Over and Under Commitment to a Course of Action in Decisions from Experience

Doron Cohen, Ido Erev

Author Note

Doron Cohen, Economic Psychology, Department of Psychology, University of Basel, Switzerland; Ido Erev, Max Wertheimer Minerva Center for Cognitive Studies, Faculty of Industrial Engineering and Management, Technion, Israel. Some of this work was done when Doron Cohen was at the Max Wertheimer Minerva Center for Cognitive Studies, Technion.

This research was supported by the I-CORE program of the Planning and Budgeting Committee and the Israel Science Foundation (grant 1821/12). This work was presented at the Economic Science Association (ESA) in 2020.

All data, model code, and stimuli are freely available online at https://osf.io/gbaen/.

Correspondence concerning this article should be addressed to Doron Cohen, Center for Economic Psychology, Faculty of Psychology, University of Basel, Missionsstrasse 62a, 4055 Basel, Switzerland. E-mail: doronco30@gmail.com.

Abstract

Many natural activities involve "stopping dilemmas": Situations that require a repeated decision between investing effort to achieve some valued goal and stopping that effort to try something else. Previous research into these problems highlights two contradicting biases. While one class of studies suggests a tendency to stop too late (e.g., Escalation of commitment), another class of studies suggests a tendency to give up too early (e.g., Learned helplessness). Our paper clarifies the conditions that trigger these biases by focusing on two factors: The decision mode (ongoing decisions vs. planning in advance) and the probability each search effort will be costly. We find that experience with stopping dilemmas produces a reversed sunk-cost effect: Most participants stop too early when search is frequently costly but stop too late when search is usually rewarding. This effect can be explained by assuming that stopping decisions reflect reliance on small samples of past experiences with similar stopping dilemmas. Comparison of ongoing and planning decisions reveals an interaction: Planning in advance increased search when searching was frequently costly, but decreased search when most search efforts were rewarding. This interaction can be explained by assuming a contingent re-evaluation process: Recent losses increase the tendency to reevaluate a plan to continue the search, and recent gains increase the tendency to reevaluate a plan to stop. In addition, we observe a preference for stopping strategies that imply maximal search. We assume this reflects an attempt to explore the full problem space.

Keywords: Decisions from experience, Search, Secretary Problem, Early abandonment, Description-experience gap

Over and Under Commitment to a Course of Action in Decisions from Experience

Many natural activities involve decisions between continuing an effort to achieve some valued goal and stopping that effort to try something else. For example, reading the current text implies the repeated decision by the reader to either continue the effort to discover interesting insights, or to try a different activity. Another example is online dating – although a usually rewarding activity with high potential gains, dating without caution might imply considerable hazards (e.g., Jerin & Dolinsky, 2001).

Previous research into these “stopping dilemmas” highlights two contradicting biases. One class of studies suggests that human and non-human subjects tend to over-commit resources to failing ventures. For example, Over-commitment was found to be particularly likely after large initial investments. Arkes and Blumer (1985, Experiment 3) asked participants to play the role of an airline executive and decide whether to invest in an already failing project. In one condition, participants were told they have already invested 9 million dollars in the project. In a second condition, participants were told no previous investments have been made. While about 80% of participants chose to invest the last one million dollars in the former condition, only about 17% chose to do so in the latter. Common names for this bias are "Sunk cost effect" (Arkes & Blumer, 1985; Thaler, 1980), "Escalation of commitment" (Staw, 1976) and the "Concorde fallacy" (Dawkins & Carlisle, 1976; Navarro & Fantino, 2005).

A second class of studies suggests a tendency to give up too early. For example, in one of the experimental conditions run by Zikmund-Fisher (2004, Study 1), participants were asked to decide when to stop an incremental investment in search for a profitable goal. Participants were presented with an urn containing 100 balls. They were told the urn has 85% chance to hold 25 red winning balls, and 15% chance to hold no winning balls. At each trial, participants had to decide whether to draw one more ball or quit the “game”. Each draw implied a small cost, and finding a winning ball ended the game. Although the optimal strategy was to quit the game at trial 10 (if the treasure was not found earlier), 75% of participants quitted by trial 6, demonstrating a tendency to under-invest. Common names for these phenomenon include “Learned helplessness" (Seligman, 1972; Teodorescu & Erev, 2014a), “De-escalation of commitment” (Heath, 1995; McCain, 1986), “Reverse sunk cost effect” (Zeelenberg & Van Dijk, 1997) and early termination in optional stopping tasks (e.g., in the “Secretary problem”, Seale & Rapoport, 1997; and see also Brockner, Shaw, & Rubin, 1979; Hoelzl & Loewenstein, 2005; Lejuez et al., 2002; Wallsten et al., 2005 for a similar pattern of results).

Part of the conflicting pattern can be described as the product of a description-experience gap (DE-gap, see Fantino & Navarro, 2012; Friedman et al., 2007; and related research in Hertwig & Erev, 2009). The DE-gap refers to the finding that behavior differs when decisions are based on described information rather than experienced outcomes. Indeed, most studies that found over-commitment relied on eliciting the participants' stated stopping preferences in reaction to some description-based investment dilemma. In contrast, most studies that found under-commitment relied on observed stopping preferences in sequential search tasks with feedback concerning the outcomes of previous investments. This explanation implies that stopping tends to be later in description-based decisions, and earlier in experience-based decisions.

Bearden and Murphy (2007) proposed a simple model that implies a sufficient condition for the emergence of a DE-gap in stopping dilemmas. Their model assumes the use of cutoff strategies. A generalization of this model to the current context implies agents invest effort until their "investment cutoff" is reached.[1] In addition, the model assumes a stochastic cutoff selection process; it assumes that agents repeatedly redraw and "reevaluate" their cutoff from a distribution of cutoffs during search. Assuming variability in the drawing process, this model implies that the mere opportunity to change the cutoff increases the probability of early stopping. To clarify this "constant re-evaluation" hypothesis, consider an agent that uses a rule of the type "stop after x failed tries", with x drawn from a logistic distribution (as assumed in Bearden & Murphy, 2007). For our example, let us assume this distribution has a mean ($\overline{x}$) of 10 failed tries and scale (parameter proportional to the standard deviation) of 3. Assuming an optimal stopping rule of $x^* = 9$, eliciting participant's discrete stopping decisions from description will likely reflect stopping too late (as $\overline{x} > x^*$). However, if the agent redraws a cutoff after each try, the actual (observed) average stopping will be only 6.54 failed attempts.

While Bearden and Murphy's (2007) abstraction of the DE-gap in stopping dilemmas is elegant, it appears to be inconsistent with the best known examples of over-commitment. Prominent examples include irrational persistence in failing ventures such as the Concorde (Dawkins & Carlisle, 1976) and Taurus (Drummond, 1999) projects, and the US involvement in the Vietnam war (Staw, 1976). However, these examples of over commitment unmistakably come from the analyses of repeated decisions from experience (e.g., Brockner, 1992; McCain, 1986; Staw, 1981). Further, sunk cost-like behaviors were observed in studies with non-human

---

[1] Bearden and Murphy's analysis focused on a generalized secretary problem. The cutoff in their analysis was a function of the estimated quality of the next candidate. The current analysis focuses on a simpler situation in which the decision maker knows with certainty when the goal is reached. Both Bearden and Murphy's and our analyses assume that the cutoff strategies represent the number of accepted failures by each agent.

subjects such as pigeons and rats, where description is (usually) irrelevant (see Magalhães & Geoffrey White, 2016 for a review).

One likely contributor to this apparent inconsistency is that these examples reflect situation-specific overgeneralizations. For example, Arkes and Ayton (1999) suggest that the sunk cost effect can be a product of an overgeneralization of a “Do not waste” rule (see also Arkes, 1996; Arkes & Blumer, 1985). This rule predicts a preference by participants for continued investments in response to losses, so that the sunk costs will not be (or perceived to be) wasted. This explanation builds on the observation that a “do not waste” rule can be highly adaptive: There are many natural situations in which there are good reasons to avoid waste and justify costs (Arkes & Ayton, 1999). For example, it is natural to assume that top officials in the US government (during the Vietnam war) believed a withdrawal from Vietnam without an achievement will end their political career. It is also possible that the continued investments in the Concorde airplane project reflected an attempt to enhance its future reputation, rather than escalation in response to sunk costs (Magalhães & Geoffrey White, 2016). The mixed evidence for the sunk cost effect in controlled repeated-choice experiments demonstrate the significance of this explanation (e.g., see Ashraf, Berry, & Shapiro, 2010; Drummond, 2014; Friedman et al., 2007; McCain, 1986; Singer & Singer, 1985).

The current paper focuses on the role of a second contributor for the co-existence of over and under commitment in repeated search tasks. Specifically, we consider the possibility that stopping decisions from experience, like other decisions from experience (see Fiedler, Brinkmann, Betsch, & Wild, 2000; Kareev, Lieberman, & Lev, 1997; Plonsky, Teodorescu, & Erev, 2015), reflect reliance on small samples of similar past experiences. Because rare events are underrepresented in most small samples, reliance on small samples implies underweighting of (long term) rare events and high sensitivity to (short term) frequent outcomes. In the current context, this implies high sensitivity to the frequency with which search for the goal is rewarding (all else being equal). Specifically, the “small samples” hypothesis predicts situation specific deviations from optimal investment: Over-investment/commitment when the probability of benefiting from further investing is high (e.g., search is mostly enjoyable), and a tendency to give up too early when the probability of losing from investing is high (e.g., search is mostly unpleasant/costly, see Teodorescu & Erev, 2014a, 2014b).

A second goal of the current research is to test the constant re-evaluation hypothesis (implied by Bearden & Murphy’s 2007 model). To do so, we explore the difference between two distinct stopping situations. In the first, the decision makers must select their commitment policy in advance (i.e., make only one choice of cutoff). In the second, decision makers must repeat the ongoing decision between continuing and stopping further investment. These situations abstract the difference between stopping decisions made by policy makers (e.g., a corporation’s board of directors) and executors of a policy (e.g., a corporation’s executive officer). While the former might convene at intervals and make strategic long-term decisions, the latter is typically in charge of managing more ongoing, mundane activities. The constant re-evaluation hypothesis (Bearden & Murphy, 2007) implies earlier stopping in the latter setting.

Overview of Studies 1 and 2

In natural investment decisions, such as development IT projects, startup firms or romantic dating, decision makers invest (time, effort and/or money) until a treasure is found (i.e., a major success). Yet, the outcomes one experiences usually differ from day to day. Some projects are rewarding in most days, while other projects are only rewarding in their final stages.

Our paradigm simulates such situations, and our analysis focuses on the impact of the experienced costs on the decision to stop during the project's early stages. Our hypotheses imply that the deviations from optimal stopping are likely to be sensitive to two factors. The first is the probability of a loss directly resulting from the decision to continue the search (keeping fixed the expected cost and probability of success). The second is the opportunity to update the cutoff during the effort to reach the goal. To study the impact of these factors we developed the "urn search" paradigm described in Figure 1. The top and middle panels present typical screenshots from Studies 1 and 2 described below (Study 3 used the same design, with a different incentive structure). The screenshots present the on-screen instructions that include a full description of the incentive structure.

The initial instructions for all the reported studies are presented in the online supplemental (https://osf.io/gbaen/). The Participants were informed they would be presented with one independent urn in each of the 30 experimental rounds. The instructions explain that an urn can contain, with equal probability, either 20 white balls, or 15 white balls and 5 red balls. The instructions informed participants that to maximize their earnings, they had to find as many red balls as they could in the 30 rounds (each red ball paid 30 points), while minimizing the number of white balls they draw (the expected loss from drawing a white ball was one point).

The participant's task was to determine when to stop searching. Balls were drawn one at a time, and after each draw the participant was shown the ball's color, before it was returned to the urn (i.e., drawing is with replacement). Although participants were fully aware of the odds determining the possible states of nature (i.e., the types of urns), they could not know definitively in which state they were in unless they found a red ball.[2] Each round terminated either when the first red ball was found, the participant decided to stop, or a predefined maximum number of sixteen draws had been reached (we chose this number arbitrarily to impose a clear boundary in the task).

Our full information optional stopping paradigm can be seen as a variant of the secretary problem (see Seale & Rapoport, 1997; Zwick et al., 2003) and the Balloon Analogous Task (see Lejuez et al., 2002; Wallsten et al., 2005). The main differences are that in our design, the optimal stopping strategy is easier to compute (simplifying inferences regarding stopping tendencies), and the probability of success (while keeping the optimal strategy fixed) is easier to manipulate. This allows comparing conditions with a different probability of success, while fixing expected returns and the level of complexity. In addition, although it is not easy (from the participant's perspective) to compute the optimal strategy, the full information design clarifies the expected costs of each unsuccessful effort (i.e., drawing a white ball).

The bottom panel in Figure 1 presents the expected values (EV) of adopting different cutoff strategies of the type "Leave after drawing x white balls" for the feasible values of x. The unimodal distribution reflects the fact that earlier stopping reduces the probability of finding a treasure if the state is winnable (i.e., the urn includes red balls) while later stopping increases the expected costs. The optimal cutoff assuming risk neutrality with the current parameters is $x^* = 7$ (see online supplemental for a description of the calculation and R code for a simulated solution,

---

[2] For a more intuitive description of the task, consider the following equivalent yet simpler scenario. In each round you are handed one of two completely indistinguishable coins. You know that one of the coins never falls on "heads", while the other falls on "heads" with only 25% probability. You do not know which type of coin you are handed. You can either flip it until it falls on heads or quit the game. Each "tails" implies, for example, a cost of 1 point, and a "heads" outcome gives you 30 points.

at https://osf.io/gbaen/). Thus, in our task, value maximization implies escalating investments until reaching the optimal cutoff, and de-escalating investments immediately after. We use this simplified rule as a benchmark to observe the direction in which participants in our task deviate from optimal choice.

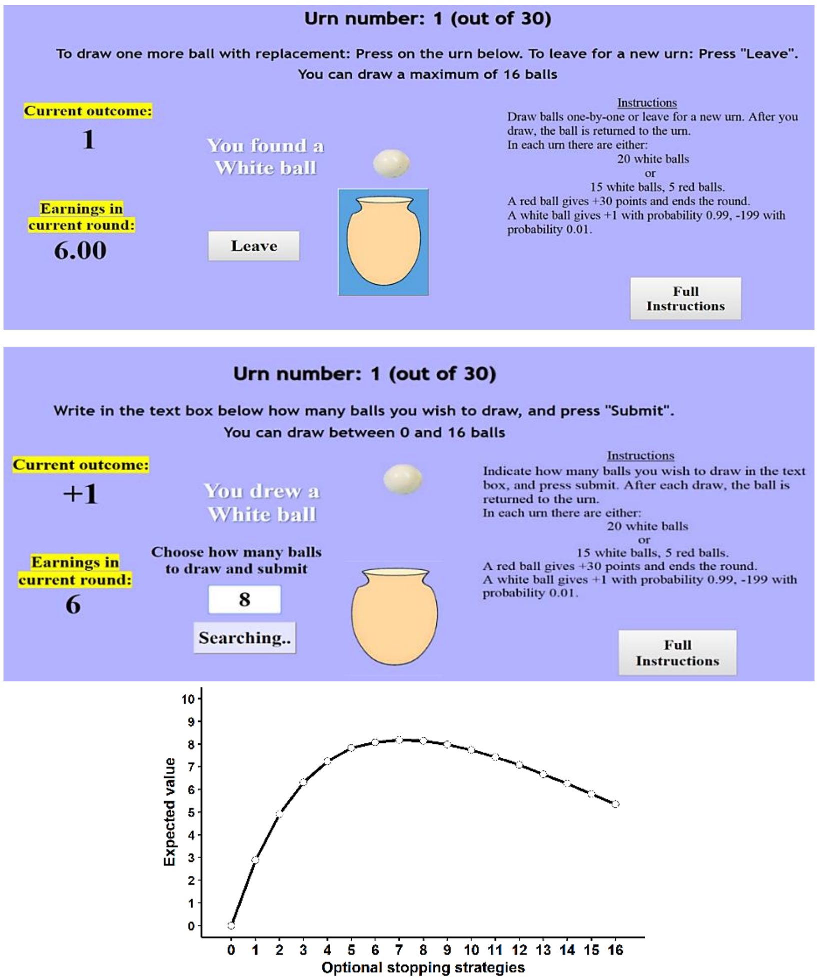

Figure 1. Top panel: A typical screen shot from Study 1 (“Ongoing”) in which participants had to make ongoing draw/leave decisions. Middle panel: A typical screen shot from Study 2 (“InAdvance”), participants had to plan and submit their stopping policy for each round in advance. In both studies, a virtual ball appeared on screen after each draw, displayed its color, and then physically dropped down back to the urn. Bottom panel: Expected values of choosing different stopping strategies of the type “Leave the round after drawing *x* white balls”.

The impact of the probability each search effort will be costly was evaluated by comparing two experimental conditions. In Condition “Cost”, the impact of drawing a white ball was a loss of 1 point with certainty. In Condition “RareLoss”, drawing a white ball implied a gain of 1 point with $p = .99$, or a loss of 199 points otherwise. Note that in both conditions, the expected value of drawing a white ball was -1 points. In Condition Cost, the probability of a loss from a white ball was 1 (i.e., search is usually costly), while the probability of loss was 0.01 in Condition RareLoss (i.e., search is usually rewarding). The left-hand side of Table 1 summarizes the incentive structures our conditions imply.

Table 1

Incentive structure, comparison to optimal strategy and mean stopping in Studies 1-3.

Decision mode	Impact of white ball	$p(S_{Red})$	#Red, V(Red)	V(White)	Optimal stopping rule	% Stopping relative to optimal: Early	% Stopping relative to optimal: Late	Mean stopping across experimental rounds
								Mean [95% CI]
Study 1								
Ongoing	Cost	.5	5, 30	-1	7	**61%**	33%**	6.6 [6.3, 6.9]
Ongoing	RareLoss	.5	5, 30	+1, .99; -199	7	26%	**67%****	10.2 [9.9, 10.5]
Study 2								
InAdvance	Cost	.5	5, 30	-1	7	47%	47%	7.9 [7.7, 8.1]
InAdvance	RareLoss	.5	5, 30	+1, .99; -199	7	37%	**59%***	9.3 [9.0, 9.5]
Study 3								
Ongoing	Cost	.1	1, 250	$+0.2 - 0.2*(t-1)$	6	**64%**	33%**	5.6 [5.3, 5.91]
Ongoing	RareLoss	.1	1, 250	$+1, .99; L_t,$ $L_t = -79 - 20*(t-1)$	6	24%	**72%****	10.8 [10.6, 11.1]
InAdvance	Cost	.1	1, 250	$+0.2 - 0.2*(t-1)$	6	55%	41%†	6.7 [6.3, 7.0]
InAdvance	RareLoss	.1	1, 250	$+1, .99; L_t,$ $L_t = -79 - 20*(t-1)$	6	42%	53%†	8.0 [7.6, 8.4]

Note. $p(S_{Red})$ = Probability the urn contains red balls. #Red = Number of red balls in winnable urns (out of 20 balls). V(Red) = Value of red ball. V(White) = the payoff from each white ball, t = trial number, L_t = Rare loss at trial t. Notice that in Study 3 the expected cost increases with trial, but the probability of loss remains constant. “Early” and “Late” columns refer to the observed proportion of rounds stopped before or after the optimal stopping rule (7 in studies 1 and 2, 6 in study 3). Significance levels represent paired sample t-test for the difference between Early and Late proportions. CI = confidence interval.

$^\dagger p < .10$. * $p < .01$. ** $p < .001$.

The current work explores the impact of the opportunity to re-evaluate and update the cutoff during the effort to reach the goal. Study 1 focuses on ongoing stopping decisions and Study 2 focuses on planning in advance. In Study 1, each round was divided into trials, and each trial started with a choice between "Leave" and "Draw one ball". In contrast, in Study 2 the participants were asked to plan and submit their cutoff (the number of balls to be drawn) before the beginning of each round. Study 3 uses a similar design to directly compare ongoing and planning decisions.

Study 1 – Ongoing stopping decisions

In this study, we consider situations in which agents must repeatedly decide (in each round) whether to continue or stop searching for the goal. Each "continue" decision implies a draw of one more ball with replacement from the urn (and incurring its costs or reward). A "stop" decision implies quitting the current urn and moving to the next round (to face a new urn). Participants played a total of 30 rounds. We pre-registered our "reliance on small samples" hypothesis (https://osf.io/5hwzj). This hypothesis predicts that frequent losses (i.e., Condition Cost) will lead to earlier stopping compared to frequent gains (i.e., Condition RareLoss). This assertion relies on the fact that reliance on small samples implies preference for the drawing policy that gave the best outcome in most previous experiences (see Cohen, Plonsky, & Erev, 2019; Erev, Ert, Plonsky, Cohen, & Cohen, 2017).[3] Note this is the opposite of the prescription of the sunk cost effect, i.e., that higher sunk costs would imply later stopping.

[3] For example, consider the comparison of cutoffs 5 and 9. In Condition Cost, stopping after 9 balls will yield a better outcome than a cutoff of 5 in only 5% of rounds (i.e., no Red by draw 5, and Red in draws 6-9). In the same condition, a cutoff of 5 will be better than 9 in 54% of rounds (otherwise they yield the same result). In Condition RareLoss, a cutoff of 9 will be better in 62% of rounds, while a cutoff of 5 will be better in only 4% of rounds. Reliance on small samples implies sensitivity to these frequencies.

Method

Participants. A (pre-registered) power analysis based on a preliminary study indicated a sample size of 65 participants in each of the two groups is suitable.[4] We slightly oversampled to ensure at least that number is assigned to each group with a random allocation mechanism, while maintaining roughly equal group sizes. One-hundred and forty participants were recruited using Prolific Academic (https://prolific.ac). One participant's data was omitted before any analysis took place, due to missing data. The final sample size consisted of 139 participants (63 females, M_{age} = 27.8; SD_{age} = 7.1).[5]

Participants were informed they will earn a fixed show-up fee of 1.4£ (about 1.83$) and will also receive a performance-dependent bonus payment: The sum of points earned in the experiment (total of the 30 rounds), added to a performance-independent endowment of 500 points (this endowment was meant to cover possible losses accrued by the participants) and converted with a rate of 1 experimental point = 0.0005£. The mean bonus was 0.35£ (about 0.46$). The experimental session lasted 14.05 minutes on average. Ethical approval was provided by the Social & Behavioral Sciences Institutional Review Board (IRB), Technion.

Procedure. Participants were fully informed of the underlying rules, and the full instructions were available throughout the experiment. Participants received feedback after each

[4] The preliminary study used a variation of the current design without replacement. For a full report of this study and data, see the online supplemental material, at https://osf.io/gbaen/. The results of the preliminary study are similar to the result of Study 1. We chose to focus on Study 1 to simplify the presentation.

[5] Before running these 139 participants, we ran a similar study with 138 participants. When analyzing the latter study, we found a minor programing error that affected the recording of the true state of nature in each round. As this error may introduce a bias toward shorter cutoffs, we ran an exact replication after correcting the mistake in the code and present the data of this second replication here. A full report and data of the original study (*N* = 138), which yielded essentially identical results, appears in the online supplemental material (at https://osf.io/gbaen/).

trial regarding the type of ball drawn, trial number, and payoff. The feedback disappeared after 0.9 seconds, after which time a choice could be made again. Each round terminated when either a red ball was found, the participant decided to stop and move to the next round, or more than 16 balls had been drawn. After a round ended, a new one immediately began, and the participants were presented with a new urn. We used a between-subject design, with 70 and 69 participants in Conditions Cost and RareLoss, respectively.

Results

The upper rows in Table 1 summarize the main results. In Condition Cost, participants exhibited a tendency to stop too early, while in Condition RareLoss participants exhibited a tendency to stop too late. The left panel in Figure 2 presents the cumulative density functions (CDFs) for the probability of stopping further investment after each draw, across the 30 rounds. Only non-winnable rounds (i.e., urns without red balls, this was done to avoid bias in the results) were entered in the analysis. A non-parametric bootstrapped Kolmogorov-Smirnov test (using R package Matching, Sekhon, 2008) rejects the hypothesis that the distributions are equal, D = 0.53, bootstrap p = .010 (Naïve p = .017).

In addition, across both conditions, in 23% of the no-treasure rounds (i.e., urns with no Red treasure balls) participants did not explicitly indicate they wish to stop drawing. That is, they behaved in these rounds as if they selected the cutoff 16 (the round was stopped automatically after 16 balls were drawn). This proportion is significantly higher than 1/16, the proportion expected assuming a cutoff is randomly chosen (95% CI [18%, 28%], $t(138) = 6.46$, $p < .001$). Figures A1.a and A1.b in Appendix A presents individual-level stopping decisions, distributed across participants (Figure A1.a) and as a function of rounds for each participant (Figure A1.b).

The right panel in Figure 2 presents the average trial in which participants stopped (in non-winnable rounds) as a function of round number. To test the statistical significance we used a linear mixed-effects model (using R packages lme4, Bates, Maechler, Bolker, & Walker, 2015; and lmerTest, Kuznetsova, Brockhoff, & Christensens, 2017). We set random intercepts for participants, and fixed effects for Condition (2 levels: Cost and RareLoss) and Round number (rounds 1-30, treated as a continuous variable) along with their interaction. We used Satterthwaite approximation to estimate degrees of freedom. Restricted maximum likelihood estimations are reported. Consistent with our preregistered hypotheses, we find a significant main effect for Condition ($\chi^2(1) = 350.15$, $p < .001$), for Round number ($\chi^2(1) = 22.27$, $p < .001$), and for their interaction ($\chi^2(1) = 10.04$, $p = .002$). A post-hoc analysis shows the mean cutoff decreased as a function of round number in Condition Cost, $b = -.080$, 95% CI [-.111, -.048], $t(5750) = -4.9$, $p < .001$, while the decrease observed in Condition RareLoss was not significant, $b = -.019$, 95% CI [-.051, .012], $t(5750) = -1.19$, $p = .232$.

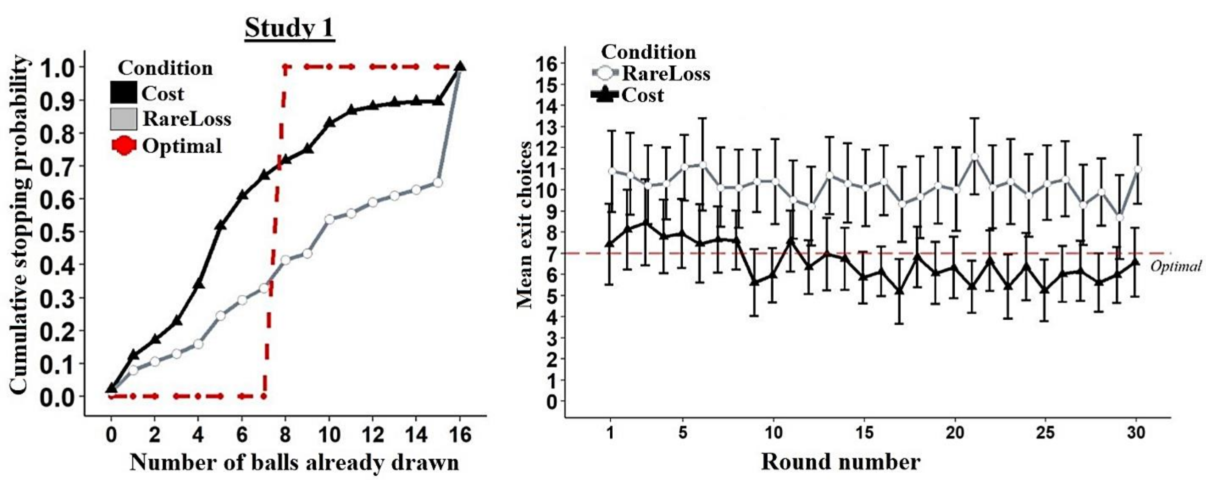

Figure 2. Left panel: Cumulative density of stopping for each feasible number of draws (across the 30 rounds). The red line represents the optimal strategy. Right panel: Average exit choices across the 30 rounds, as a function of the different experimental conditions in Study 1. Error-bars represent the 95% CIs.

The top row in Table 2 summarizes two analyses that clarify the impact of experiencing a loss of 199 points in Condition RareLoss in the current study. The first analysis examines the impact of a rare loss on subsequent decisions within the same round. It focuses on the choices of 34 participants who experienced a loss of 199 points in a no-treasure round (urn without red balls; we excluded 16 participants that never experienced a rare loss, and 19 participants that experienced a rare loss in a treasure round). We denote the trial number in which the first rare loss was experienced as t^* and compute two scores for each of these 34 individuals. The first is the stopping trial in the round with the first rare loss. The second is the average stopping trial over all remaining no-treasure rounds in which the participant drew at least t^* white balls. The difference between these scores was not statistically significant, $M_{One\ loss}$ = 13.9, 95% CI [12.6, 15.1] and $M_{No\ Rare\ loss}$ = 13.8, 95% CI [12.7, 14.9], $t(33) = 0.09$, $p = .931$.

Table 2.
The effect of experiencing the large loss in Condition RareLoss on stopping.[a]

		Impact of outcome in trial t^*, stopping in the same round[b]		Mean stopping relative to round with the first large loss experience[c]		
Study	Stopping mode	Large loss	+1	Before	After	Never experienced
1	Ongoing	13.9 [12.6, 15.1]	13.8 [12.7, 14.9]	10.8 [9.55, 12.1]	10.6 [9.42, 11.7]	7.84 [5.36, 10.3]
2	InAdvance	-	-	9.99 [8.96, 11.0]	9.96 [8.93, 11.0]	6.76 [3.04, 10.5]
3	Ongoing	13.2 [12.4, 14.1]	12.85 [11.5, 14.2]	11.9 [10.6, 13.2]	12.5 [11.2, 13.7]	5.6 [2.6, 8.54]
3	InAdvance	-	-	9.2 [8.24, 10.1]	8.5 [7.46, 9.54]	5.25 [2.66, 7.84]

[a] Values present average stopping and 95% CIs in rounds without a red ball.
[b] This analysis focuses on 34 (in Study 1) and 52 (Study 3) participants in Condition RareLoss with observations before and after experiencing exactly one large loss outcome; t^* denotes the trial number of this experience.
[c] This analysis focuses on rounds before and after experiencing only the first rare loss (excluding the rounds in which a rare loss appeared). Number of participants who never experienced a rare loss were 16 in Study 1 and 8 in Study 2, and 8 (Ongoing) and 7 (InAdvance) in Study 3.

The second analysis examines the impact of a rare loss on the decisions made in the preceding and the following rounds. It compares the means of stopping trial (in rounds without a red ball) in the rounds played before and after the round in which the first rare loss has been experienced. We also compare these observations to the average stopping by the 16 participants that never experienced the -199 outcome. The results reveal a non-significant difference between rounds played before and after the first rare loss, $t(48) = 0.42$, $p = .674$ (we excluded 4 participants who experienced the rare loss in the first round, and also the 16 participants who never experienced the rare loss). Conversely, participants that never faced a rare loss tended to stop significantly earlier than participants that experienced at least one rare loss, $t(21.59) = 2.15$, $p = .043$. At first glance, the latter comparison appears to reflect a sunk cost effect: People that experienced a large cost search more. However, our results suggest that this interpretation is incorrect. Specifically, we find that experiencing the rare loss did not affect stopping choices either in subsequent rounds or within the same round. This suggests that the observed correlation between experiencing losses and stopping position reflects the opposite causal relation: People who search more tend to experience more losses.

Discussion

The results of Study 1 support the small samples hypothesis that implies underweighting of rare events: Our participants behaved as if they underweight the rare loss, searching more in Condition RareLoss than in Condition Cost.[6] In the current setting, this tendency implies a

[6] At first glance, this pattern is inconsistent with the results reported by Zwick, Rapoport, Lo, and Muthukrishnan (2003), as they reveal participants searched too much when the search was costless and did not search enough when the search was costly. We believe that the difference reflects the fact that in Zwick et al.'s study, the manipulation of the cost also changed the optimal cutoff. This change, which was

reversed sunk cost effect: The participants were more likely to stop when experiencing mostly losses, and more likely to continue when experiencing mostly gains.

In addition, the results reveal two observations that we did not initially predict. The first is the difference observed between the two conditions on the very first round. The second observation is the high prevalence of stopping at the maximal cutoff of 16. Under one explanation, these observations reflect the joint impact of two hypothetical tendencies: (1) A tendency to explore and sample as many balls as possible, and (2) Higher probability to re-evaluate the initial cutoffs (and stop) after a loss. While the first assumed tendency implies a preference for exploration, as suggested in previous studies (e.g., Teodorescu & Erev, 2014), the second implies a new "re-evaluation after losses" hypothesis. Study 2 explores this hypothesis.

Study 2 – Explicit cutoff selection in advance

Study 2 was identical in design to Study 1 with one exception: The participants were required to explicitly state their cutoff in advance (i.e., in the beginning of each round). They did not get an opportunity to change their stated cutoffs once they submitted it. The middle panel in Figure 1 presents the game instructions (the online supplemental presents the initial, full instructions, https://osf.io/gbaen/). After submitting their preferred number (between zero and sixteen), the participants watched the computer draw that number of balls, one by one (with the same presentation method used in Study 1). If a red ball was found, further drawing immediately stopped, and a new round automatically began.

---

avoided in our study, together with regressive cutoff placement, can explain the apparent inconsistency with Zwick et al.'s results.

Study 2 was designed to test our pre-registered hypothesis (aspredicted.org/2ie8z.pdf) regarding the effect of removing optional stopping within each round. Specifically, we hypothesized this will not affect participants' average choices in the RareLoss Condition (i.e., they will still stop too late). Conversely, we predicted it will reduce the participants' sensitivity to losses in the Cost Condition. Specifically, we predicted that the decision mode manipulation would drive participants in Condition Cost to stop later than in Study 1. This is because our manipulation removes the opportunity to re-evaluate and update the current strategy when it yields negative results. This prediction implies a reduction in the difference between the two conditions, compared to the difference observed in Study 1.

Method

Participants. One-hundred and thirty-three participants (52 females, M_{age} = 27.9; SD_{age} = 8.4), were recruited using Prolific Academic (https://prolific.ac). Participants were informed they will earn a fixed show-up fee of 1£ (about 1.23$) and will also receive a performance-dependent bonus payment: The sum of points earned in the experiment (total of the 30 rounds), added to a performance-independent endowment of 500 points and converted with a rate of 1 experimental point = 0.0005£. The mean bonus was 0.34£ (about 0.42$). The experimental session lasted 10.2 minutes on average. Ethical approval was provided by the Social & Behavioral Sciences Institutional Review Board (IRB), Technion.

Procedure. The design of the task was identical to the design used in Study 1, except that at each round, participants made only one choice. They were first asked to submit the number they wished to draw from the current urn (between zero and sixteen), and then watched that number of draws made automatically, one by one (drawing was with replacement). If a red ball was found, further drawing in that round stopped, and a new round immediately began. We used a between-subject design, with 68 and 65 participants in Conditions Cost and RareLoss, respectively.

Results

The middle rows in Table 1 summarize the main results. It shows that the requirement to select the cutoff in advance reduced the deviations from the optimal cutoff. In Condition Cost, the optimal cutoff was also the median observed choice (47% of choices were below and 47% were above optimal). In Condition RareLoss, 59% of the observed cutoffs were higher than optimal.

The left panel in Figure 3 presents the cumulative density functions (CDFs) for choosing each possible search strategy across the 30 rounds. The observed distributions of strategies are consistent with our prediction that making stopping choices in advance will reduce the differences between the conditions. A non-parametric bootstrapped Kolmogorov-Smirnov test could not reject the hypothesis that the distributions are equal, $D = 0.29$, bootstrap $p = .414$.

As in Study 1, the modal cutoff is 16: In 18.2% (95% CI [14%, 23%]) of the rounds without a red ball, the participants sampled as much as possible, a rate that is significantly higher than 1/16 ($t(131) = 5.38$, $p < .001$). Appendix A presents individual level stopping decisions, distributed across participants (Figure A2.a) and as a function of rounds for each participant (Figure A2.b).

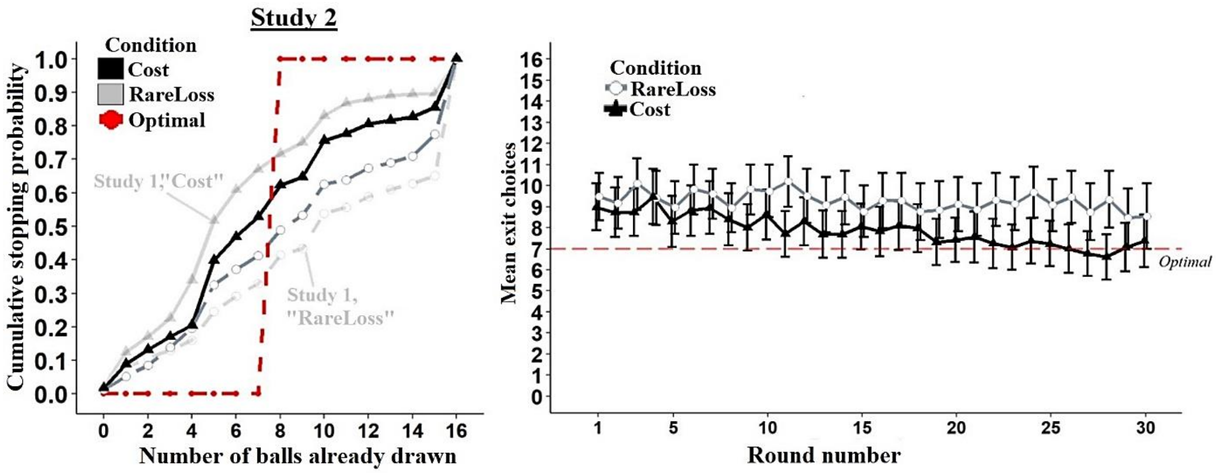

Figure 3. Left panel: Cumulative density of stopping for each feasible number of draws (across the 30 rounds) in Study 2. The red line represents the optimal strategy. The transparent lines show for comparison the CDFs observed in Study 1, by condition. Right panel: Average exit choices across the 30 rounds, as a function of the different experimental conditions in Study 2. Error-bars represent the 95% CIs.

The right panel in Figure 3 presents the mean of the stated cutoffs as a function of round number. The results are partially consistent with our pre-registered hypotheses: Figure 3 suggests requiring participants to explicitly choose a cutoff at the beginning of each round reduced the difference between the two conditions at the start of the experiment. Yet, the similarity disappeared after a few rounds, as the preferences of the two groups diverged as a function of experience. We tested the statistical effects using the same linear mixed-effects model reported for Study 1. We found a significant main effect for Condition ($\chi^2(1) = 95.94$, $p < .001$), for Round number ($\chi^2(1) = 40.56$, $p < .001$) and a significant interaction between them ($\chi^2(1) = 10.52$, $p = .001$). A post-hoc analysis shows that the mean cutoff decreased faster as a function of round number in Condition Cost, $b = -.076$, 95% CI [-.095, -.058], $t(5750) = -8.4$, $p < .001$, compared to Condition RareLoss, $b = -.025$, 95% CI [-.043, -.007], $t(5750) = -2.76$, $p = .006$. This asymmetric learning effect suggests participants in general tended to reduce their stated cutoffs as a function of experience, but this effect was stronger in Condition Cost.

Analysis of the impact of experiencing a -199 outcome (in Condition RareLoss, see Table 2) reveals the pattern documented in Study 1: Non-significant differences between rounds played before and after the first -199 outcome, and a large difference between participants that experienced this outcome, and the participants that never experienced it (second row of Table 2).

Discussion

Comparison of the results of Study 2 and Study 1 highlights two similarities and two differences. One similarity involves the impact of the experienced costs. In both studies, experiencing frequent costs (i.e., in Condition Cost) led decision makers to stop earlier, relative to Condition RareLoss. This pattern suggests a reversed sunk cost effect. A second similarity involves the observation that the highest cutoff (16) was the modal choice in both studies. As noted above, this observation can be explained as a product of exploration.

The differences between the two experiments involves the timing and magnitude of the difference between the two conditions. While this difference was observed already in the very first round in Study 1, it emerged later in Study 2. The difference in magnitude (i.e., Study 1/Cost < Study 2/Cost, Study 1/RareLoss > Study 2/RareLoss) suggests a two-way interaction: When the probability of loss was high (i.e., Condition Cost), planning in advance increased search. Yet, when the probability of loss was low (i.e., search is usually rewarding, as in Condition RareLoss) planning in advance decreased search. These differences can be captured by refining the constant re-evaluation hypothesis (Bearden & Murphy, 2007) presented above. The refinement implies a contingent re-evaluation rule: Re-evaluation is more likely after experiencing a loss when the current plan is to "continue", but also after experiencing a gain when the current plan is to stop (see similar idea in Yechiam & Hochman, 2013). This assumption can explain why the option to re-evaluate the cutoff (i.e., in Study 1) led to earlier stopping in Condition Cost but had the opposite effect in Condition RareLoss (when comparing the results of Studies 1 and 2).

Taken together, the main results observed in Studies 1 and 2 can reflect the joint impact of three behavioral tendencies: Reliance on small samples (implying underweighting of rare events and sensitivity to the frequent reward from search), a tendency to explore the problem space (implying a modal response of maximal search) and contingent re-evaluation rules.

Study 3 –Generality of the current results

Our summary of Studies 1 and 2 can be criticized on several grounds. First, while we observe clear differences between ongoing decisions (Study 1) and planning decisions (Study 2), these differences were observed in different studies, disallowing evaluation of their significance. Second, the observation that the ongoing opportunity to revise the cutoff increased over-commitment in Condition RareLoss was not predicted a priori. Thus, it might be a chance result. Third, it is possible that the effect of round number on stopping decisions reflects the impact of the cumulative payoff rule we used (e.g., a house money effect).

Study 3 was designed to address these critics. Specifically, it tests if the three-tendencies summary of Studies 1 and 2 can be used to predict behavior in a new setting. We have pre-registered these predictions (https://aspredicted.org/fz8ab.pdf).[7] Study 3 examines variants of the conditions investigated in Studies 1 and 2 in a 2x2 design (to allow clear tests of statistical significance). Each participant was randomly allocated to one of the four cells presented in the

[7] In addition to the qualitative hypotheses, we also pre-registered quantitative predictions based on a model that generalizes Bearden and Murphy's (2007) abstraction. This model can be found online at https://osf.io/gbaen/. A simpler model that provides similar predictions is discussed below, and described in Appendix C.

lower left side of Table 1. Two cells focused on ongoing decisions and used a variant of Study 1's task, and two cells focused on explicit cutoff policies and used a variant of Study 2's task. The main difference between the conditions examined in Study 3 and the previous studies involved the incentive structure (composition of the urns, and the implication of drawing White and Red balls, as described in Table 1), and the final payment rule. Specifically, to eliminate a possible impact of payoff accumulation, the final payoffs in Study 3 were determined by one randomly selected round.

Method

Participants. Two-hundred and eighty-seven participants, (98 females, 187 males, M_{age} = 26.7; SD_{age} = 8.41), were recruited using Prolific Academic (https://prolific.ac). Participants were informed they will earn a fixed show-up fee of 1.4£ (about 1.83$) and will also receive a performance-dependent bonus payment. Participants were informed the probability to win a 0.35£ bonus (which was the average bonus payout over Studies 1 and 2) will be determined by their earnings in one randomly selected round, using the equation: $\frac{(Sum\ of\ points\ in\ one\ randomly\ selected\ round + endownment\ of\ 400\ points)}{665\ \text{(650 in the Cost Conditions)}}$. One hundred and three participants (64%) received the bonus payoff of 0.35£. The experimental session lasted 15.42 minutes on average. Each participant was randomly assigned to one of the four groups: 73 to Group Ongoing/Cost, 70 to Group Ongoing/RareLoss, 73 to Group InAdvance/Cost, and 71 to Group InAdvance/RareLoss. Ethical approval was provided by the Social & Behavioral Sciences Institutional Review Board (IRB), Technion.

Procedure. The procedure was identical to those of Studies 1 (Ongoing cutoff mode) and 2 (InAdvance cutoff mode) except for the different incentive structure. All groups faced 30 investment periods, and drawing was with replacement. In each period, one urn was presented. The participants were informed that the prior probability of a “non-winnable” urn (i.e., the probability the urn contains 20 white balls) was 0.9. Otherwise, the urn contains 19 white balls and only 1 red ball. In all conditions, a red ball awarded participants a payoff of +250 points. As in Studies 1 and 2, drawing a red ball terminated the round.

The outcome of drawing a white ball differed across two incentive levels. In the Cost groups, a white ball implied a payoff of 0.2 - 0.2(t - 1) where t is the trial (draw) number. That is, a white ball gave participants 0.2, 0, -0.2, -0.4, -0.6…-2.8 in trials 1 to 16. In the RareLoss groups, each draw of a white ball implied a gamble that paid “1 with probability .99; L_t otherwise” where $L_t = -79 - 20(t - 1)$. That is, a white ball gave +1 with probability .99. Otherwise, if the rare negative event happened (with probability .01) a white ball would cost participants -79, -99, -119, -139, -159…-379 as a function of the trial in which it occurred (in trials 1-16).

Importantly, the expected values of drawing each white ball were equal across the two incentive conditions. As in Studies 1 and 2, the expected values of the different cutoffs and the implied optimal strategy did not differ between the groups. The optimal cutoff (assuming risk neutrality) is $x^* = 6$ (see Appendix B for EV curve). Also note that these distinct incentives implied a different probability of experiencing a loss from search in each trial. In the Cost groups, following the second draw (and assuming a red ball was not found), the probability of experiencing a cost from search was 1 (each white ball implied an increase of 0.2 in cost). In the RareLoss groups, the probability of experiencing a cost was only .01.

Results

The average cutoffs, presented in the lower right rows of Table 1, reflect the 2-way interaction suggested by Studies 1 and 2: In the Cost groups, when the probability of loss was high, planning in advance increased the average cutoff from 5.6 (in ongoing decisions) to 6.7 (decisions in advance). But in the RareLoss groups, when the probability of loss was low, planning in advance decreased the cutoff from 10.8 (for ongoing decisions) to 8.0 (decisions in advance). This interaction was highly significant ($\chi^2(1) = 19.39$, $p < .001$). Analysis of the common deviation from the optimal cutoff (Table 1 and the left-hand side of Figure 4) highlights a similar pattern.

As in the previous studies the modal cutoff is 16. Over the four experimental conditions, participants sampled as much as they were allowed in 23.4% (95% CI [20%, 27%]) of the rounds (we include in this analysis only rounds with no red balls). This is significantly higher than 1/16 (expected assuming random draw of cutoffs), $t(287) = 9.67$, $p < .001$. The modal cutoff of 16 was observed in 45.3% (95% CI [37%, 53%]), 20% (95% CI [14%, 26%]), 17% (95% CI [11%, 23%]) and 12.6% (95% CI [7.1%, 18%]) in Conditions Ongoing/RareLoss, InAdvance/RareLoss, InAdvance/Cost and Ongoing/Cost, respectively.

An analysis including the effect of experience (using the same statistical method as in Studies 1 and 2, see right-hand side of Figure 4) suggests the 3-way interaction between Incentive structure (Cost/RareLoss), decision Mode (Ongoing/InAdvance) and Round number (1-30) was not significant ($\chi^2(1) = 0.12$, $p = .734$). Analysis of the 2-way interactions and main effects reveal the pattern predicted in our pre-registration: While Mode significantly interacted with Round number ($\chi^2(1) = 24.96$, $p < .001$), the simple main effect of Mode was not significant ($\chi^2(1) = 2.61$, $p = .106$). Also consistent with the predicted results, while the interaction between

Incentive and Round number was not significant ($\chi^2(1) = 0.65$, $p = .418$), the simple main effect for Incentive was highly significant ($\chi^2(1) = 49.25$, $p < .001$). A simple slope analysis clarifies this pattern: While the decrease in mean cutoff with rounds was much steeper in Conditions InAdvance/Cost ($b = -.097$, 95% CI [-.117, -.078], $t(7864) = -9.73$, $p < .001$) and InAdvance/RareLoss ($b = -.084$, 95% CI [-.104, -.065], $t(7864) = -8.48$, $p < .001$), the decrease was much slower in Conditions Ongoing/Cost ($b = -.042$, 95% CI [-.062, -.022], $t(7864) = -4.11$, $p <$ 001) and Ongoing/RareLoss ($b = -.037$, 95% CI [-.058, -.015], $t(7864) = -3.41$, $p < .001$). This suggests the between-round experiences did not moderate the effect of the incentive structure.

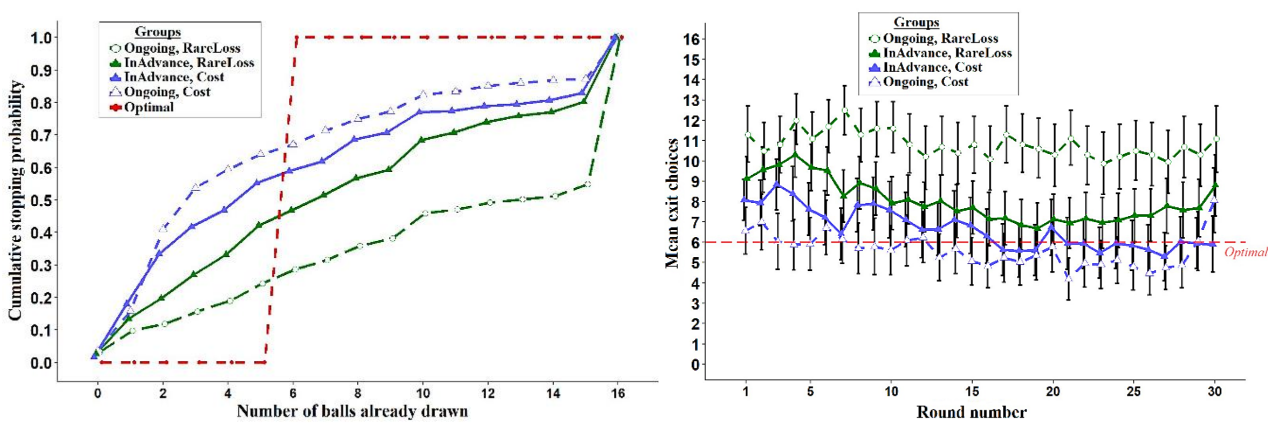

Figure 4. Left panel: Cumulative density of stopping for each feasible number of draws (across the 30 rounds) in the four conditions of Study 3. Right panel: Average exit choices across the 30 rounds, as a function of the different experimental conditions in Study 3. The red line represents the optimal strategy. Error-bars represent the 95% CIs.

The lower rows in Table 2 show that the impact of rare losses in the current study is similar to what we observed in Studies 1 and 2. As in our analysis in Study 1, we focus on the impact of a rare loss on subsequent decisions within the same round in the Ongoing/RareLoss Condition. This analysis includes 52 participants who experienced a large loss (i.e., 79 or more points) in a no-treasure round (we exclude from this analysis 8 participants that never experienced the large loss, and 10 participants that experienced the large loss in a treasure round). It suggests that experiencing a large loss had no immediate impact on stopping within the same round: Stopping in rounds where a rare loss was experienced at trial t^* was not different to rounds where the participants drew a White ball at trial t^* (see Table 2), $t(51) = 0.90$, $p = .370$. Analysis of stopping in rounds before and after experiences of a large loss reveals a similar pattern to that observed in Studies 1 and 2. We find non-significant differences between rounds played before and after the large loss both in the Ongoing and InAdvance Conditions. In addition, we again find a large difference between participants that experienced the large loss and the participants who did not experience it throughout the experiment.

Discussion

Study 3 highlights the predictive value of the three-tendencies explanation motivated by the results of Studies 1 and 2. First, in agreement with the reliance on small samples hypothesis, the results reveal a reversed sunk cost effect. Second, the results reveal again a modal cutoff of 16 draws that can be described as reflecting a tendency to explore the problem space. Finally, the results also document a significant interaction between Cost and Decision mode, as predicted by assuming Contingent re-evaluation rules: When the probability of loss was high, planning in advance increased search. Yet, when the probability of loss was low (i.e., search is usually rewarding) planning in advance decreased search.

Appendix C presents a model that quantifies the three-tendencies explanation. It shows that best fit of the data is obtained with the assumption that before each choice, the agents rely on their outcomes in an average of 8.5 previous rounds (i.e., a small number of past experiences). In addition, this modeling exercise clarifies the importance of distinguishing between learning from two types of experiences: Learning between rounds (i.e., investment episodes, the focus of Study 2) and learning within rounds. Our model suggests that assuming reliance on small samples and a tendency to fully explore the problem space imply sufficient conditions for the observed effect of the former. The latter is captured with an assumption of contingent re-evaluation. Thus, a relatively simple abstraction of the three-tendencies explanation is sufficient to capture the behavioral stopping effects we observe.

General discussion

The current analysis was motivated by the observation that many natural activities involve stopping dilemmas: Repeated decisions between searching for a desired goal and stopping that search. Previous studies focusing on this type of tasks documented two contradicting biases. While some studies document a tendency to give up too early, other studies find the opposite bias: Giving up too late (e.g., over-investment of resources in failing projects). The current work clarifies the conditions that trigger the different biases.

Our results suggest that part of the apparent inconsistencies can be explained by considering three behavioral tendencies: The first is an effort to explore and avoid giving up too early. This tendency captures the modal behavior (i.e., stopping at 16) and can contribute to over-investment in one-shot decisions from description. The second involves reliance on small samples of past experiences. This tendency captures the reversed sunk cost effect that emerges from experience in our studies. The reliance on small samples hypothesis also implies high sensitivity to the frequent outcomes yielded during search for the desired goal. It suggests that experience drives people to stop too late when the common outcome from search is rewarding. The opposite bias emerges when the common outcome is unattractive: When most search efforts yield negative results, people tend to stop too early. The third tendency involves contingent re-evaluation of the planned strategy: An initial plan to continue is reconsidered after a loss, and an initial plan to stop is reconsidered after a gain. This tendency can explain why the impact of the probability of experiencing costs or rewards from search is reduced when the decision makers select cutoffs in advance.

The current explanation distinguishes between two classes of experiences that can affect behavior in stopping dilemmas. The first class involves local experiences. These include the previous outcomes from the current effort to reach the present goal (e.g., complete the current project). These trial-by-trial, “local” experiences are relevant only when decision makers can make ongoing investment decisions. The second class of experiences are more global; they involve the outcomes of similar projects in the past (e.g., previous urns in our study). Our results suggest the impact of these between-project experiences can be captured with the reliance on small samples hypothesis. That is, they imply insufficient sensitivity to rare events.

The apparent inconsistency between the current evidence for a reversed sunk cost effect and previous studies documenting a sunk cost effect can be attributed to five different factors. First, as noted in the introduction, the clearest evidence for the sunk cost effect were documented in studies of decisions from description. In such settings, additional investments could be used to justify the initial investment and avoid waste of the already expanded resources (e.g., Arkes, 1996; Arkes & Ayton, 1999; Arkes & Blumer, 1985). Importantly, under one interpretation, the predicted impact from generalizations of a “do not waste” rule in the current design is unclear. Specifically, it is possible that participants perceived forgoing draws of more White balls in the RareLoss Conditions as wasteful, and for that reason were more inclined to escalate investments.

A second (and related) factor involves the difference between progress and utilization decisions (Moon, 2001; Roth et al., 2015). Progress decisions imply a sequence of economically linked, small-impact choices between continuing and stopping further investment (Staw & Ross, 1989). In contrast, utilization decisions imply a choice between two equally attractive options, e.g., in situations in which allocation of funds have already been made. In such cases, the option associated with sunk costs is likely to be preferred (Roth et al., 2015). While much of the indication of the sunk cost effect involves utilization decisions, the current analysis focuses on progress decisions. We hope to address utilization decisions in future research. For example, previous studies have suggested overgeneralizations of the “do not waste” rule may have a larger impact on utilization decisions (e.g., Roth et al., 2015). Thus, one important question involves how this class of overgeneralizations interact with the effects of experience on stopping decisions.

A third factor involves the possibility of a selection bias of the type illustrated in Table 2. Specifically, participants that stopped later in our study were also more likely to experience a rare extreme loss. When looked at in isolation, this suggests that higher losses can produce later stopping (as the sunk cost effect predicts). Yet, our results show that the causal relationship is in fact reversed: Later stopping increase the likelihood of experiencing a loss, rather than losses increase the likelihood of stopping later. This suggests that a positive relationship between escalation and sunk costs can reflect an artefact stemming from the outcome generating process. One example that clarifies the importance of this observation is the analysis of gambling behaviors. While sunk costs are usually present when people exhibit problematic gambling behaviors (e.g., Fantino et al., 2005; Griffiths & Wood, 2001), our results suggest they are not necessarily the root cause for such behaviors. Instead, our findings emphasize the importance of a high frequency of small rewards in escalating different gambling choices (see also Dixon et al., 2006; Griffiths, 1999; Haruvy et al., 2001; Parke & Parke, 2013 for similar observation).

A fourth factor involves the possibility of a conflict of interests between the decision makers and the public good. For example, it is possible that this conflict partly contributed to the escalation of US involvement in the Vietnam war (Staw, 1976). Specifically, it is possible that US Presidents had good reason to believe that the decision to withdraw from Vietnam without a clear achievement will be personally too costly (see also Friedman et al., 2007). The current design eliminates the impact of conflicts of this type.

The fifth factor is the focus of the current analysis, the probability of costly outcomes. In many natural escalation commitment problems, most isolated investment efforts are successful from the decision maker's point of view. That is, the outcome (e.g., destroying another enemy stronghold) is perceived as advancing the decision maker toward the desirable goal. This explanation implies that the decision to continue in these situations appear to reflect a sunk cost effect only because they are analyzed as an isolated episode (and many times in retrospect). If compared to corresponding examples with more frequent costs, a reversed sunk cost effect would emerge.

To demonstrate the practical significance of the current results, we chose to conclude with a hypothetical example. Imagine a non-profit organization focused on offering technological education for senior adults (https://oats.org/ is one example). One key factor for the hypothetical organization's success involves maximizing students' commitment to learning the new technologies, despite considerable barriers (Broady et al., 2010; Charness & Boot, 2009). Assuming a sunk cost effect, the best policy calls for charging a large enrollment fee. Yet our analysis suggests a very different approach; the best policy would be to ensure costs are only rarely experienced, and that interactions with the new technology are rewarding most of the time.

Context paragraph

Our original goal was to clarify the effectiveness of friction nudges (e.g., Mazar et al., 2020). Specifically, we hypothesized that keeping unhealthy snacks far away will be more effective the farther away they are. This will be because the effort to consume the unhealthy snack can be re-evaluated more times (the farther they are) before reaching the end state. For example, when the snacks are located away from a sitting researcher (who plans to consume them), she can change her mind even after standing up with the intention to reach the snacks (and perhaps go for a walk outside instead). Yet we later realized that our explanation was already developed by Bearden and Murphy (2007) in their model explaining insufficient search. We then considered the option of stopping this line of research but felt that we have already invested too much to stop. Simultaneously, we also noticed that while our initial investment can be described as a sunk cost, it is not clear that this description is accurate. Specifically, we noticed our costs only appear to be sunk, and that our incentive to continue was even larger. This insight helped us address a larger puzzle than we initially planned.

References

Arkes, H. R. (1996). The psychology of waste. *Journal of Behavioral Decision Making*, *9*(3), 213–224.

Arkes, H. R., & Ayton, P. (1999). The sunk cost and Concorde effects: Are humans less rational than lower animals? *Psychological Bulletin*, *125*(5), 591.

Arkes, H. R., & Blumer, C. (1985). The psychology of sunk cost. *Organizational Behavior and Human Decision Processes*, *35*(1), 124–140.

Ashraf, N., Berry, J., & Shapiro, J. M. (2010). Can higher prices stimulate product use? Evidence from a field experiment in Zambia. *American Economic Review*, *100*(5), 2383–2413.

Bates, D., Maechler, M., Bolker, B., & Walker, S. (2015). *lme4: Linear mixed-effects models using Eigen and S4. R package version 1.1–7. 2014*.

Bearden, J. N., & Murphy, R. O. (2007). On generalized secretary problems. In *Uncertainty and risk* (pp. 187–205). Springer.

Broady, T., Chan, A., & Caputi, P. (2010). Comparison of older and younger adults' attitudes towards and abilities with computers: Implications for training and learning. *British Journal of Educational Technology*, *41*(3), 473–485.

Brockner, J. (1992). The escalation of commitment to a failing course of action: Toward theoretical progress. *Academy of Management Review*, *17*(1), 39–61.

Brockner, J., Shaw, M. C., & Rubin, J. Z. (1979). Factors affecting withdrawal from an escalating conflict: Quitting before it's too late. *Journal of Experimental Social Psychology*, *15*(5), 492–503.

Charness, N., & Boot, W. R. (2009). Aging and information technology use: Potential and barriers. *Current Directions in Psychological Science*, *18*(5), 253–258.

Cohen, D., & Erev, I. (2021, February 5). The effects of experience on commitment to a course of action - an exploration into cutoff decisions in dynamic and static environments. https://doi.org/10.17605/osf.io/gbaen

Cohen, D., Plonsky, O., & Erev, I. (2019). On the Impact of Experience on Probability Weighting in Decisions Under Risk. *Decision*. https://doi.org/10.1037/dec0000118

Dawkins, R., & Carlisle, T. R. (1976). Parental investment, mate desertion and a fallacy. *Nature*, *262*(5564), 131–133.

Dixon, M. R., MacLin, O. H., & Daugherty, D. (2006). An evaluation of response allocations to concurrently available slot machine simulations. *Behavior Research Methods*, *38*(2), 232–236.

Drummond, H. (1999). Are we any closer to the end? Escalation and the case of Taurus. *International Journal of Project Management*, *17*(1), 11–16.

Drummond, H. (2014). Escalation of commitment: When to stay the course? *Academy of Management Perspectives*, *28*(4), 430–446.

Erev, I., Ert, E., Plonsky, O., Cohen, D., & Cohen, O. (2017). From anomalies to forecasts: Toward a descriptive model of decisions under risk, under ambiguity, and from experience. *Psychological Review*, *124*(4), 369–409. https://doi.org/10.1037/rev0000062

Fantino, E., & Navarro, A. (2012). Description-experience Gaps: Assessments in Other Choice Paradigms. *Journal of Behavioral Decision Making*. https://doi.org/10.1002/bdm.737

Fantino, E., Navarro, A., & O'daly, M. (2005). The science of decision-making: Behaviours related to gambling. *International Gambling Studies*, *5*(2), 169–186.

Fiedler, K., Brinkmann, B., Betsch, T., & Wild, B. (2000). A sampling approach to biases in conditional probability judgments: Beyond base rate neglect and statistical format. *Journal* *of Experimental Psychology: General*, *129*(3), 399.

Friedman, D., Pommerenke, K., Lukose, R., Milam, G., & Huberman, B. A. (2007). Searching for the sunk cost fallacy. *Experimental Economics*, *10*(1), 79–104.

Griffiths, M. (1999). Gambling technologies: Prospects for problem gambling. *Journal of Gambling Studies*, *15*(3), 265–283.

Griffiths, M., & Wood, R. (2001). The psychology of lottery gambling. *International Gambling Studies*, *1*(1), 27–45.

Haruvy, E., Erev, I., & Sonsino, D. (2001). The medium prizes paradox: Evidence from a simulated casino. *Journal of Risk and Uncertainty*, *22*(3), 251–261.

Heath, C. (1995). Escalation and de-escalation of commitment in response to sunk costs: The role of budgeting in mental accounting. *Organizational Behavior and Human Decision Processes*, *62*(1), 38–54.

Hertwig, R., & Erev, I. (2009). The description–experience gap in risky choice. *Trends in Cognitive Sciences*, *13*(12), 517–523.

Hoelzl, E., & Loewenstein, G. (2005). Wearing out your shoes to prevent someone else from stepping into them: Anticipated regret and social takeover in sequential decisions. *Organizational Behavior and Human Decision Processes*, *98*(1), 15–27.

Jerin, R., & Dolinsky, B. (2001). You've got mail! You don't want it: Cyber-victimization and on-line dating. *Journal of Criminal Justice and Popular Culture*, *9*(1), 15–21.

Kareev, Y., Lieberman, I., & Lev, M. (1997). Through a narrow window: Sample size and the perception of correlation. *Journal of Experimental Psychology: General*, *126*(3), 278–287.

Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest package: tests in linear mixed effects models. *Journal of Statistical Software*, *82*(13).

Lejuez, C. W., Read, J. P., Kahler, C. W., Richards, J. B., Ramsey, S. E., Stuart, G. L., Strong, D. R., & Brown, R. A. (2002). Evaluation of a behavioral measure of risk taking: the Balloon Analogue Risk Task (BART). *Journal of Experimental Psychology: Applied*, *8*(2), 75.

Magalhães, P., & Geoffrey White, K. (2016). The sunk cost effect across species: A review of persistence in a course of action due to prior investment. *Journal of the Experimental Analysis of Behavior*, *105*(3), 339–361.

McCain, B. E. (1986). Continuing investment under conditions of failure: A laboratory study of the limits to escalation. *Journal of Applied Psychology*, *71*(2), 280.

Moon, H. (2001). Looking forward and looking back: Integrating completion and sunk-cost effects within an escalation-of-commitment progress decision. *Journal of Applied Psychology*, *86*(1), 104.

Navarro, A. D., & Fantino, E. (2005). The sunk cost effect in pigeons and humans. *Journal of the Experimental Analysis of Behavior*, *83*(1), 1–13.

Parke, J., & Parke, A. (2013). Does size really matter? A review of the role of stake and prize levels in relation to gambling-related harm. *The Journal of Gambling Business and Economics*, *7*(3), 77–110.

Plonsky, O., Teodorescu, K., & Erev, I. (2015). Reliance on small samples, the wavy recency effect, and similarity-based learning. *Psychological Review*, *122*(4), 621.

Roth, S., Robbert, T., & Straus, L. (2015). On the sunk-cost effect in economic decision-making: a meta-analytic review. *Business Research*, *8*(1), 99–138.

Seale, D. A., & Rapoport, A. (1997). Sequential decision making with relative ranks: An experimental investigation of the" secretary problem". *Organizational Behavior and Human Decision Processes*, *69*(3), 221–236.

Sekhon, J. S. (2008). Multivariate and propensity score matching software with automated balance optimization: the matching package for R. *Journal of Statistical Software*.

Seligman, M. E. P. (1972). Learned helplessness. *Annual Review of Medicine*, *23*(1), 407–412.

Singer, M. S., & Singer, A. E. (1985). Is there always escalation of commitment? *Psychological Reports*, *56*(3), 816–818.

Staw, B. M. (1976). Knee-deep in the big muddy: A study of escalating commitment to a chosen course of action. *Organizational Behavior and Human Performance*, *16*(1), 27–44.

Staw, B. M. (1981). The escalation of commitment to a course of action. *Academy of Management Review*, *6*(4), 577–587.

Staw, B. M., & Ross, J. (1989). Understanding behavior in escalation situations. *Science*, *246*(4927), 216–220.

Teodorescu, K., & Erev, I. (2014a). Learned helplessness and learned prevalence: Exploring the causal relations among perceived controllability, reward prevalence, and exploration. *Psychological Science*, *25*(10), 1861–1869.

Teodorescu, K., & Erev, I. (2014b). On the decision to explore new alternatives: The coexistence of under-and over-exploration. *Journal of Behavioral Decision Making*, *27*(2), 109–123.

Thaler, R. (1980). Toward a positive theory of consumer choice. *Journal of Economic Behavior & Organization*, *1*(1), 39–60.

Wallsten, T. S., Pleskac, T. J., & Lejuez, C. W. (2005). Modeling behavior in a clinically diagnostic sequential risk-taking task. *Psychological Review*, *112*(4), 862.

Yechiam, E., & Hochman, G. (2013). Losses as modulators of attention: review and analysis of the unique effects of losses over gains. *Psychological Bulletin*, *139*(2), 497.

Zeelenberg, M., & Van Dijk, E. (1997). A reverse sunk cost effect in risky decision making:

Sometimes we have too much invested to gamble. *Journal of Economic Psychology*, *18*(6), 677–691.

Zikmund-Fisher, B. J. (2004). De-escalation after repeated negative feedback: emergent expectations of failure. *Journal of Behavioral Decision Making*, *17*(5), 365–379.

Zwick, R., Rapoport, A., Lo, A. K. C., & Muthukrishnan, A. V. (2003). Consumer sequential search: Not enough or too much? *Marketing Science*, *22*(4), 503–519.

Appendix A

Figure A1. a) Individual level stopping decisions, across participants and rounds, and b) Individual level stopping decisions by rounds in Study 1. These figures show only non-winnable rounds (to avoid bias in the results).

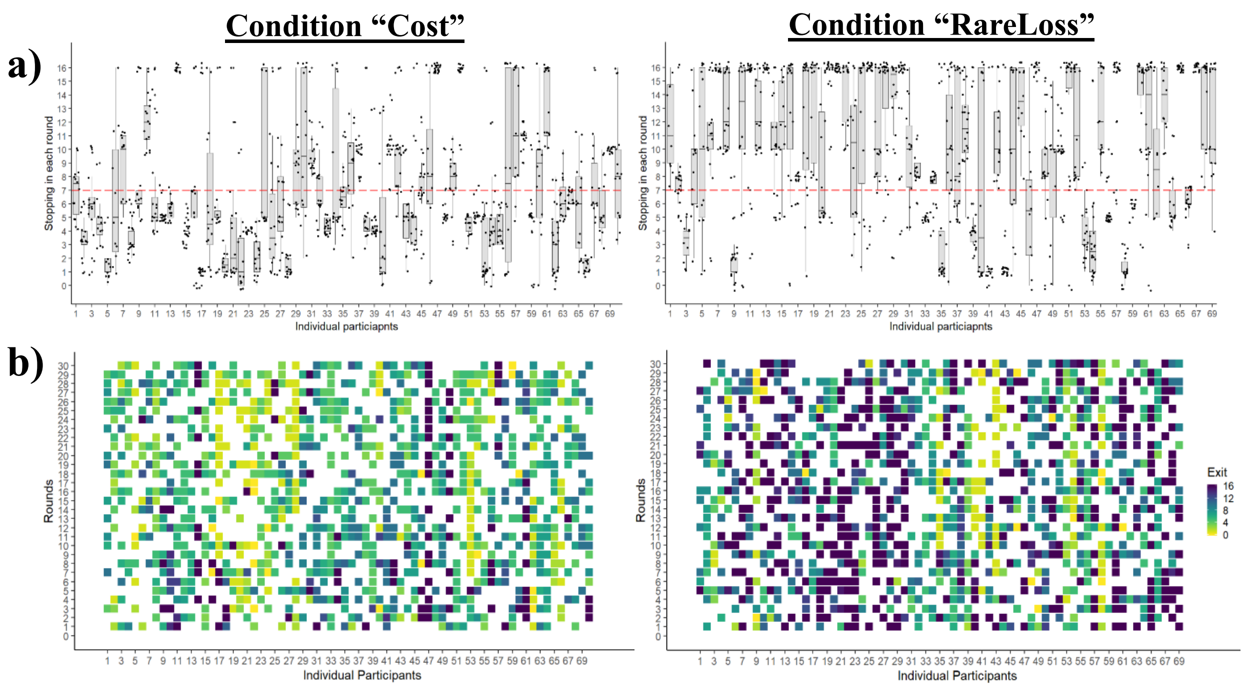

Figure A2: a) Individual level stopping decisions, across participants and rounds, and b) Individual level stopping decisions by rounds in Study 2.

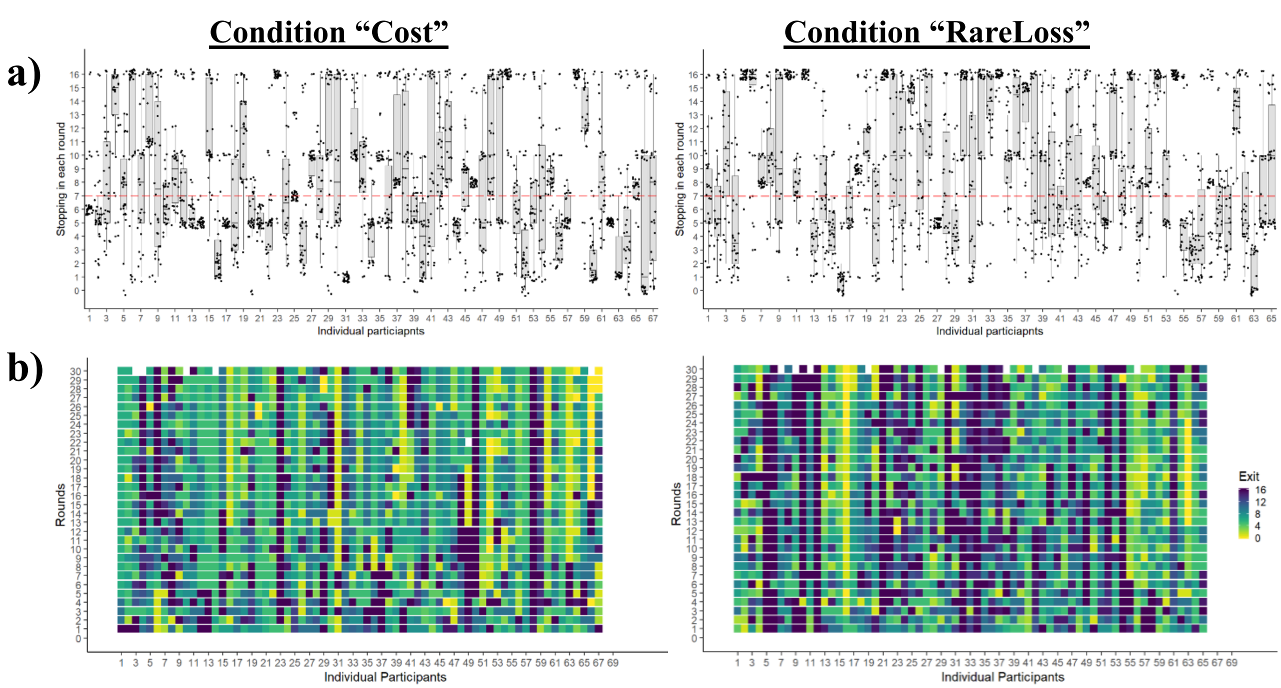

Appendix B

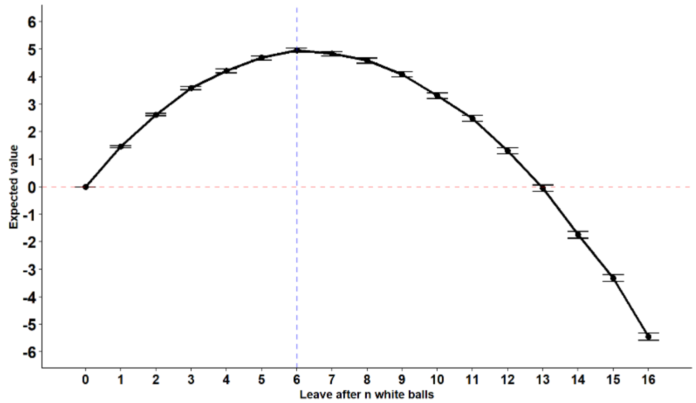

Figure B1. Expected values of choosing different stopping strategies of the type "Leave the round after drawing *x* white balls" in Study 3.

Appendix C
The Stopping Thresholds and Re-evaluation (STaR) model

The current Stopping Thresholds and Re-evaluation (STaR) model presents one feasible quantification of the three-tendencies explanation of the results: A tendency to fully explore the problem space, reliance on small samples and contingent re-evaluation rules. The code for this model can be found at https://osf.io/gbaen/.

The tendency to explore (that implies high proportion of drawing all possible balls) is captured by assuming that in addition to the (16) cutoff strategies, the decision maker considers one additional strategy that implies maximal search.

*Reliance on small sample*s is captured, as in the model BEAST (Erev et al., 2017), by assuming that the subjective value of each strategy reflects high sensitivity to a small sample of past experiences. Specifically, the model assumes the selection of the strategy with the highest subjective value. The following computation determines the subjective value of strategy j:

$$SV_j = AB_j + e_j + Sample_j$$

Where the term AB_j is the anticipated benefit from strategy j. For the 16 cutoff strategies, AB_j equals the expected value from selecting cutoff j (e.g., see Figures 1 and B1). The anticipated benefit from the additional "exploration" strategy is assumed to be equal to the maximal value of all the other anticipated benefits (i.e., $\max(AB_j)$). e_j is a noise term drawn from a normal distribution with a mean of 0 and standard deviation of σ (a free parameter). $Sample_j$ is the mean outcome of cutoff j in a small sample of k_i past experiences. The sample is drawn with replacement from all previous rounds in which the payoff from choosing this cutoff could be computed. We assume k_i, a free parameter, is an attribute of agent i, and is drawn from the uniform distribution U[1, 2, … κ]. Note that when k is small, this mechanism also implies underweighting of rare events, as rare events are under-represented in a small subset of observations.

Contingent re-evaluation is abstracted by assuming that the decision maker forms an initial opening cutoff at the beginning of each new round (by choosing the cutoff with the highest SV_j), but certain contingencies trigger re-evaluation. Re-evaluation implies independent selection of a new cutoff. Before reaching the planned cutoff, re-evaluation is possible only when the most recent draw of a White ball led to a loss. The probability of this re-evaluation is $P_{upd.}$(*before*), a free parameter. After reaching the planned cutoff, re-evaluation occurs only if the most recent draw of a White ball led to a gain. The probability of this re-evaluation is $P_{upd.}$(*after*), a free parameter.

Parameters. Our model implies four parameters: k (the subset size of previous rounds the agent considers), σ (estimation error), $P_{upd.}$(*before*) and $P_{upd.}$(*after*). We fitted these parameters on the results of our three studies. Figure C1 shows the fit of the current model with the aggregate results both between and within rounds of each condition. Best fit was obtained with the parameters k = 16*,* σ = 10, $P_{upd.}$(*before*) = .3 and $P_{upd.}$(*after*) = .6.

Studies 1 & 2

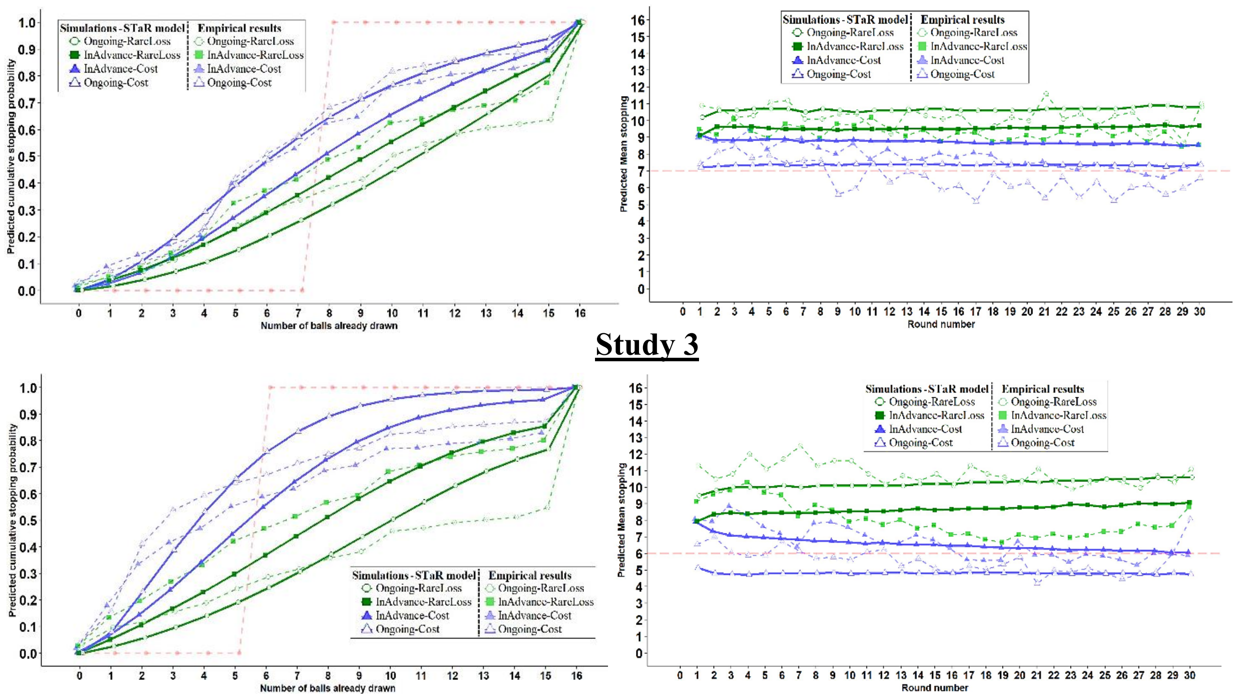

Figure C1. Shaded lines represent the empirical results in Studies 1 & 2 (upper panels) and Study 3 (bottom panels). Solid lines represent simulation results of the model for each condition in each study. Left panels: Cumulative density of stopping for each feasible number of draws (across the 30 rounds) in Studies 1 & 2 (upper left panel) and Study 3 (bottom left panel). Right panels: Average exit choices across the 30 rounds, as a function of the different experimental conditions in Studies 1 & 2 (upper right panel) and in Study 3 (bottom right panel).