

WORKING PAPER · NO. 2024-93

Limits on Regret as a Tool for Incentive Design

Felipe Araujo, Alex Imas, and Alistair Wilson

JULY 2024

Limits on Regret as a Tool for Incentive Design
Felipe Araujo, Alex Imas, and Alistair Wilson
July 2024
JEL No. D0,D03

ABSTRACT

We demonstrate the pitfalls when extrapolating behavioral findings across different contexts and decision environments. We focus on regret theory and the use of “regret lotteries” for motivating behavior change. Here, findings from one-shot settings have been used to promote regret as a tool to boost incentives in recurrent decisions across many settings. Using theory and experiments, we replicate regret lotteries as the superior one-shot incentive; however, for repeated decisions the comparative static is entirely reversed. Moreover, the effects are extremely sensitive to details of regret implementation. Our results suggest caution should be used when designing incentive schemes that exploit regret.

Felipe Araujo
Department of Economics
College of Business
Lehigh University
f.araujo@lehigh.edu

Alistair Wilson
Department of Economics
University of Pittsburgh
Posvar Hall
Pittsburgh, PA 15213
alistair@pitt.edu

Alex Imas
Booth School of Business
University of Chicago
5897 S. Woodlawn Avenue
Chicago, IL 60637
and NBER
alex.imas@chicagobooth.edu

A online appendix is available at <http://www.nber.org/data-appendix/w32759>

1 Introduction

Regret has been well-established in psychology as a powerful negative emotion (for a review see Zeelenberg and Pieters, 2007). It is activated when someone makes a choice and receives counterfactual information that a different decision would have resulted in a better outcome. While the experience of regret is backwards-looking, the *anticipation* of future regret has also been shown to impact decision-making (Zeelenberg, 1999; Janis and Mann, 1977).

For example, consider the choice to purchase a simple lottery that offers a big prize with a small probability, yielding nothing the large majority of the time. You are offered this lottery in one of two cases. In one case (call this the *standard* version), if you buy the lottery you will find out if the realization, whether you won or lost; but if you do not buy the lottery you will not learn about the realization. In the other case (the *regret* version of the lottery), if you buy the lottery everything proceeds as before, but if you do not buy the lottery you will still find out the realization. This counterfactual information generates the scope for regret as an incentive tool: finding out you would have won had you bought the lottery would make you regret a decision to play it safe—a scenario that is impossible in the standard version of the lottery. Anticipation of this future regret may then spur the purchase of the regret lottery, where the standard lottery purchase would be rejected.

Regret theory (Bell, 1982; Loomes and Sugden, 1982) proposes a simple modification to expected-utility (EU) theory to capture this phenomenon. In the model, individuals derive utility not only from the direct choice, but also indirectly through their knowledge of the outcome in counterfactual choices. Specifically, while a decision-maker considers the benefits from each possible realization from a choice (weighted by its likelihood, as in EU), additional costs/benefits are derived from a comparison to the realizations from unchosen options. Regret theory makes two assumptions: (i) people experience feelings of regret (rejoicing) when their chosen option does worse (better) than a counterfactual; and (ii) these feelings are anticipated ex-ante. For example, a person choosing a safe option over a risky lottery may experience regret if they find out the lottery had a good realization; anticipation of the potential regret then leads her to increase the ex-ante valuation for the risky lottery. The theoretical consequence for decision-making is that individuals may opt away from an EU-maximizing choice towards alternatives that reduce the prospective regret.

A recent body of work has sought to utilize regret aversion as a non-pecuniary boost to the design of incentives through regret lotteries (Zeelenberg and Pieters, 2004; Volpp et al., 2008b; Milkman et al., 2021).¹ In the case of incentive design, regret lotteries differ from standard lotteries

¹A prominent example of regret lotteries is the Dutch postcode lottery, which is used to incentivize the collection of public revenues. To realize the lottery outcome, a winning postcode is drawn from the entire population of Dutch

by ensuring that the decision-maker is aware of lottery realizations regardless of her decision to comply with the incentivized activity. In a typical regret lottery set-up, an entire population of possible entrants is first assigned a fixed and commonly-known winning state (an employee number, a zip code, their driver's license number, etc.). Their choice on whether to enter the lottery is then decided through compliance with the incentivized activity (exercising, car-pooling to work, taking a prescribed drug, getting vaccinated against COVID-19, etc.). Regardless of their compliance, the lottery state is drawn and the entire population is informed of the realization. If the realized state matches an individual's state *and* they complied, a corresponding prize is awarded. If they either failed to embrace the incentivized activity, or their particular winning state was not drawn, then no prize is awarded. This type of lottery seeks to exploit regret aversion by providing counterfactual information regardless of compliance. A regret-averse decision-maker, knowing that feedback will be provided, acts as if the lottery incentive has a larger effective value. The anticipated regret from missing out on a large prize can therefore push those on the margins to comply with the incentivized activity where they otherwise would not.

The argument for the increased incentive power of regret lotteries is based on a theoretical and empirical literature on anticipated regret that focuses on one-shot decisions, with the majority of evidence coming from lab experiments. Though many choices are indeed one-shot (e.g., entering a single lottery), a large fraction of meaningful decisions occur in recurrent, dynamic settings (e.g., whether to exercise). However, both the anticipated and realized aspects of regret change substantially in dynamic settings. Unlike one-shot decisions, repeated choices offer individuals the opportunity to learn about the incentive structure and respond accordingly. With standard lotteries that provide no counterfactual information, a learning motive might push decision makers to enter the lottery to learn about the incentives through feedback, generating greater compliance with the incentivized activity in the early stages. In contrast, individuals facing a sequence of regret lotteries can freely learn about the offered incentives through the provided counterfactuals, without having to comply. This difference in the anticipated learning can lead to *less* engagement with regret lotteries than standard ones, even if regret aversion is present. Additionally, real-world settings differ from the lab in many of ways. This makes it critical that regret as a construct be relatively robust to implementation details such as the population used or the basis of the counterfactual information. .

Our paper demonstrates that while regret exists as a psychological mechanism, its utility as an

postcodes. Each individual living within the drawn postal region that bought a ticket wins a large cash prize. Those in the winning neighborhood who did not buy a ticket do not get the prize but do observe that their postcode was chosen and that their neighbors who did purchase a ticket have won. Zeelenberg and Pieters (2004) argue that regret aversion should make this lottery program more successful than a national lottery in which the counterfactual (for example, an assigned winning state) is not publicly announced to those not buying tickets.

incentive is highly sensitive to implementation details in both the choice setting (e.g., static versus dynamic decisions) and context (e.g., target population, small changes to method of inducing regret). We use a series of controlled laboratory experiments that compare participants' valuations for identical lotteries with and without counterfactual feedback.

Our results corroborate the directional effect of higher valuations for regret lotteries in a lab setting without feedback (though the effect size is modest). However, when we examine the valuations in a dynamic setting—for sequences of independent and identically distributed (iid) decisions with feedback—the difference between the standard and regret lotteries is not only reduced, it is *reversed* in direction. We show that participants initially provide higher valuations for standard lotteries than those providing counterfactual information, though the difference dissipates as experience is gained. Moreover, replicating the design outside of the lab with an online subject pool eliminates any differences between standard and regret lotteries. Together, our results suggest caution should be applied when attempting to use regret for incentive design or as a policy tool. Seemingly innocuous implementation details eliminate the benefits of regret and, in the case of dynamic environments, may even reserve them.

Why would the predicted effects of regret be sensitive to the setting and implementation details, even reversing for repeated decisions? The one-shot aspect of regret aversion that motivate the provision of counterfactual information focuses entirely on an anticipatory feature which is isolated in the static decision. However, in a dynamic context, the provision of counterfactual information also affects decision-making via the potential for learning.^{2,3} As such, while anticipating the adaptation of decisions through time will not be an issue when the goal is to incentivize a one-time decision—such as opting into a 401-K plan or a single large lottery for COVID vaccinations—it will be when the decisions are recurrent (such as exercise, recycling, or car-pooling) or when the incentives are repeated (a smaller lottery for COVID vaccinations each week or the postcode lotteries).

In repeated settings the regret lotteries generate an anticipatory double-edged sword. On the

²In contrast, the focus within the learning literature on regret ignores the anticipatory effects, and focuses on the backward-looking realized effects on a similar choice (see for instance Erev and Barron, 2005). Our paper examines repeated decisions in the sense of many identical, statistically independent choices. In contrast, Strack and Viefers (2015) examine a dynamic decision to divest from a risky asset with a persistent state, finding that counterfactual information helps through a correction for excessive risk aversion.

³With the large-prize–low-probability gambles that are typically implemented as part of regret lotteries, the vast majority of participants will lose, and those that chose to comply with the incentivized activity will experience realized regret. Feelings of ex-post regret have been shown to lead individuals to move away from the actions that produced them (Ku, 2008), while positive emotions such as rejoicing tend to reinforce that same action in subsequent decisions (Keltner and Lerner, 2010). See Marchiori and Warglien (2008) and Hart (2005) for the use of regret as a process to predict subsequent choices.

one hand, the regret aversion can increase the value of the prospect, but on the other, any learning benefit can now be acquired ‘for free.’ Participants can adopt a wait-and-see approach without undertaking the incentivized activity. If they are uncertain about the value and structure of the offered incentives, they can learn about them through the provided counterfactual information. For example, take the case for the regret lotteries used to incentivize COVID-19 vaccinations in Milkman et al. (2021). Rather than having to take the vaccine in order to learn whether one had won or lost, as in standard lotteries, a regret lottery allows the potentially non-vaccinated to wait and see realizations (both their own, and those in their social network) without having to vaccinate. Repeated draws therefore serve to reduce the likelihood of entering the lottery even before the first outcome is realized, simply because of the learning value/curiosity that the standard lottery can induce.

Given the two countervailing forces, the predictions of regret theory are far from clear in a dynamic setting, and even less so once we allow anticipated and realized regret to interact. It is possible that anticipated regret aversion outweighs the anticipated learning effects, and the regret lottery will continue to be superior to the standard lottery in the sequential environment. However, if the learning effect dominates, the valuation for regret lotteries can be reduced below that of a standard lottery. In Section 2.1 we outline a simple model that formalizes the anticipated regret and learning effects. The model makes a series of predictions about lottery valuations across several treatment pairs in our study: across both the standard and regret lottery comparisons, but also between the static and dynamic settings. Our experimental design examines these predictions, with a particular focus on how the comparative static relevant to policymakers (regret vs. standard implementation) is affected by the move from a one-shot decision to a repeated one. Our main analyses compare decisions over standard versus regret lotteries in one-shot and repeated settings. Participants appear to value regret lotteries over standard ones in one-shot environments, but prefer the latter to the former when decisions are repeated.

However, it is also important to study the sensitivity of regret to other, seemingly innocuous changes in implementation. In the lab, we explore other design adaptations that are likely to occur in practice. These variations include allowing people to choose their winning numbers, adding a social component similar to the Dutch postcode lottery, drawing tickets randomly in each round instead of fixing a winning number, and modifying the learning opportunities in a practice stage before the main decisions. We find consistent evidence for the diminished effectiveness of regret aversion in repeated settings. Importantly, the details of the regret implementation significantly affect behavior in some cases. For example, randomly assigning a winning number each round completely eliminates the difference between standard and regret lotteries.

An important aspect of implementing a psychologically-motivated incentive is making sure that its effects are robust within the target population. Recent work has shown that many concepts from behavioral economics such as loss aversion are robust across populations.⁴ Online subject pools have increasingly been used in experimental social science, particularly in economics and psychology. As part of examining the robustness of the regret effects we documented in the lab, we replicated the same design in an online sample (Prolific). We also included a number of extensions designed to address potential concerns. We used a larger sample and pre-registered the experiments.⁵ However, where all but one of our theory-driven hypotheses were supported in the lab sample, every single comparison in our online replications had a close-to-zero effect size, and was nowhere close to significant.

Given two distinct populations with different results, we therefore decided to add a pre-registered replication of our prior lab study, trying to understand whether the initial lab results were a false positive.⁶ Data from the lab replication mirrors all of our original findings; as per the pre-registration, the combined results within this population are very similar to our initial experiment. This suggests a systematic difference between populations, which serves as another note of caution in using regret as a tool for behavior change.⁷

Our results add a new facet to the emerging literature on the potential for regret to be used to increase worker motivation (Babcock et al., 2012), policy (Madrian, 2014), improving health outcomes (Kessler and Zhang, 2014), and as a tool in development programs (Datta and Mullainathan, 2014). In much of this discussion, regret lotteries are advocated as a tool for incentivizing recurrent, ongoing choices. However, the evidence supporting regret lotteries as the superior incentive tool relative to standard lotteries comes primarily from lab experiments examining one-shot decisions (Loomes and Sugden, 1987; see Zeelenberg, 1999 for a review). While some studies do examine repeated settings, they typically compare regret lotteries to the absence of an extrinsic incentive (Volpp et al., 2008a; Milkman et al., 2021) or to fixed payments that are below the lottery's expected value (Volpp et al., 2008b). For example, Volpp et al. (2008a) examine the effectiveness of regret lotteries in incentivizing adherence to a prescription-drug regimen, where the authors report

⁴For example, Ruggeri et al. (2020) replicated all of the features of prospect theory, such as loss aversion, diminishing sensitivity, and the overweighting of low probability events, across a representative sample of more than 4,000 subjects across 19 different countries. See also Brown et al. (2024) for a recent meta analysis.

⁵Online sample pre-registration can be accessed in the following link: https://aspredicted.org/QSJ_93Z.

⁶The lab replication pre-registration can be found here: https://aspredicted.org/XQX_PVG.

⁷In line with recent work documenting increased noise in online environments (Gupta et al., 2021), we conjecture that lower attention levels are driving the online results, where regret (both static and dynamic) seems to require a great deal attentiveness to the incentives. While exploring the exact driver of instability between populations is beyond the scope of the current paper, understanding the limits of what we can ask of online participants is a pressing methodological question.

greater adherence in the lottery treatment than in the no-incentive control (see Volpp et al., 2008b; Kimmel et al., 2012; Haisley et al., 2012, for applications to other settings). On the other hand, some studies such as Milkman et al. (2021), who examine the effectiveness of regret lotteries in motivating vaccine take-up, find no improvement over the control group. Our results suggest that standard lotteries may be superior tools for encouraging compliance in repeated settings.⁸

Together, our findings highlight the caution necessary when extrapolating behavioral results from static one-shot settings to repeated, dynamic settings, as well as highlight the need to test for robustness to implementation details. Beyond just an attenuation of an effect size, the results show that the direction of the behavioral comparative static can be entirely reversed as we shift to repeated decision-making. We organize the paper as follows: Section 2 discusses the experimental design and derives testable predictions from a simple decision-making model. Section 3 presents results for our main treatments and extensions, while section 4 concludes.

2 Experimental Design

Our lab experiments were conducted at the Pittsburgh Experimental Economics Laboratory (PEEL) with participants recruited from the general undergraduate population of the University of Pittsburgh. We had a total of 486 participants (320 in our core sessions, and 166 in our robustness treatments). Each experimental session ran with 10 to 20 unique participants. Sessions lasted approximately one hour and the average payment was \$27.60, including a participation fee of either \$6 (for the first wave) or \$8 (for the second wave).^{9,10}

The first part of the experiment was common to all treatments. This consisted of a multiple price list task where participants made a series of decisions between a $\frac{1}{2}$ chance of winning \$10 and a certain amount. For each of 21 binary questions, the fixed option was a \$10 lottery while the certain amount varied from \$0.00 to \$10.00, in \$0.50 increments. While the initial task is a standard risk elicitation, its primary purpose was to familiarize participants with the subsequent valuation elicitation. That is, after participants completed the price list, we used it to motivate the BDM mechanism we use to elicit lottery valuations in the main experimental treatments.¹¹

⁸Beyond the lottery implementation question, an additional question is over the efficacy of using lotteries rather than fixed/certain payments, where Campos-Mercade et al. (2021) show fixed monetary payments are effective for increasing COVID-19 vaccine take-up. Matching this, average lottery valuations in our experiment are substantially below the lottery's expected value.

⁹The experimental design and procedures for both waves of the laboratory experiments were identical other than the participation fee. In 3.3.3 we outline the details of our online experimental sessions, and how the design changed.

¹⁰The Online Appendix includes full instructions for the first task and main experiments, as well as representative screenshots of the experimental interface.

¹¹We thank P.J. Healy for this suggestion. His paper *Explaining the BDM—Or Any Random Binary Choice Elicitation*

In all treatments, the lottery is implemented using an assigned ticket: three distinct numbers $\{A, B, C\}$ between 1 and 50. The lottery is then realized through a physical draw of three numbered balls $\{a, b, c\}$, without replacement, from a bingo cage containing fifty balls numbered from 1 to 50. Lottery prizes are determined by the number of matches on the assigned ticket, $\{A, B, C\} \cap \{a, b, c\}$. Matching one ball wins a prize of \$2.50, two a prize of \$25, while matching all three numbers yields a prize of \$250. The expected value of the lottery is:

$$\begin{aligned} \mathbb{E}V &= p_1 \cdot \$2.50 + p_2 \cdot \$25 + p_3 \cdot \$250 \\ &= \frac{3,243}{19,600} \cdot \$2.50 + \frac{141}{19,600} \cdot \$25 + \frac{1}{19,600} \cdot \$250 = \$0.61. \end{aligned}$$

Instructions for all treatments outline the prize lottery and how the realizations are determined. Moreover, in explaining how an entry ticket translates into the three different prizes, we explicitly provide the chance of winning each separate prize (\$2.50, \$25 and \$250). Lottery valuations are then elicited by asking for the maximum amount the participant would turn down to enter the lottery—using a BDM mechanism with a uniform offer draw. If the offer is at or below the elicited value, then the participant enters the lottery; if the offered amount is greater, then the participant gives up the lottery and takes the offer.

Each experimental session consists of 30 lottery draws. At the end of each round, the physical bingo cage is spun several times to mix the balls, where the experimenter then draws three balls in turn. The numbers on the three selected balls are publicly called out as they are drawn and entered into the monitor computer. Participants' screens then informed them of their current earnings for the round. Final payment is the sum of the earnings in each of the 30 rounds plus a participation fee. Additionally, in each session, one participant is randomly selected to get paid based on their choices in the initial price-list task.

Our core treatments implement a 2×2 between-subject experiment (summarized in Table 1) that varies participants' counterfactual information on the realizations of non-entered lotteries and their ability to change their valuation decisions over time in response to feedback. The availability of counterfactual information on the lottery realization is controlled through the lottery type, either *Regret* or *Standard*. In the Regret lotteries, participants' tickets are fixed at the start of the session: $\{A_i, B_i, C_i\}$ for participant i . The Regret lottery tickets are printed out and given to participants at the start of the session. Hence, in every round t , as the draw $\{a_t, b_t, c_t\}$ is made from the bingo cage, the participant knew whether or not they would have won. In contrast, in the Standard lotteries, participants only received a ticket $\{A_{it}, B_{it}, C_{it}\}$ if their round valuation led them to enter

Mechanism—To Subjects (Healy, 2018) outlines this approach.

the lottery. As such, counterfactual information was unavailable to participants who did not enter. Our second treatment dimension is over the dynamics of choice. In our *Simultaneous* treatment, participants made a single valuation decision V_i , which they knew would bind for all thirty rounds. For our *Sequential* treatment, participants submitted a new valuation V_{it} every round after the past round's lottery outcome was realized.

Table 1: Core Experimental Design

Decision Type	Standard Lottery	Regret Lottery
Simultaneous:	– Random ticket on entry	– Fixed ticket, printed at desk
	– Single valuation	– Single valuation
	– $N = 80$ participants	– $N = 80$ participants
Sequential:	– Random ticket on entry	– Fixed ticket, printed at desk
	– Sequential valuations	– Sequential valuations
	– $N = 80$ participants	– $N = 80$ participants

This design allows us to examine the following main question: do the results from the static settings generalize to the repeated setting? Participants in the Regret conditions can anticipate the regret they will feel if they do not enter the lottery but learn they would have won a large prize. In contrast, participants who do not enter in the Standard condition have no information on the realization, and thus cannot anticipate regret. In contrast, the Simultaneous and Sequential conditions are designed to examine a different anticipatory channel: whether valuations could be adapted to information over time. In the Simultaneous treatments, participants know that their first valuation commits them for thirty rounds, whereas in the Sequential treatments they can anticipate the flexibility to adapt their subsequent choices. An overly high or low valuation in round one could be changed in subsequent rounds in response to realizations. Note, however, that this also means that their valuations can respond to the differential learning opportunities available in each condition.

2.1 Model & Predictions

To model our setting and derive predictions, we consider an additively separable framework that incorporates a payoff term (based on realized outcomes) and a regret term (based on observable counterfactual outcomes). As we are interested in the potential for learning/adaptation when choices are sequential, the utility function $u(\cdot)$ for round t takes an implicit parameter ω_{t-1} , reflecting the agent's prior experiences with the lottery. Given a valuation V_t for the lottery in round t , a realized certain offer v_t , and the random prize outcome π_t , the payoff-utility for round t is given by:

$$\text{Round payoff-utility} = U(V_t; \omega_{t-1}) = \begin{cases} \mathbb{E}u(\pi_t; \omega_{t-1}) & \text{if } v_t \leq V_t \text{ (enter lottery),} \\ u(v_t; \omega_{t-1}) & \text{otherwise (take certain amount).} \end{cases}$$

The entry decision in round t is based on the agent's lottery valuation V_t , which is the maximum amount the agent will forgo in order to take part in the lottery. The function $u(\cdot; \omega_t)$ is a standard Bernoulli function, allowing for risk preferences. We implicitly assume that the function $u(\cdot; \omega_0)$ is initially more concave, reflecting both risk aversion and uncertainty over the lottery's attributes. We assume that this function becomes *less concave* in expectation as the agent gains experience with the lottery. A direct interpretation is that the agent is both risk averse and holds some subjective/cognitive uncertainty over the expectation of lottery value. While the risk aversion may remain, the subjective uncertainty over the lottery is dissipated through experience.¹²

While our specification allows for a non-standard response through time, in the absence of any regret terms the optimal entry decision in any particular round is to enter the lottery whenever the expected value satisfies $u(v_t; \omega_{t-1}) < \mathbb{E}u(\pi_t, \omega_{t-1})$. As such, the valuation V_t in this model is simply the certainty equivalent of the lottery at time t , reflecting all the available information.

On top of the round payoff utility, we also consider an additive regret term that is anticipated when making the valuation V_t . This realized regret term stems from the observed counterfactual prize outcome π_t^I , where we make the simplifying assumption that regret is driven entirely through observation of a counterfactual win on the lottery.

$$\text{Regret disutility} = R(V) = \begin{cases} 0 & \text{if } v_t \leq V \text{ (enter) or } \pi_t^I < \$2.50 \text{ (lose lottery),} \\ \mathbb{E}r(\pi_t^I) & \text{otherwise.} \end{cases}$$

Assuming that $r(\cdot)$ is convex over the three possible realizations—\$2.5, \$25 and \$250—, our regret model implies that there is no rejoicing over an ex-post correct entry decision, and that the

¹²One could think of this formulation as modeling the reduction in model uncertainty as in Routledge and Zin (2009).

negative regret effects are more substantive the greater the foregone prize. This is consistent with the standard assumption that regret is concave over the amount while also simplifying the analysis.

Aggregate preferences over the entire experimental session are modeled through the sum of the individual rounds:

$$W(\mathbf{V}; \boldsymbol{\omega}) = \sum_{t=1}^T [U(V_t; \omega_{t-1}) - R(V_t)].$$

For our Simultaneous setting, participants specify a fixed V across all of the rounds and do so in the first round with no experience (ω_0). Without counterfactual information on the lottery realizations, the regret term has no impact on decisions. This leads to a clear prediction for the Standard lotteries in the Simultaneous setting: the agent maximizes $T \cdot U(V; \omega_0)$. The optimal valuation V_{Sim}^S is therefore the agent's certainty equivalent for the offered lottery at time zero.

In contrast, for the Simultaneous regret lotteries, the valuation takes into account the possible regret in each of the thirty rounds, so that the agent maximizes $T [U(V) - R(V)]$. The aggregate anticipated regret is the product of the number of rounds the agent is expected to not enter the lottery (linearly decreasing in V) and the expected regret cost from not entering:

$$\begin{aligned} W(V) &= T \cdot U(V) - T \cdot \Pr \{ \text{Don't Enter} \} \cdot [p_1 \cdot r(\$2.50) + p_2 \cdot r(\$25) + p_3 \cdot r(\$250)], \\ &= T (U(V) - (1 - V) \cdot R_0) \end{aligned}$$

The anticipated regret $R_0 > 0$ from not entering in a round leads the agent to adjust her valuation V_{Sim}^R upwards relative to the standard lottery. This leads to our first hypothesis, that without the scope for adjusting choices in response to information, anticipated regret will increase the value for the offered lottery.

Hypothesis 1. *In the Simultaneous environment, valuations will be larger for regret lotteries than for standard lotteries: $\Delta V_{\text{Static}} = V_{\text{Sim}}^{\text{Reg.}} - V_{\text{Sim}}^{\text{Std.}} > 0$.*

As we move to the Sequential environment, we allow the ω_{t-1} term to evolve over time in response to prior realizations of π_{t-1} . As participants gain experience with the lottery, their uncertainty over the incentive structure will decrease (at least in expectation). This implies that their choices will be more in line with their underlying risk preferences over objective risk as the subjective uncertainty is dissipated. Because the environment is sequential, decision-making in period t for the standard lottery is governed by the following recursive equation:

$$W_t^{\text{Std}}(\omega_{t-1}) = \max_V \left(U(V; \omega_{t-1}) + \mathbb{E} [W_{t+1}^{\text{Std}}(\omega_t) | V, \omega_{t-1}] \right),$$

where the final-round's continuation value is $W_{T+1} = 0$.

The core assumption here is that the agents internalize a future benefit from observing each lottery realization π_t . When they choose not to enter in the standard lottery, their experience with the lottery remains constant ($\omega_t = \omega_{t-1}$). On the other hand, entry increases the continuation value because the expected future decisions (the $\mathbb{E}W_{t+1}^S$ term) benefit from the reduced subjective uncertainty, so that agents anticipate the value of learning. Entering the standard lottery therefore provides a future benefit—which is largest when the agent has the least experience—which leads the initial lottery valuations to go up.

Hypothesis 2. *For standard lotteries, anticipated learning increases initial valuations for the lottery in the Sequential environment relative to the Simultaneous environment. Namely, $\Delta V_{Lrn.} = V_{Seq.}^{Std.} - V_{Sim.}^{Std.} > 0$.*

Finally, we consider the case of regret in the Sequential environment. Since the regret lottery releases counterfactual information about the incentive structure regardless of the entry decision, there is no longer a learning motive to enter. That is, while the continuation value *will* change with the realization, this is unaffected by the entry decision, as information is provided regardless of entry. Valuation choices are governed by:

$$W_t^{Reg.}(\omega_{t-1}) = \max_V \left(U(V, \omega_{t-1}) - (1 - V) \cdot R_0 + \mathbb{E} \left[W_{t+1}^{Reg.}(\omega_t) | \omega_{t-1} \right] \right).$$

While there are differences in the details of how the optimal valuation is chosen over time,¹³ the model predicts that initial valuations in the Sequential Regret treatment are identical to those in the Simultaneous Regret treatment. The main insight is that because the regret lottery provides information on the outcome regardless of the decision, there is no anticipatory benefit from learning (as there is in the Standard Sequential treatment); the agent's information about the lottery ω_{t+1} in the next period is the same regardless of her entry decision.¹⁴

Hypothesis 3. *Initial valuations for regret lotteries in the Sequential setting will be the same as in the Simultaneous setting, $V_{Seq.}^{Reg.} - V_{Sim.}^{Reg.} = 0$. This is due to no differential learning opportunities across the settings.*

¹³In particular, the Simultaneous environment scales up the problem by repeating the same decision 30 times. In contrast, for the Sequential decision, while subsequent periods' valuation decisions may be different, if ω_{t+1} is only affected by realized information on the lottery realization, then the continuation value is invariant to the current period's entry/exit decision.

¹⁴The assumptions that the regret component is fully separable by round is admittedly a strong one. For example, the agent may exhibit less regret over not entering the lottery in a particular round and losing \$2.50 if the same decision had saved them from losing an aggregate of \$5.00 in prior rounds.

Together, Hypotheses 1, 2 and 3 imply a possible reversal of the regret effect when moving from the one-shot to the repeated setting. While the regret lottery can still generate an anticipatory increase in valuation, the presence of counterfactual information lessens this anticipatory gain in values due to the learning benefit from entering that is only present for the standard lottery. Given that the evidence for anticipatory regret in standard one-shot experimental settings is relatively weak, our final hypothesis formalizes the overall effect by assuming the sequentiality effect (Hypothesis 2) dominates.

Hypothesis 4. *Initial valuations will be lower for Regret lotteries in the Sequential setting than for Standard lotteries.*

Note that our model is meant to be illustrative rather than comprehensive. It is designed to highlight and structure a potential mechanism for a reversal of the regret effect when moving from a static to a dynamic setting. It is not meant to capture all potential interactions between regret and the various forces that are introduced with repeated choice.


3 Main Results

We now present the main results of the lab experiment, starting with the average valuations by treatment. We then turn to an examination of our four hypotheses, using both pairwise comparisons motivated by the theory and in a series of joint regressions. As per the pre-registration, our main analyses combine data from the first and second waves of lab experiments.¹⁵ As a preview, while we do find a small positive impact of regret in the static setting (Hypothesis 1), this effect is not statistically significant. In contrast, in the dynamic setting we find a large and significant decline in initial valuations for the regret lottery relative to the standard lottery (Hypothesis 4). The mechanics parallel the theory, where a comparison of the initial valuation for the standard lottery in the sequential and simultaneous settings reveal a (quantitatively large) increase in valuations (Hypothesis 2). While we still find a sequentiality premium within regret lotteries (against Hypothesis 3), this premium is smaller than in the case of standard lotteries.


3.1 Average results

Figure 1 illustrates our core results. In the left panels we illustrate the lottery valuations by lottery type and decision environment. Panel (a) uses data from the very first valuation, while Panel (c)


¹⁵See Section D.1 in the Online Appendix for a comparison of both waves. Crucially, the main findings are all directionally replicated in second wave, where and we cannot reject differences in joint tests of equality across the two waves.




(a) Average valuation (first round)



(b) Effects (first round)



(c) Average valuation (all rounds)



(d) Effects (all rounds)

Figure 1: Average lottery valuations and regret premium, all rounds

looks at the average across all rounds—where this distinction only affects the Sequential treatments. For ease of interpretation, each valuation is re-scaled to a relative valuation using the expected prize from the lottery.¹⁶ As such, a 100 percent valuation (\$0.61) would be the risk-neutral choice, and risk aversion is indicated by a relative valuation below 100 percent.

In the Simultaneous environment, the average valuation is equal to 58.5 percent of the lottery’s expected value for the Standard lottery and 64.5 percent for the Regret lottery. Per Hypothesis 1, this represents a 6 percentage-point valuation increase for the Regret lottery. A direct comparison between the Standard and Regret lotteries in the Simultaneous treatment is illustrated in the first bar in Figure 1 panel (b). This directional difference is, however, not statistically significant ($p = 0.304$ for a paired t -test, $p = 0.530$ for a Mann-Whitney test).

As we move to the Sequential treatments, we observe an across-the-board increase in valuations—which we refer to as the *sequentiality premium*. Panel (a) of Figure 1 illustrates the relative valuations in the very first round for both the Standard and Regret lotteries. The Standard lottery is initially valued at 92.3 percent of the lottery’s expected value, a sharp increase from the Simultaneous setting (58.5). A smaller, though still economically large, increase is observed for the Regret lotteries, with valuations going from 64.5 to 81.3 percent. Panel (c) of Figure 1 shows the same comparison, but now averaging valuations across all thirty rounds. The Standard lottery’s average valuation stands at 88.0 percent of the lottery’s expected value, while the Regret lottery’s valuation is at 83.5 percent.

The significant sequentiality premium is illustrated in the second bar of Panel (b) of Figure 1, indicating the pairwise comparison from our second hypothesis.¹⁷ Examining Hypothesis 2, we can compare the Standard lottery’s valuation in the Sequential and Simultaneous environments. The treatment difference represents a 34 percentage-point increase ($p < 0.001$, both paired- t and Mann-Whitney tests), which the theory indicates is driven by the flexibility in adjusting future decisions in response to feedback.

Finally, Figure 1 Panel (b) also displays the initial valuations of the Regret lottery in the dynamic setting, which increase by 16.8 percentage points when we compare the Simultaneous and Sequential treatments ($p = 0.004$ paired t , $p = 0.003$ Mann-Whitney). While the sequentiality effect is halved for the Regret lotteries compared to Standard ones, the premium is still substantial. This leads us to reject Hypothesis 3.

Because regret lotteries also get a partial sequentiality premium, this might wash out the regret reversal predicted in Hypothesis 4 when we compare the two sequential implementations.

¹⁶That is, for each treatment we transform the raw valuation V_t^i to a relative valuation $\hat{V}_t^i = \frac{V_t^i}{\mathbb{E}\pi}$.

¹⁷For completeness, we also provide the same pairwise treatment comparisons in Panel (d) of Figure 1 using the averages across all rounds. However, we note that Hypotheses 2 and 3 are developed over the very first round.

The final bar in Panel (b) of Figure 1 indicates the difference in first-round valuation between Standard-Sequential and Regret-Sequential lotteries and shows this is not the case. Here we find the hypothesized reversal, where Regret lotteries decrease the valuation by 11 percentage points ($p = 0.040$ paired- t , $p = 0.025$ Mann-Whitney) over Standard. The anticipatory effect of learning in the Sequential treatments (Hypothesis 2) dwarfs the anticipated regret effect in the Simultaneous lotteries (Hypothesis 1). While we do see a substantive gain for the Sequential Regret lottery over the Simultaneous version, our results suggest that Regret lotteries are less effective than Standard lotteries in the Sequential setting.

Looking at valuations across all rounds (Figure 1, panels c and d) we see a very similar pattern of results. While the regret effect in the Sequential treatment (Hypothesis 4) is no longer significantly different from zero, we can reject even a very small positive effect of regret. For example, a Wald test of the estimated regret effect in sequential against the point estimate of the effect in the static treatment yields a p -value of 0.001 ($F_{1,159} = 6.89$).

We have thus far compared the valuations in the treatment pairs specified by the hypotheses. Table 2 reports results on the hypothesis tests from a joint regression over all four treatments. We do this for both the first round (which our hypotheses are focused on) and across all rounds to study the longer-run effects. In addition to the difference in averages (columns I and III) we also provide results where we control for the gender of the decision-maker.¹⁸ The results from the joint regression mirror those from the pairwise comparisons, though controlling for gender increases the precision of the coefficient estimates.

Across the pairwise treatment comparisons—using both paired- t , Mann-Whitney, and regression-based tests—the takeaways are the following:

Result 1 (Static regret effect). *Testing Hypothesis 1, while the directional effects do move in the predicted direction, we do not find strong evidence that regret increases lottery valuations.*

Result 2 (Anticipated learning effect). *Our experiments provide strong evidence for Hypothesis 2. When participants know that they will be able to adjust their choices to information, the valuation for the standard lottery increases substantially.*

Result 3 (Regret null in sequential). *In contrast to Hypothesis 3, we find a significant sequentiality premium in the Regret treatments.*

¹⁸The control serves to reduce the unobserved variation and is not interacted with treatment. Women consistently value the lottery less than men in all of our treatments. The average estimated effect indicates that women's relative valuations in the first (all) rounds is 11.4 (15.4) percentage-points lower than men's valuations across the four treatments. While we also have other controls (college major, the risk elicitation in the first part) none of them provide a substantive reduction in the variance.

Result 4 (Dynamic regret reversal). *Our results are consistent with Hypothesis 4, showing that regret lotteries reduce valuations for the lottery incentive in the dynamic setting.*

Table 2: Regression-based hypothesis tests

	First round		All rounds	
	(I)	(II)	(III)	(IV)
Hypothesis 1: ($\hat{V}_{\text{Sim.}}^{\text{Reg.}} - \hat{V}_{\text{Sim.}}^{\text{Std.}}$)	6.0 ($p = 0.304$)	6.8 ($p = 0.237$)	6.0 ($p = 0.302$)	6.8 ($p = 0.235$)
Hypothesis 2: ($\hat{V}_{\text{Seq.}}^{\text{Std.}} - \hat{V}_{\text{Sim.}}^{\text{Std.}}$)	33.8 ($p < 0.001$)	34.6 ($p < 0.001$)	29.6 ($p < 0.001$)	30.6 ($p < 0.001$)
Hypothesis 3: ($\hat{V}_{\text{Seq.}}^{\text{Reg.}} - \hat{V}_{\text{Sim.}}^{\text{Reg.}}$)	16.8 ($p = 0.004$)	16.1 ($p = 0.006$)	19.0 ($p < 0.001$)	18.1 ($p = 0.001$)
Hypothesis 4: ($\hat{V}_{\text{Seq.}}^{\text{Reg.}} - \hat{V}_{\text{Seq.}}^{\text{Std.}}$)	-11.0 ($p = 0.040$)	-11.4 ($p = 0.032$)	-4.5 ($p = 0.325$)	-5.6 ($p = 0.206$)
Gender control	<i>No</i>	<i>Yes</i>	<i>No</i>	<i>Yes</i>
N	160	160	4,800	4,800

Note: Significance tests p -values are given in parentheses. Coefficients in columns (I) and (II) are derived from OLS regressions of lottery valuations using only the first-round valuation (one observation per participant). Columns (III) and (IV) report results from estimations allowing for a participant random-effect in the sequential treatments (4,800 sequential choices, compared to 160 for the simultaneous treatments).

3.2 Dynamics

Our framework assumes that agents have some subjective uncertainty about the lottery incentive, and that an ability to learn and respond can generate additional value. However, we also predict that participants will respond to experienced outcomes in the Sequential treatment, and move their valuations in subsequent rounds. In this section, we provide evidence that these learning mechanisms are at play in the data via the dynamic responses.¹⁹

Figure 2 presents the evolution of average relative valuations across the 30 rounds, with each average representing a 5-round block (400 total observations per point). The figure makes clear that both the sequentiality premium and the regret reversal hold across the 30 decisions. Clearly,

¹⁹While the evidence does point to the learning channel, related motives are also possible. For example, it is possible that an increased curiosity is at play, with a desire to find out the assigned numbers. However, such arguments would have to generate an interaction with sequentiality, as curiosity could equally apply in the static setting.

the gap in valuations between Standard and Regret lotteries becomes smaller in the later rounds, though this closing of the gap would also be predicted by the model. That is, more information on lottery realizations is available for the regret lotteries (29 realizations with certainty at the last round) than for the standard (number of realization equal to the number of rounds the participant entered). Indeed, once all learning is concluded, the model would predict that in the limit the Regret lotteries would again have greater value than the standard, as the R_0 term will remain after all learning has taken place.

To examine the extent to which the participants' experience with the lottery are indeed shifting valuations, in Table 3 we present results from random-effects regressions where we explain the round-to-round changes in valuations. The dependent variable in these regressions is the change in lottery valuation relative to the previous round, $V_t - V_{t-1}$. On the right-hand side of the regression equation we include indicators for realized gains and losses and, for the regret treatment, also indicators for counterfactual gains and losses.




Figure 2: Average lottery valuations (five-round blocks)

The results indicate a strong negative reaction to entering and losing the lottery in the previous round (the modal outcome), estimated to move the valuations downward by approximately 5 (5.5) percentage-points in the Standard (Regret) lottery treatment, where this move is significant in all estimates ($p < 0.001$). In contrast, there is no significant response when the lottery is entered and won, where the participants on average choose a similar lottery valuations than the previous

round. The final three rows represent the response to not entering the lottery. In both the standard and regret lotteries, participants increase their valuations significantly ($p < 0.001$), by about 5 percentage points on average.

Table 3: Experience effects on lottery valuations

Last Rd. Outcome	Change in relative valuation ($\hat{V}_t - \hat{V}_{t-1}$)		
	Standard	Regret	
		(I)	(II)
Entered, Loss	-4.9 (0.4)	-5.4 (0.6)	-5.4 (0.6)
Entered, Win	0.0 (0.9)	0.0 (1.3)	0.0 (1.3)
Did Not Enter	4.5 (0.4)	5.2 (0.5)	
Counterfactual Win			8.1 (1.2)
Counterfactual Loss			4.6 (0.5)
<i>N</i>	2,320	2,320	2,320

Note: Standard errors given in parentheses. Each column represents a random-effects regression estimate on the change in the relative valuation across the 80 participants in each sequential session over 29 separate rounds.

While there is no further information available to participants that do not enter the Standard lottery, a core mechanic in the theory is that regret lotteries allow for counterfactual learning without entry. To examine this, in the last column of the table, we separately estimate the dynamic response to not entering the lottery when the last round's opt out decision was ex-post the 'correct' thing to do (a counterfactual loss), or ex-post 'incorrect' (a counterfactual win). We find a statistically significant valuation gap of 3.5 percentage points in favor of a counterfactual win ($p = 0.001$), though smaller than the gap in response to a realized win or loss (5.4 percentage-points, $p < 0.057$ each).²⁰ The conclusion from this analysis suggests that, while counterfactual information is used, information from realized outcomes is more important.²¹

²⁰In one difference across the waves, in the first sample, we noted that participants responded differentially to larger wins (\$25 vs \$2.50). However, we do not replicate this finding in the pooled sample.

²¹This also suggests a potential reason for why we observe a larger gap between the Simultaneous and Sequential regret lotteries, where Hypothesis 3 indicates we should expect a null effect.

3.3 Robustness and Extensions

Our results outline that the use of regret lotteries in repeated settings may have unintended consequences. We show that the anticipatory force of learning from counterfactual information decreases valuations of the regret lottery relative to the standard lottery; that is, the core mechanism of a regret lottery (revealing the realizations regardless of engagement) creates a perverse incentive effect in the sequential environment. This leads the regret lottery to have a negative net effect in the sequential implementation. In this section, we report on a series of robustness exercises involving (a) design changes in how the notion of regret was implemented, (b) an examination of the impact of different learning opportunities in a practice stage, and (c) an attempted online replication. The main take-away across all of these iterations is the sensitivity of regret to implementation details. We now consider each one in turn.

3.3.1 Changes to the regret implementation

In our core Regret treatments, participants are given a fixed lottery ticket at the very beginning of the session. Much like a postcode, social-security number, date of birth, or an employee ID, this ticket value was then held constant across all of the rounds. In this section, we consider a series of alternative implementations of the regret lottery, as follows:

1. A *Social Regret* ticket, mirroring the idea of a postcode lottery. Here, two groups of five participants were given a common ticket. While the ticket was printed on their individual desk and appeared on their screens as before, a larger print-out was positioned above each group for all to see. As such, regret might plausibly have greater power through social pressure by, for example, not entering the lottery but knowing that others in your group ended up winning ($N = 30$ participants).
2. A *Choice Regret* ticket. Here, participants themselves choose the three entry numbers for their ticket at the beginning of each round before choosing a valuation (and so no ticket was printed at their desks). Allowing participants to choose their own entry ticket allowed them to select ‘lucky’ or personally significant numbers, as allowed for in many state/national lotteries ($N = 30$ participants).
3. A *Random Regret* ticket, where each participant was *assigned* a fully randomized ticket at the beginning of the round (where again, no ticket was printed at their desks). Here participants were given three randomly drawn numbers for their ticket at the beginning of each round, and then asked to provide a valuation. The treatment therefore removes the persistent realization

common to many lotteries ('lucky' numbers, a drivers' license number, or a zip code) and replaces it with a random draw ($N = 46$ participants).

Given the independence of the ticket assignment and the draws from the bingo cage, all three alternatives offer identical expected values, possibilities for anticipatory regret, and learning about the lottery through the free counterfactual information.²² As such, there are no clear theoretical predictions on these alternative implementations, as none have elements that would distinguish them in the model from the standard Sequential Regret treatment. Rather, the aim of these treatments is exploratory, to study the robustness of the regret component to implementation details in the design. Results for these three robustness treatments are illustrated in Figure 3.

Looking at the initial behavior, the results indicate similar qualitative responses to the core Regret treatment for both the Choice Regret and Social Regret implementations. First round valuations are not significantly different from the original Regret lotteries with a fixed ticket ($p > 0.500$ for all paired- t and Mann-Whitney comparisons), but are significantly different from the Standard lottery ($p < 0.05$ for paired- t and $p < 0.057$ for Mann-Whitney tests). As Figure 3 makes clear, the time-path across the session for the Choice Regret lotteries is nearly identical to that of the core Regret experiment. For the Social Regret implementation, however, the late-session behavior is distinct. In fact, looking at all the rounds, behavior in the Social Regret treatment is much more in line with Hypothesis 3, where there is no substantive sequentiality premium over the Simultaneous Regret treatment ($p = 0.423$), while there is a significant sequentiality premium for all of the other regret treatments ($p < 0.040$). While our results here cannot speak to the specifics of the mechanism, the experiment suggests that a fixed joint realization across an entire group (as in the postcodes lottery) is the least effective sequential implementation for a regret lottery.

The one outlier is the Random Regret treatment, where an exogenously assigned ticket is randomly provided before each valuation decision. For this treatment we do not detect a significant difference in valuations relative to the Standard-Sequential implementation, neither for initial valuations ($p = 0.843$) nor over all rounds ($p = 0.571$). As such, while there is no behavioral benefit from regret in the Random Regret implementation over a Standard lottery, this treatment also yields no significant penalty.

Together, our results suggest that moving from a static to a dynamic environment generally either eliminates the effectiveness of regret in increasing valuations or can actually *lower* them.²³

²²Note, however, that opportunities for learning are slightly increased in the Social Regret implementation, since the other groups' ticket is also commonly visible.

²³While the results in Random-Regret do suggest that the pernicious effect of regret in the sequential setting are related to persistence of the ticket—the ecologically valid case with zip codes, employee numbers, SSNs, etc—it still underlines the core idea of the paper. Once we move outside of the static setting, the best case for anticipatory regret

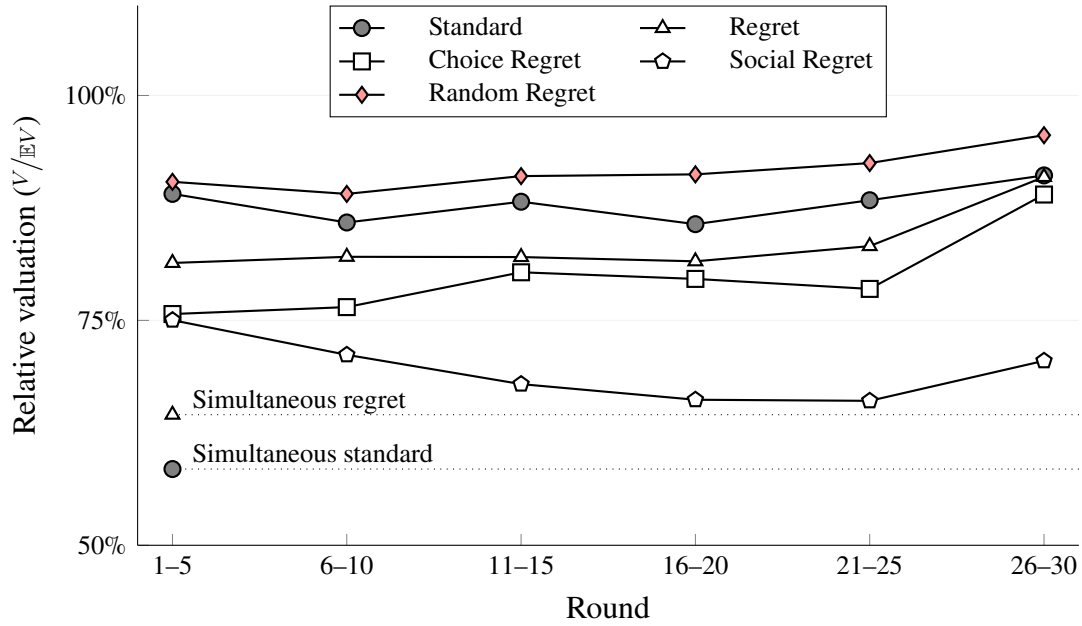


Figure 3: Sequential Robustness Treatments

3.3.2 Changes to the practice stage

As stated in Hypothesis 2, the option value of learning is one reason why valuations of standard lotteries might be higher in sequential compared to the simultaneous decisions. Another reason, also related to learning, has to do with the practice session prior to the experiment.²⁴ After reading the instructions, but before starting the experiment, participants were given 2 minutes to practice with the lottery-valuation task. The screen allowed them to select different values on a slider and to click a button to generate random offers.²⁵ For each value chosen and random offer generated, participants learned if they would have entered the lottery or earned the offer amount.

In an effort to closely match it to the remainder of the experiment, the practice stages were different in the Sequential and Simultaneous treatments. In Sequential treatments, a single random offer was generated at each click, and a unique decision—lottery entry or not—was shown. In Simultaneous treatments, on the other hand, one click of the offer-generating button resulted in 30 random offers, and participants were given feedback on lottery entry for all 30 rounds. Even though there were no limits on the number of random offers that could be generated, this difference in design could have resulted in different learning experiences. For example, consider a participant that selected a low value in the practice stage and generated three random offers. In Sequential

in a dynamic environment is that it does not actively reduce the incentive effects.

²⁴We thank an anonymous referee for highlighting this channel.

²⁵See Online Appendix Section B.1 for screenshots of the practice stage.

treatments, this participant would have a high probability of not entering the lottery in all three instances. For Simultaneous treatments, however, the same low value would almost certainly result in a few entries over the 30 randomly-generated offers. If participants in Sequential treatments (incorrectly) conclude that chances of entering even a few times are very low, then they would have an incentive to choose higher valuations in initial rounds.

To test this idea, we ran additional sessions of the Sequential treatments with a modified practice stage, where participants are given feedback about lottery entry over 30 rounds based on 30 randomly generated offers. We collected data on an additional 60 subjects, half for Regret and half for Standard lotteries. In line with the idea that extended feedback reduces uncertainty, and hence the option value from engaging, first-round valuations for Sequential treatments were lower for both the Standard (74% versus 92% of EV; $p = 0.017$ two-sided t-test) and Regret (67% versus 81% of EV; $p = 0.070$ two-sided t-test) lotteries. Differences in valuations over all rounds are much smaller and not statistically significant, again both for Standard (83% versus 88% of EV; $p = 0.465$) and Regret (80% versus 83% of EV; $p = 0.677$) lotteries.

Though most of the main results from Table 2 are replicated considering the treatments with the modified practice rounds, the effect sizes are smaller.²⁶ The results from this robustness check highlight the fact that, besides the importance of considering static versus dynamic environments, other changes in implementation and framing can also affect the relative valuation of Regret versus Standard lotteries in both static and dynamic environments.

3.3.3 Online replication

As part of an effort to explore the generalizability of the regret effects we documented in the lab, we pre-registered a large-scale replication ($N = 981$, including additional extension treatments) using the online platform Prolific. The design of the experiment was similar to the original lab study, but with a reduced number of rounds (from 30 to 20) and smaller incentives for the online sample (\$1.50, \$15, and \$150 for one, two, and three matches respectively). While many noteworthy behavioral effects have replicated across populations, including in online samples (Gupta et al., 2021; Ruggeri et al., 2020), our online results indicate unambiguous null effects for regret. Figure 4 depicts average valuations in our four main treatments, where panel (a) reports on first-round valuations and panel (b) on average valuations across all twenty rounds. As Figure 4 makes clear, there are no substantive valuation differences between Standard and Regret lotteries, neither in the Simultaneous nor the Sequential treatments. That is, we find a complete absence of any evidence in our online sample for: (i) static regret aversion, (ii) the dynamic reversal, (iii) nor any differences in

²⁶See Online Appendix Table D.3 for details.

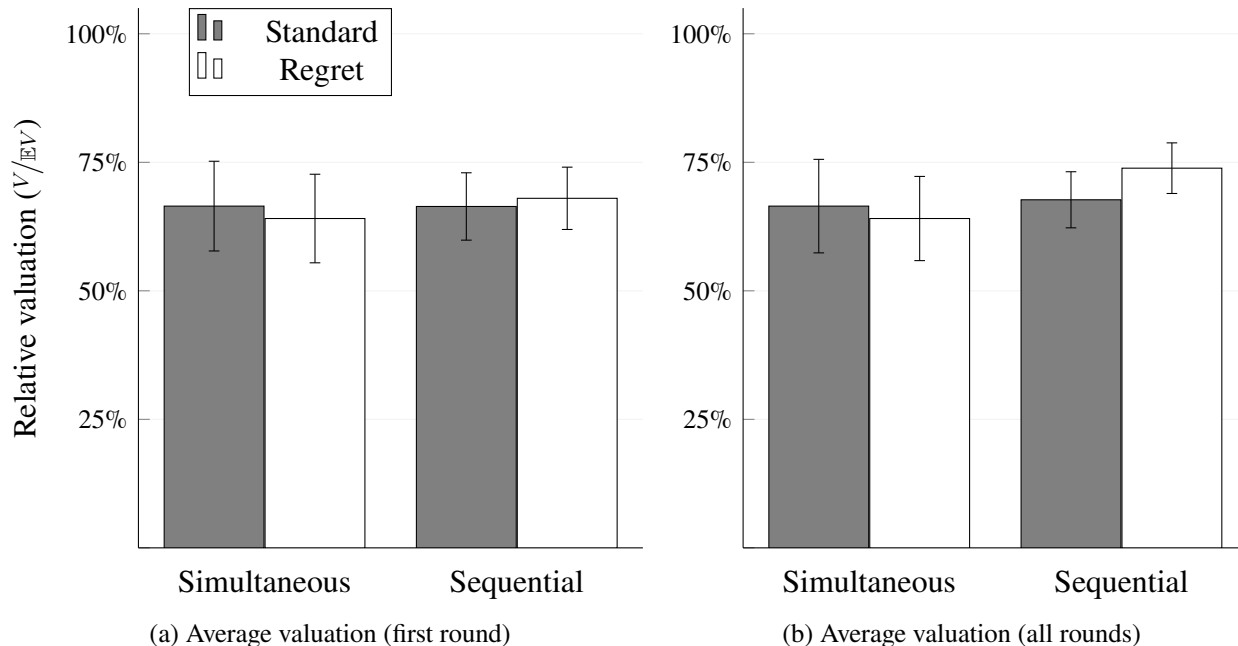


Figure 4: Average lottery valuations and regret premium with data from online experiments

the valuations between a single valuation applied to 20 rounds, or across 20 sequential valuations with feedback.

On average, participants value the lottery at 65% of the expected value. Moreover, the online sample fails to replicate the largest effect from the lab sample, the sequentiality premium for standard lotteries. Table D.4 in the Online Appendix provides regression coefficients on our four main hypotheses—similar to Table 2 in the paper. However, the only result that replicates at the 5 percent level is the sequentiality premium over all rounds for regret lotteries; all other predictions are well-powered nulls effects in the Prolific sample.

We also used the online laboratory to replicate four of our previous robustness treatments from the lab and to conduct a new treatment to test for mechanisms.²⁷ Specifically, we ran sessions of the Sequential-Regret treatment where tickets were chosen by the participant in each round (*Choice Regret*) and sessions in which participants were assigned a new random ticket at the start of each round (*Random Regret*). With the exception of first-round valuations in the Regret Choice treatment, coefficients are not statistically different than the Sequential-Regret treatment. We also collected data for both Sequential treatments (i.e., Regret and Standard) with a modified practice stage (as discussed in Subsection 3.3.2). As with the other robustness checks, the extension results

²⁷Section D.3 in the Online Appendix reports regression results on the robustness and the additional treatment for the online sessions.

were statistically indistinguishable from the core treatments (Table C.4 in the Online Appendix).

Lastly, we conducted a version of the Sequential-Standard treatment in which entry outcomes are realized each round (just as in the original standard sequential treatment), but where lottery results (ticket numbers, winning numbers, and matches) are not realized until after the last round (round 20).²⁸ However, mirroring the null results in the other comparisons of online data, we find that first-round lottery valuations are not statistically different ($p = 0.205$).

4 Discussion & Conclusion

Academics, policy makers, and practitioners have increasingly attempted to import findings from laboratory studies to incentive design. One of the most prominent examples is the use of regret to induce behavior change (Milkman et al., 2021; Volpp et al., 2008b). Our results suggest caution in this extrapolation. While we directionally replicate the finding that anticipated regret can increase the efficacy of a lottery incentive in a static setting—although the effect is small and not statistically significant—this effect is significantly reversed in a repeated setting. Consistent with a model that considers the value from anticipated learning about the incentive structure, we find that standard lottery valuations are much higher in the dynamic context; which we term the “sequentiality premium.” However, because the same information is provided without having to enter the lottery in the regret case, this leads to a relative valuation drop in the dynamic setting. This is exactly what we find.

Our findings also highlight the sensitivity of the effects based on regret to the particular details of the implementation. For example, in our Regret-Random specification, where participants are given a new random lottery ticket every round instead of a fixed ticket at the beginning, results are much more similar to those for standard lotteries. Moreover, in settings where attention is likely to be lower, as in our online studies, regret effects are greatly diminished.

In providing guidance to policymakers on incentive design, a first-order question is whether lotteries should be used at all. Risk-aversion over the offered lotteries suggests that paying a constant non-stochastic payment may be preferable to a lottery incentive in many situations. This would be the case in our data, where, in all of the treatments, the average valuations fall below the lottery’s actuarial value. This is also supported by the contrasting findings of Milkman et al. (2021) and Campos-Mercade et al. (2021), who provided regret lotteries and fixed payments, respectively, as

²⁸This new treatment was devised as an out-of-sample test of the anticipated-learning hypothesis. Specifically, if anticipated learning were an important channel, we would see a reduction in first-round valuations in the new treatment compared to first-round valuations in the original Standard-Sequential treatment, as the anticipated-learning channel was removed.

incentives for vaccine take-up. While the latter found a positive effect of incentives, the former did not.

However, even where lotteries are valued less than their actuarial cost, policymakers might still want to deploy them. For example, if making payments to every individual has large administrative costs, lotteries can help reduce these as only a few receive payment. In situations where lotteries are the focal incentive, our results provide guidance on implementation options. First, if repeated incentives need to be used for an ongoing choice, our findings imply standard rather than regret lotteries may in fact be the superior tool. Second, where policymakers have options over both an ongoing dynamic implementation and a static one (for example, a single incentive at the end of a prolonged period versus many smaller incentives at the end of each sub-period), our results suggest a substantial “sequentiality premium.”

Our findings caution against the untested extension of behavioral phenomena from static to dynamic contexts. Learning, risk aversion over final wealth, and lottery realizations can all contribute towards making a sequence of decisions interrelated, even when the decisions are statistically independent. In some settings, this non-separability might exacerbate behavioral effects, and in others it may attenuate them. As in our case study of regret, it can even reverse the direction of the comparative static. Our paper therefore provides a clear example of a worst-case scenario from an incentive-design perspective, such that the direction of the policy’s intended behavioral effect is entirely reversed in the dynamic setting. That being said, our results should also be interpreted with caution: our laboratory environment may abstract from factors that are important for activating regret in the field, and more work is needed to examine the ecological validity of our findings. Finally, an interesting area for future research would be to explore the interaction between uncertainty and the anticipated learning channel. In particular, do we observe a smaller regret effect in dynamic settings where the decisions are less uncertain, and hence where there is a lower learning value, compared to situations with higher uncertainty?

References

- Babcock, Linda, William J Congdon, Lawrence F Katz, and Sendhil Mullainathan**, “Notes on behavioral economics and labor market policy,” *IZA Journal of Labor Policy*, 2012, 1 (2), 1–14.
- Bell, David E.**, “Regret in Decision Making under Uncertainty,” *Operations Research*, 1982, 30 (5), 961–981.
- Brown, Alexander L, Taisuke Imai, Ferdinand M Vieider, and Colin F Camerer**, “Meta-analysis of empirical estimates of loss aversion,” *Journal of Economic Literature*, 2024, 62 (2), 485–516.
- Campos-Mercade, Pol, Armando N Meier, Florian H Schneider, Stephan Meier, Devin Pope, and Erik Wengström**, “Monetary incentives increase COVID-19 vaccinations,” *Science*, 2021, 374 (6569), 879–882.
- Datta, Saugato and Sendhil Mullainathan**, “Behavioral design: a new approach to development policy,” *Review of Income and Wealth*, 2014, 60 (1), 7–35.
- Erev, Ido and Greg Barron**, “On adaptation, maximization, and reinforcement learning among cognitive strategies,” *Psychological review*, 2005, 112 (4), 912.
- Gupta, Neeraja, Luca Rigotti, and Alistair Wilson**, “The experimenters’ dilemma: Inferential preferences over populations,” *arXiv preprint arXiv:2107.05064*, 2021.
- Haisley, Emily, Kevin G Volpp, Thomas Pellathy, and George Loewenstein**, “The impact of alternative incentive schemes on completion of health risk assessments,” *American Journal of Health Promotion*, 2012, 26 (3), 184–188.
- Hart, Sergiu**, “Adaptive heuristics,” *Econometrica*, 2005, 73 (5), 1401–1430.
- Healy, Paul J**, “Explaining the BDM-Or any random binary choice elicitation mechanism-To Subjects,” Technical Report, mimeo 2018.
- Janis, Irving L and Leon Mann**, *Decision making: A psychological analysis of conflict, choice, and commitment.*, Free press, 1977.
- Keltner, Dacher and Jennifer S Lerner**, “Emotion,” in “Handbook of social psychology” 2010.
- Kessler, Judd B and C Yiwei Zhang**, “Behavioral Economics and Health,” in “Oxford Textbook of Public Health” 2014.
- Kimmel, Stephen E, Andrea B Troxel, George Loewenstein, Colleen M Brensinger, Jane Jaskowiak, Jalpa A Doshi, Mitchell Laskin, and Kevin Volpp**, “Randomized trial of lottery-based incentives to improve warfarin adherence,” *American Heart Journal*, 2012, 164 (2), 268–274.
- Ku, Gillian**, “Learning to de-escalate: The effects of regret in escalation of commitment,” *Organizational Behavior and Human Decision Processes*, 2008, 105 (2), 221–232.
- Loomes, Graham and Robert Sugden**, “Regret theory: An alternative theory of rational choice under uncertainty,” *Economic Journal*, 1982, pp. 805–824.
- and —, “Testing for regret and disappointment in choice under uncertainty,” *Economic Journal*, 1987, pp. 118–129.
- Madrian, Brigitte C**, “Applying insights from behavioral economics to policy design,” *Annual Review of Economics*, 2014, 6, 663.
- Marchiori, Davide and Massimo Warglien**, “Predicting human interactive learning by regret-driven neural networks,” *Science*, 2008, 319 (5866), 1111–1113.
- Milkman, Katherine L, Linnea Gandhi, Sean Ellis, Heather Graci, Dena Gromet, Rayyan Mobarak, Alison Buttenheim, Angela Duckworth, Devin G Pope, Ala Stanford et al.**, “A 500,000-Person Experiment Evaluating the Impact of Geographically-Targeted, High-Payoff Vaccine Lotteries,” *High-Payoff Vaccine Lotteries (August 13, 2021)*, 2021.
- Routledge, Bryan R. and Stanley E. Zin**, “Model uncertainty and liquidity,” *Review of Economic Dynamics*, 2009, 12 (4), 543–566.
- Ruggeri, Kai, Sonia Alí, Mari Louise Berge, Giulia Bertoldo, Ludvig D Bjørndal, Anna Cortijos-Bernabeu, Clair Davison, Emir Demić, Celia Esteban-Serna, Maja Friedemann et al.**, “Replicating patterns of prospect theory for decision under risk,” *Nature human behaviour*, 2020, 4 (6), 622–633.
- Strack, Philipp and Paul Viefers**, “Too Proud to Stop: Regret in Dynamic Decisions,” November 2015. Berkeley working paper.
- Volpp, Kevin G, George Loewenstein, Andrea B Troxel, Jalpa Doshi, Maureen Price, Mitchell Laskin, and Stephen E Kimmel**, “A test of financial incentives to improve warfarin adherence,” *BMC Health Services Research*, 2008, 8 (1), 1.

- , **Leslie K John, Andrea B Troxel, Laurie Norton, Jennifer Fassbender, and George Loewenstein**, “Financial incentive-based approaches for weight loss: a randomized trial,” *Journal of the American Medical Association*, 2008, 300 (22), 2631–2637.
- Zeelenberg, Marcel**, “Anticipated regret, expected feedback and behavioral decision making,” *Journal of Behavioral Decision Making*, 1999, 12 (2), 93.
- **and Rik Pieters**, “Consequences of regret aversion in real life: The case of the Dutch postcode lottery,” *Organizational Behavior and Human Decision Processes*, 2004, 93 (2), 155–168.
- **and —** , “A theory of regret regulation 1.0,” *Journal of Consumer Psychology*, 2007, 17 (1), 3–18.