45%	57%-75%	40%-57%
Proportion of participants choosing at least 5 left-biased rules (out of 6)	77%	40%	45%	23%	64%	38%
Proportion of participants choosing at least 5 right-biased rules (out of 6)	2%	26%	23%	47%	15%	15%

Table 5: Each type of problem consists of 6 questions. The first row presents, for each type of problem, the maximal and minimal proportion of left-biased choices per question. The second (third) row presents the proportion of participants who tend to choose left-biased (right-biased) rules for each type of problem in Part A and B of T_p.

The above observations suggest that the patterns of behavior in T_p were similar to those in T_0. A comparison of the two treatments indicates that neither of the two measures (number of left-biased choices and total score) differs significantly between the treatments. Furthermore, there are no significant differences between the treatments in the number of left-biased choices for any of the six types of questions.[13]

Type Classification. In both parts, the proportion of participants classified as types is slightly higher in T_0 than in T_p (the differences are highly significant when using the conservative measure and significant at the 10% level when using the two other definitions). Nevertheless, resolving the prize-probability tradeoff in a consistent manner in terms of right-/left-biased rules appears to explain the behavior in T_p better than the alternative theories discussed in Section 2.2. Of those theories, loss aversion with diminishing sensitivity to gains/losses best explains the participants' behavior (the complete analysis appears in Appendix A).

In sum, although the behavior in T_p is not identical to that in T_0, the observed patterns are quite similar. Thus, knowing the stopping rules' induced probabilities had only a minor effect on the participants' behavior.

6. Pure preference for skewness

Our notion of biasedness corresponds to the skewness of the stopping rules' induced lotteries, i.e. the more right-biased a rule is, the greater is the skewness of its induced lottery.[14] Thus, a pure taste for positively or negatively skewed prospects might affect a participant's preferences over stopping rules. In Part D, we examined whether the participants have such a taste for either positive or negative skewness.

The participants were presented with 18 problems in which they chose between a pair of binary lotteries. In each pair, the two lotteries had an expected value of 0, the same variance, and the same kurtosis. The key difference between each pair of lotteries was that one was positively skewed and the other was negatively skewed. The prizes were chosen such that (given a baseline winning probability p=0.5) the positively skewed lottery could be induced by a right-biased rule and the negatively skewed lottery by a left-biased rule. The participants' decisions in this part of the experiment are simpler than those in Part A and B in two main dimensions: the winning probabilities are given and the lotteries are not presented as stopping rules.

[13] The only significant difference in behavior between the treatments is the abovementioned tendency to choose more extreme stopping rules (i.e., choices 1 and 5 are more common than 2 and 4). This decreases the sum of choices in problems 1-6 (fixed loss) of Part A and problems 7-12 (fixed gain) of Part B and increases the sum of choices in problems 1-6 (fixed loss) of Part B (see Table S1 in Appendix A).

[14] It should be noted that a left-biased rule may induce a positively skewed lottery when p<0.5 and a right-biased rule may induce a negatively skewed lottery when p>0.5.

On an aggregate level, 49% of the choices in Part D were of negatively skewed lotteries. In almost all of the 18 problems, the distribution of choices was relatively balanced, i.e. between 40% and 60% choices of negatively skewed lotteries, where the highest proportion was 71% (which occurred in the first problem in Part D). On an individual level, the number of choices of negatively skewed lotteries (which could potentially range from 0 to 18) is, on average, 8.75 with a median of 8. Among the 114 participants in the two treatments, we could classify 31% of the participants as L types and 38% as R types using Definition 1. Despite the slightly different choice pattern, there is a significant correlation between the individuals' total score in Part D and their total score in Part A and B of T_0 (Pearson's r=0.27, p=0.025 and Pearson's r=0.35, p=0.003, respectively).[15] The behavior in Part D is also correlated with the number of left- and right-biased choices in Part A and B .

In conclusion, at the aggregate level, the participants did not exhibit a taste for negatively skewed prospects to an extent that could explain the general tendency to choose left-biased rules in Parts A and B. Nonetheless, the significant correlation suggests that pure preference for negative or positive skewness is related to the individual tendency to choose left- or right-biased stopping rules.

Comment: the literature on skewness-seeking and prudence

The literature on skewness-seeking and prudence documents a taste for positively skewed lotteries. The results in Part D are closer to those in that strand of the literature than the results in Part A and B. Nonetheless, the proportion of positively skewed choices is still lower than in the literature. The lower proportion of positively skewed choices may be the result of the different type of lotteries we used (i.e., the two lotteries were mirror images of each other which perhaps emphasized the direction-bias) and an order effect (Part D was played last).

7. Concluding remarks

The contribution of the present paper is threefold. First, we examined individuals' preferences over stopping rules when they have commitment power, such that they can express their true preferences. Second, we found that individuals tend to trade off between the size of the prizes and the winning probability in a consistent manner - in favor of either right-biased stopping rules or left-biased stopping rules. The participants' choice patterns depend on the favorability of the baseline lottery and

[15] The total score in Part D is calculated as follows: each choice of positively (negatively) skewed lottery is translated to 2 points (1 point). The choices in Part D are correlated with those in Part A and B in T_p as well (Pearson's r=0.49, p<0.001 and Pearson's r=0.43, p=0.003, respectively).

cannot be explained by leading theories of decision under risk. In fact, the consistency of the participants' choices enabled us to classify two-thirds of them into types according to the way they trade off between prizes and winning probabilities. Finally, we found a significant tendency to choose stopping rules that induce a relatively large winning probability even at the cost of a large downside risk, at both the aggregate and individual levels.

The prize-probability tradeoff appears in various real-world decision problems. For example, consider a job-seeker who chooses whether to prepare for a job-interview for position A or to prepare for a job-interview for position B. Suppose that position A entails a greater compensation/status than position B such that the competition for position A is more intense and, as a result, the job-seeker's chances of converting the interview to a job-offer are lower. While the downside (not getting a job despite the effort exerted in preparing for the interview) is fixed, the job-seeker has to settle the tradeoff between the higher conversion chance at the interview for position B with the higher reward offered at position A. Our results suggest that in such contexts, where it is difficult to evaluate the exact probabilities, individuals tend to settle the tradeoff qualitatively and consistently either in favor of the probability of success or in favor of the rewards.

The results of Part C establish that almost none of the participants estimated the stopping rules' induced probabilities correctly. The participants' estimates in Part C corresponded to a probability of winning a single baseline lottery p that is closer to 0.5 than it really was. For example, in the first problem of Part C, the participants had to estimate the likelihood of finishing the game with a gain when the baseline lottery's winning probability is 18/37 and the stopping rule is (-25,+25). While the rule's induced winning probability is roughly 0.205, over 85% of the participants in the main treatment overestimated this probability, with 50% of the participants estimating that it is in the range of 0.45-0.5 and 69.9% estimating that it is in the range of 0.4-0.5.

A possible interpretation that is consistent with this overestimation is that the participants did not take into account the fact that small differences in p lead to large differences in the rule's induced probabilities. In particular, the fact that the baseline lottery's winning odds are close to 0.5 does not imply that the rule's induced lottery is close to being fair. This pattern may have significant implications for situations in which processes are perceived to be "almost fair". Thus, 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 it is slightly positive (e.g., stock market trading).

References

1. Alaoui, L. and Fons-Rosen, C. (2017): “Know When To Fold 'em: The Grit Factor,” *Barcelona GSE Working Paper*.
2. Barberis, N. (2012): ”A model of Casino Gambling,” *Management Science*, 58, 35-51.
3. 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.
4. Barkan R. and Busemeyer, J. (2003): ”Modeling Dynamic Inconsistency with a Changing Reference Point,” *Journal of Behavioral Decision Making*, 16, 235-255.
5. Bleichrodt and van Bruggen (2018): “Reflection for Higher Order Risk Preferences,” Mimeo.
6. Brunner, T., Levínský, R. and Qiu, J. (2011): ”Preferences for Skewness: Evidence from a Binary Choice Experiment,” *The European Journal of Finance*, 17, 525-538.
7. 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.
8. Deck, C. and Schlesinger, H. (2010): “Exploring Higher Order Risk Preferences,” *Review of Economic Studies*, 77,1403-1420.
9. Deck, C. and Schlesinger, H. (2014): “Consistency of Higher Order Risk Preferences,” *Econometrica*, 82,1913-1943.
10. Ebert, S. (2015): “On Skewed Risks in Economic Models and Experiments.,” *Journal of Economic Behavior and Organization*, 112, 85-97.
11. Ebert, S. and P. Strack (2015): “Until the Bitter End: On Prospect Theory in a Dynamic Context,” *American Economic Review*, 105, 1618-1633.
12. Ebert, S. and D. Wiesen (2011): “Testing for Prudence and Skewness Seeking,” *Management Science*, 57, 1334-1349.
13. Ebert, S. and D. Wiesen (2014): “Joint Measurement of Risk Aversion, Prudence, and Temperance,” *Journal of Risk and Uncertainty*, 48, 231-252.
14. 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.
15. Golec, J. and Tamarkin, M. (1998): “Bettors Love Skewness, Not Risk, at the Horse Track,” *Journal of Political Economy*, 106, 205-225.

16. Greiner, B. (2004): “The Online Recruitment System ORSEE 2.0 – A Guide for the Organization of Experiments in Economics,” University of Cologne, Department of Economics.
17. Grossman, P. and Eckel, C (2015): “Loving the long shot: Risk taking with skewed lotteries,” *Journal of Risk and Uncertainty*, 51, 195-217
18. Halevy, Y. (2007): “Ellsberg Revisited: an Experimental Study,” *Econometrica,* 75, 503-536.
19. Kahneman, D. and Tversky, A. (1992): “Advances in Prospect Theory: Cumulative Representation of Uncertainty,” *Journal of Risk and Uncertainty, 5, 297-323.*
20. Magnani, J. (2017): “Reference-Dependent Preferences and Optimal Stopping Decisions: a Laboratory Investigation,” *Mimeo*.
21. Maier, J. and Rüger, M. (2012): ”Experimental Evidence on Higher-Order Risk Preferences with Real Monetary Losses,” *working paper*.
22. Noussair, C., Trautmann S., and G. van de Kuilen (2014): “Higher Order Risk Attitudes, Demographics, and Financial Decisions,” *Review of Economic Studies*, 81 (1), 325-355.
23. Oprea, R., Friedman, D., and Anderson, S. (2009): “Learning to Wait: A Laboratory Investigation,” *Review of Economic Studies,* 76, 1103–1124.
24. Plonner, M. (2017): ”Hold on to it? An Experimental Analysis of the Disposition Effect,” *Judgment and Decision Making*, 12, 118-127.
25. Rubinstein, A. (2016): “A Typology of Players: Between Instinctive and Contemplative,“ *The Quarterly Journal of Economics*, 131, 859–890.
26. Sandri, S., Schade, C., Muβhoff, 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.
27. 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, 777-90.
28. Strack, P. and Viefers, P. (2018): “Too Proud to Stop: Regret in Dynamic Decisions,” *Mimeo*.