REVERSALS BETWEEN ONE-SHOT AND REPEATED DECISIONS IN INCENTIVE DESIGN: THE CASE OF REGRET

ALEX IMAS, DIEGO LAMÉ, AND ALISTAIR J. WILSON

Abstract. We demonstrate potential pitfalls when extrapolating behavioral findings identified from one-shot choices to repeated settings. As a case study, we examine the use of "regret lotteries" as a behavioral tool to boost motivation. Based on findings from one-shot settings, presenting counterfactual information generates the potential for regret, which can be used to increase a lottery's value. This result has motivated the increasing use of "regret lotteries" in the field to incentivize recurrent decisions like exercise and compliance with company directives. Using a controlled experiment we show that while regret lotteries are the superior motivational tool in one-shot decisions, for repeated decisions the effect is entirely *reversed*. These findings have implications for incentive and policy design, highlighting the scope for error when extrapolating one-shot findings to inherently repeated settings.

1. Introduction

Though many important choices are one shot, a large fraction of meaningful decisions are recurring. Many of these smaller day-to-day activities are often the target of interventions by policy makers and managers. For example, choosing whether to comply with a desired activity such as engaging with a workplace wellness program, completing a customer-experience survey, or complying with a new pro-environment cost-saving initiative. All of these are best-thought of as ongoing, recurring decisions, and as these activities are desirable from a welfare or organizational perspective, they have been the frequent target of incentive programs (see, for example, Jones et al. 2019, Iyer and Kashyap 2007, and Gosnell et al. 2020). Yet many of the findings in behavioral economics that inform incentive design come from studies conducted in one-shot, static settings, where anticipatory features of choice are paramount. However, recurrent decisions typically introduce feedback and opportunities for learning that are absent in the case of one-shot choice. Individuals may adjust their behavior both in anticipation of feedback

Date: August, 2020.
Imas: University of Chicago, Booth School of Business, 5807 S. Woodlawn Ave, Chicago, IL; alex.imas@chicagobooth.edu. *Lamé*: University of Pittsburgh, Department of Economics, 230 Bouquet Street, Pittsburgh, PA; djw63@pitt.edu. *Wilson*: University of Pittsburgh, Department of Economics, 230 Bouquet Street, Pittsburgh, PA; alistair@pitt.edu. Our thanks to the following for their feedback: Linda Babcock, Julie Downs, Uri Gneezy, Judd Kessler, PJ Healy, David Huffman, George Loewenstein, Graham Loomes, Ryan Oprea, Erkut Ozbay, Lise Vesterlund and Stephanie Wang.

(Gneezy and Potters, 1997) and in response to it (Cohen et al., 2007). As such, though well-identified for the one-shot settings, these findings may miss behavioral responses that play a significant and ecologically valid role in repeated-choice settings. In this paper, we demonstrate the pitfalls in extrapolating from findings in a one-shot setting to a recurrent one, showing that anticipation and feedback interact to fully *reverse* the policy-relevant directional effect of one psychological force that has been leveraged to increase the motivational effects of incentives—regret aversion.

Regret theory (Bell, 1982; Loomes and Sugden, 1982) proposes a simple modification to expected-utility (EU) theory through which individuals derive utility not only directly from choice outcomes, but also indirectly through their knowledge of counterfactual outcomes. 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 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 one may experience negative feelings of regret if provided with information on a positive realization of the lottery; anticipating this potential regret increases her ex-ante valuation of the lottery. The theoretical consequence for decision-making is that individuals may opt away from EU-maximizing choices and instead select options that minimize prospective regret.

A recent body of work has sought to utilize these forces as a non-pecuniary boost to incentives through the use of *regret lotteries* (Zeelenberg and Pieters, 2007; Volpp et al., 2008b). Regret lotteries differ from standard lotteries 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 are first assigned fixed winning states (an employee number, a zip code, *etc.*). They then make choices on whether to enter the lottery through their compliance with the incentivized activity (for example, exercising, car-pooling to work, taking a prescribed drug). 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* compliance then a corresponding prize is awarded. If either she did not embrace the incentivized choice or her particular winning state was not drawn, then no prize is awarded. Such lotteries seek to exploit regret aversion through the provision of counterfactual feedback regardless of the entry decision. A regret-averse decision-maker, knowing that feedback will be provided regardless of the entry choice, acts as if the lottery incentive has a larger prize. The anticipated regret from not entering and missing out on a large prize can push those on the margins to comply with the incentivized activity where they otherwise would not.

A prominent example from the field is the Dutch postcode lottery, designed to incentivize the collection of public revenues. To realize the lottery outcome, a winning postcode is drawn from the entire population of Dutch postcodes. Each individual living within the drawn postal region that bought a ticket wins a large cash prize. Those in the winning neighborhood that did not buy tickets do not get prizes but do observe that their postcode was chosen and 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.

The argument for the increased incentive power of regret lotteries is based on a theoretical and empirical literature on anticipated regret, which focuses on one-shot decision making without feedback. However, both the anticipated and realized aspects of regret change substantially in repeated settings. Unlike one-shot decisions, repeated choices offer individuals the opportunity to learn about the incentive structure. With standard lotteries that provide no counterfactual information, people must enter the lottery to learn about the incentives, so learning motives push towards greater compliance with the incentivized activity. In contrast, individuals facing a sequence of regret lotteries can freely learn about the offered incentives without having to engage. This difference in the anticipated learning can lead to *less* engagement with regret lotteries than standard ones. Additionally, in contrast to one-shot choices, sequential decisions are subject to both anticipated and realized regret. A number of natural questions arise. Does the realized regret from entering the lottery and losing build up over time? Can this agglomeration be anticipated?

Given the differences in both anticipated and realized aspects of the incentive structure, it is critical to examine the extent to which behavioral effects like regret aversion extend from one-shot to repeated settings. In this paper, we directly study the differential effects of counterfactual information through a series of controlled laboratory experiments that compare participants' valuations for identical lotteries with and without counterfactual feedback. Our results do corroborate the higher valuations of regret lotteries in one-shot settings. However, when we examine the valuations for sequences of statistically independent decisions, the regret effect is not only mitigated, but it also *reverses* direction—people prefer standard lotteries over those which provide counterfactual information. Our results imply that regret lotteries may be inferior to standard lotteries for incentivizing repeated decision making.

Our paper adds a new facet to the emerging economics literature focusing on the potential for regret as an incentive tool for increasing 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 a superior incentive tool relative to standard lotteries (or to fixed payments) comes primarily from 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) 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 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). While our results do not dispute the idea that regret lotteries are a positive motivator, they do imply that regret lotteries are an inferior incentive device for recurring choices and suggest that managers and policy-makers have better available options. Additionally, our findings highlight general point of caution in extrapolating behavioral results from one-shot to repeated settings. Beyond just an attenuation of an effect size, our paper shows that the direction of the behavioral comparative static can be entirely reversed as we shift to repeated decision-making.

The paper's aim is to be a concise summary of the core empirical result, and is organized as follows: we begin with Section 2, presenting the design; this is followed by the results in section 3 and a summary of some robustness treatments; finally section 4 concludes.

2. Design

Our experimental design consists of a 2×2 between-subject experiment (summarize in Table 1) examining lottery valuations across:

- (**I**) The type of lottery:
  - (*a*) *Regret lotteries* where counterfactual information is available if the lottery is not entered.
  - (*b*) *Standard lotteries* without counterfactual information for those choosing not to enter.
- (**II**) Feedback on the lottery outcome:
  - (*a*) A *simultaneous* one-shot valuation decision which is used for thirty independent, identically distributed (*iid*) lotteries carried out after the decision.

Table 1. Core Design

Decision Type	Standard Lottery	Regret Lottery
One shot:	– Random ticket on entry – Single valuation – $N = 30$ participants	– Fixed ticket, printed at desk – Single valuation – $N = 30$ participants
Sequential:	– Random ticket on entry – Sequential valuations – $N = 30$ participants	– Fixed ticket, printed at desk – Sequential valuations – $N = 30$ participants

(b) Thirty *sequential* decisions for thirty *iid* lotteries, with feedback on the realizations after every choice.

Our lotteries in all treatments are implemented through an assigned ticket: three distinct numbers $\{A, B, C\}$ between 1 and 50. Lottery realizations are determined by a common physical draw of three numbered balls $\{a, b, c\}$ without replacement from a bingo cage at the front of the room containing fifty balls, numbered from 1 to 50. Prizes for the lottery are determined by the number of matches on the assigned ticket—the cardinality of $\{A, B, C\} \cap \{a, b, c\}$. Matching one ball number wins a prize of \$2.50, matching two yields a prize of \$25, while matching all three numbers yields a prize of \$250.[1]

Given 50 balls in the bingo cage there are 19,600 possible outcomes. Examining all possible draws the expected value of the lottery is calculated as:

$$\mathbb{E}V = \frac{3,243}{19,600} \cdot \$2.50 + \frac{141}{19,600} \cdot \$25.00 + \frac{1}{19,600} \cdot \$250.00.$$

Our decision task elicits participant's lottery valuations by asking for the maximum amount of money the participant would turn down to enter the lottery. Truthful reporting is incentivized through a Becker-DeGroot-Marschak (BDM) procedure with a uniform draw of an offer amount. If the offer amount is at or below the elicited value the participant enters the lottery; if the offered amount is greater the participant gives up the lottery and takes the offer.[2]

For each lottery-type we vary whether participants receive their lottery entry tickets before the decision to enter (printed out on their desks as they are seated) or only after choosing to enter (a randomized ticket assigned by the computer). The former therefore

[1] See the Online Appendix for a more-detailed discussion of the experimental details, as well as representative instructions.

[2] In a robustness section we also outline results from an alternative one-shot standard/regret treatment pair where we elicit valuations for a single lottery (rather than thirty *iid* draws).

implements a *regret lottery*, the latter a *standard lottery*. Participants in the regret-lottery condition know that they will observe the public lottery draw and realized outcome regardless of the decision to enter the lottery. In contrast, participants in the standard-lottery treatment only find out an assigned ticket conditional on entry—so those who accept the offer cannot know whether their ticket would have won or not. For the feedback dimension we vary whether participants make a single one-time valuation decision that binds for all thirty lotteries (our *Simultaneous* treatment) or whether they make repeated decisions over the 30 rounds with lottery realizations after each decision (our *Sequential* treatment).

The literature on regret aversion makes a clear prediction in the one-shot simultaneous treatment: The standard lottery has the same prize structure as the regret lottery but removes the possibility of *anticipated* regret from not entering. A regret-averse agent can take the offered amount knowing they will not be confronted with a counterfactual realization. In contrast, regret aversion provides a motive for participants to enter the lottery at greater rates. Participants can anticipate the regret they may feel if they do not enter but would have won a large prize (a maximum of $7,500 across the thirty decisions). Regret aversion therefore predicts that participants in the one-shot treatments will on average give up larger offers to enter the regret lottery than in the standard version.

2.1. **Predictions.** Why would the predicted effects of regret change for repeated decisions? The one-shot aspects of regret aversion that motivate the provision of counterfactual information focus entirely on the impact of *anticipated* regret, which is appropriate for a one-shot decision. However, in a repeated context, counterfactual information can also affect decision-making through differential opportunities for learning and realized outcomes.[3,4] With the large-prize–low-probability gambles that are typically implemented as part of regret lotteries, the large majority of participants will lose, and those that chose to comply with the incentivized activity will experience realized regret. Feelings of 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

[3]In contrast, the focus within the learning literature on regret ignores the anticipatory effects to focus on the backward-looking realized effects on a similar choice (see for instance Erev and Barron, 2005).

[4]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.

that same action in subsequent decisions (Keltner and Lerner, 2010).[5] While realized regret or rejoicing will not be an issue when the goal is to incentivize a one-time decision—such as opting into a 401-K plan or initial portfolio allocation decisions—it will be when incentivizing recurrent decisions such as exercise, recycling, or car-pooling to work.

Importantly, in a repeated setting even the anticipatory effects become a double-edged sword through the interaction with the learning opportunities provided by counterfactual information. Participants can adopt a wait-and-see approach if they are uncertain about the value of the offered lotteries, anticipating the opportunity to learn about the incentives without actually having to enter. Anticipation of learning with repeated regret lotteries can therefore serve to reduce value of entering the lottery even before the first outcome is realized.

Given the above, the predictions of regret theory are far from clear in the repeated setting once anticipated and realized factors interact. It is possible that anticipated regret outweighs the effect of ex-post realizations and anticipated learning differences, and the regret lottery will continue to be superior to the standard lottery in the sequential environment. However, if the realization and learning effects are dominant, the valuation of regret lotteries can be reduced below that of standard lotteries.[6] The main motivation of our experiment is therefore to examine how the regret comparative static is affected by the move from a one-shot setting to a repeated one, i.e. the interaction between lottery type and decision environment.

3. Results

Figure 1 illustrates the average valuations from each of our four experimental treatments. The first figure panel provides the overall averages for participant valuations V_t^i across each participant i and every round t. For ease of interpretation, it is presented relative to the expected value as a percentage.[7] The second figure panel shows how the valuations change across the rounds in our sequential sessions.

For statistical inference, we regress participant-level average valuations on two dummies, corresponding to lottery and decision environment type, and their interaction. We also include a gender dummy to assess and control for the significantly lower lottery valuations for women. The regression results are presented in Table 2. The econometric

[5] See Marchiori and Warglien (2008) and Hart (2005) for the use of regret as a process to predict subsequent choices.

[6] We should note here that our priors before conducting the experiments were that ex-post realizations and the high probability of not winning any prize in the lottery (82.7 percent) would substantially diminish the efficacy of regret as an incentive device. We did not expect that the downward pressure from anticipatory learning would be as significant as they were.

[7] That is for each treatment we provide the rescaled relative valuation $\hat{V}_t^i = \frac{V_t^i}{\mathbb{E}V}$.

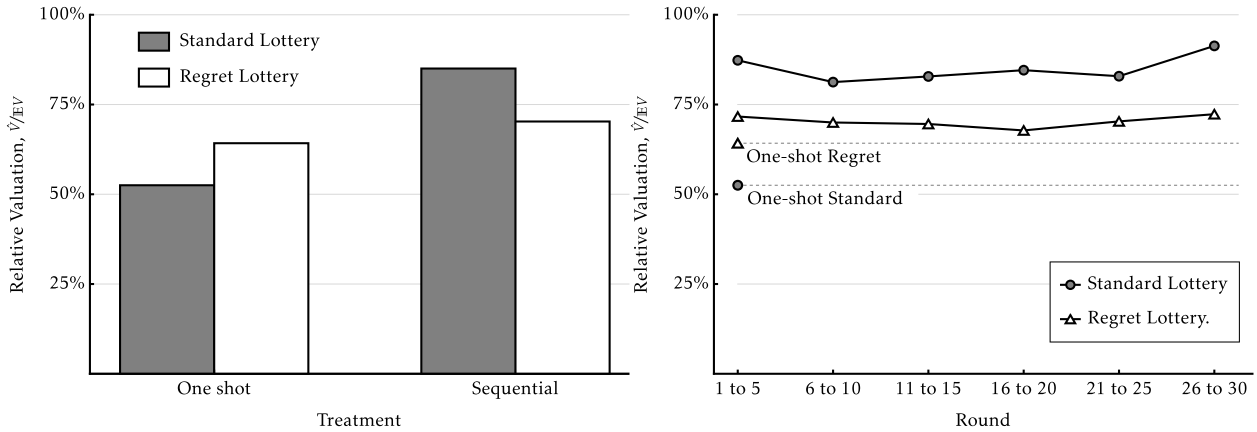

(A) Overall (B) Across time

FIGURE 1. Participant's lottery valuations, relative to actuarial value

equation is:

$$\hat{V}_i = \beta_0 + \beta_1 \cdot (\text{Seq.}) + \beta_2 \cdot \text{Std.Lot} + \beta_3 \cdot (\text{Seq.}) \times (\text{Std.Lot}) + \text{controls} + \epsilon.$$

3.1. **One-shot valuations.** The baseline behavior in the one-shot regret lottery treatment is captured by the constant coefficient, the relative effect of a standard lottery in the one-shot treatment by the *Std.Lot* term, while the interaction *Seq*× *Std.Lot* provides the estimated difference-in-difference across the lottery type and timing of decisions.

In the one-shot treatments, our experiments indicate that participants in the standard lottery value it at approximately half its expected value (52.5 percent). However, in the same one-shot choice environment participants increase their valuation of the lottery to 64 percent of its expected value when it is paired with anticipated regret. Replicating prior findings, the data indicates that when making a one-time decision, anticipated regret increases lottery valuations. The negative comparative effect from using standard versus regret lotteries in the one-shot setting is marginally significant with a two-sided test ($p = 0.084$).

The one-shot regret effect is estimated at 13.0 percentage points (of the lotteries actuarial value), reflecting an increase of 23.8 percent over the standard lottery for the average participant. The takeaway from the one-shot treatments can be summarized as follows:

Result 1 (One-Shot)**.** *Regret lotteries are effective in motivating entry in one-shot decision settings relative to standard lotteries.*[8]

[8]Another finding, though one our design did not explicitly set out to address, is that the lotteries are priced *lower* than the expected value. From a pure monetary incentive design point of view, the lottery is

Table 2. Regression results

Variable	Coeff.	Std. Err	p-Value
Constant	0.768	0.064	<0.001
Sequential	0.047	0.074	0.521
Std.Lottery	-0.130	0.074	0.084
Sequential×Std.Lottery	0.302	0.106	0.005
Female	-0.188	0.056	0.001
N	120		

3.2. **Sequential valuations.** In contrast to the finding for the one-shot environment, when looking at participants' average valuation across the 30 sequential decisions, the benefit of regret lotteries not only disappears, the direction of the effect is entirely reversed. While standard lotteries are valued at 85 percent of the expected value, average valuations for the regret lottery are 15 percentage points lower. Moreover, as Figure 1(B) illustrates by plotting the average valuation across the 30 rounds, the difference in valuations appears early on and persists across the sequential treatment.[9]

Moving to our main result which compares standard versus regret lotteries as a function of the decision environment, the regressions for participant-level average valuations in Table 2 indicate a large and substantial interaction effect. The difference-in-difference across our the two treatment pairs is 30.2 percentage points. Factoring in the negative estimated effect for the standard lottery in the one-shot treatment, the net effect in the sequential treatments is 17.3 percentage points (significantly different from zero with $p = 0.023$). As such, while regret boosts the standard lottery's incentive power in the one-shot setting by about a quarter, in the sequential environment it decreases the incentive power by a fifth.[10]

therefore inferior to paying the expected value with certainty. While this is true for men (total effect in a one-shot regret lottery of 76.7 percent the expected value, different from 100 percent with $p < 0.001$), it is exacerbated for women (total valuation of 58.0 percent for the regret lottery, different from the expected value with $p < 0.000$). However, we should note that another benefit to using low-win-probability lotteries for managers is the large reduction in administrative costs over incentive payment. Given this, it is unclear whether certain payments would dominate lottery incentives when the full scope of the costs are taken into account.

[9]Using regressions with treatment dummies we find that the difference between the standard and regret lotteries in the sequential treatment is significant in the very first round ($p = 0.001$), for the first half ($p = 0.010$ clustering for repeated observations by participant), for the final half ($p = 0.008$, again clustering), and for the last round ($p = 0.007$).

[10]The overall level effects for valuations indicate that participants are substantially more favorable towards the lottery in the sequential treatments. As our experimental design is constructed around differences we focus on the comparative static effects rather than the levels.

Though regret lotteries can be an effective incentive boost for one-time decisions, our data indicates they *reduce* the incentive efficacy in repeated settings. The takeaway from the sequential treatments can be summarized as follows:

Result 2 (Sequential Implementation)**.** *We find no benefit from regret lotteries as an incentive device in the repeated environment. In fact, the difference-in-difference is large enough to fully reverse the net effect of regret for repeated choice, where it is significantly worse than the standard lottery.*

3.3. **Robustness.** We examine the robustness of our results through an array of additional treatments, finding consistent results throughout. First, we ran a true one-shot treatment where we recruited participants to value a single lottery (30 participants for a standard lottery, 30 for a regret lottery) rather than eliciting fixed valuations for a sequence of independent lotteries.[11] Replicating the main study results for the one-shot setting, our data indicates that the standard lottery is valued 13.3 percentage points lower than the regret lottery, though the results are again only marginally significant at $p = 0.104$.[12] Pooling the two one-shot decision environments reveals that standard lotteries are valued significantly lower than regret lotteries at the five percent level ($p = 0.024$).

Second, one concern with the sequential environment is that participants' beliefs of winning may follow a law of small numbers. As such, a ticket that won last round may lead to a lower valuation the next round under the belief that it is unlikely to come up two times in a row. To examine this channel, we recruited participants ($N = 30$) to engage in a variant of the regret lottery where, rather than fixing the ticket number across all rounds (as it would be with an employee number or postcode), participants instead had to choose their own ticket before recording their valuation of the lottery. The results here indicate no significant difference ($p = 0.355$) with the regret lottery valuations in the main sequential treatment, and are lower than the standard lottery across all rounds.[13]

Finally, we also examined whether some social elements of regret that might be playing a role in field settings could increase effectiveness in sequential environments. A separate group of participants ($N = 30$) took part in a 'social' regret treatment. In this treatment, lottery tickets were fixed across groups of five and displayed prominently in the room

[11] Because participants only face one decision, we also multiplied the lottery stakes by 4 and enlarged the BDM elicitation.

[12] All subsequent p-values for tests from a regression mirroring Table 2. See Table A1 in the appendix for estimated coefficients.

[13] Because the regret-choice valuations lie strictly below the standard lottery valuations (see Figure A1 in the appendix), there is no possibility for a positive regret effect. However, in terms of the reversal, while first-round valuations for the regret lottery with ticket choice are significantly lower than standard lotteries ($p = 0.013$) matching the anticipated learning story, by the final rounds the observed decreased valuation of regret lotteries is not significant ($p = 0.435$).

so that each group could see the other group's ticket as well as their own. As such, the treatment mirrors some of the other-regarding factors in postcode lotteries, where a participant might feel increased regret if they do not win along with their neighbors when this is observable. Valuations in this treatment were similar to the main sequential regret treatment ($p = 0.923$), where the average valuation under 'social' regret is 0.7 percentage points *below* the main sequential regret-lottery treatment.[14]

4. Discussion & Conclusion

Our results replicate the finding that anticipated regret can increase the efficacy of a lottery incentive in a one-shot setting. However, when we move to a repeated setting the result is entirely reversed, with the prospect of counterfactual information driving down the lottery valuation.

In providing guidance to policy and incentive design, our results suggest that paying a constant non-stochastic payment may be preferable to a lottery incentive in many situations. However, where lotteries have other potential benefits, our results suggest two options. First, when repeated incentives have to be used, our results imply standard lotteries may be the superior tool. Second, rather than a repeated implementation for the same decision-maker, it may be more effective to instead design an incentive scheme closer to a one-shot setting. For example, a regret lottery paid one time at the end of a prolonged period—with required engagement levels across the entire period for entry—may be superior at nudging behavior towards a desired outcome until habits are formed.

Finally, our findings caution against the untested extension of behavioral decision-making phenomena from one-shot to repeated contexts. Learning, risk aversion over final wealth, and realized emotions can all contribute towards making a sequence of decisions non-separable, even when decisions are statistically independent. In some settings, this non-separability might exacerbate behavioral effects, and in others it may attenuate them. It could even reverse the direction of the effect. Our paper provides a clear example of this worst-case scenario from the perspective of incentive design or a policy intervention, such that the direction of the intended behavioral effect is entirely reversed as we move away from a one-shot to a repeated environment.

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.

[14]Figure A1 in the appendix illustrates nearly inseparable levels across the experiment between the social and standard regret lotteries. Sequential 'social' regret lotteries are valued persistently lower than sequential standard lotteries, where we can reject equivalence with $p = 0.004$ across all rounds (clustering by participant).

Cohen, J. D, S. M McClure, and A. J Yu, "Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration," *Philosophical Transactions of the Royal Society B: Biological Sciences*, may 2007, *362* (1481), 933–942.

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.

Gneezy, U. and J. Potters, "An Experiment on Risk Taking and Evaluation Periods," *Quarterly Journal of Economics*, may 1997, *112* (2), 631–645.

Gosnell, Greer K, John A List, and Robert D Metcalfe, "The impact of management practices on employee productivity: A field experiment with airline captains," *Journal of Political Economy*, 2020, *128* (4), 1195–1233.

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.

Iyer, Easwar S and Rajiv K Kashyap, "Consumer recycling: Role of incentives, information, and social class," *Journal of Consumer Behaviour: An International Research Review*, 2007, *6* (1), 32–47.

Jones, Damon, David Molitor, and Julian Reif, "What do workplace wellness programs do? Evidence from the Illinois workplace wellness study," *Quarterly Journal of Economics*, 2019, *134* (4), 1747–91.

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.

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.

Appendix A. Additional Results and Figures

A.1. **Additional Results.** Below we provide regression results including our additional robustness treatments.

Table A1. Regression results: Main + Robustness

Variable	**Coeff.**	**Std. Err**	p-**Value**
Baseline, regret one-shot decision (30 repetitions):			
Constant (Regret, One Shot)	0.772	0.058	0.000
Environment Effects:			
Sequential	0.046	0.077	0.555
True One-Shot	0.641	0.059	<0.001
Regret Effect in One Shot:			
Standard Lot.	-0.134	0.059	0.023
Lottery framing effect interactions with sequential:			
Sequential×Standard Lot.	0.307	0.102	0.003
Sequential×Choice Lot.	0.071	0.083	0.392
Sequential×Social Lot.	-0.007	0.0826	0.924
Gender:			
Female	-0.192	0.043	<0.001
N	240		

The additional treatments are:

True one-shot: An environment where we ask for a single valuation for a one-shot lottery, as opposed to a one-time valuation to be used across 30 lotteries. (30 participants in regret, 30 in standard)

Regret Choice: A sequential regret lottery where instead of a fixed ticket, we force particpants to choose a lottery ticket (three integers in 1–50) before entering their valuation. (30 participants, all non-standard)

Regret Social: A sequential regret lottery where groups of five have share a ticket that is publicly viewable to all ten participants in the session. (30 participants, all non-standard)

To estimate the effects we run a regression on the average valuations from 240 total participants where we estimate:

$$\hat{V}_i = \beta_0+\beta_1\text{SEQ}+\beta_2(\text{True1S})+\beta_3(\text{Std.Lott})+(\text{SEQ})\times(\beta_4\text{Std.Lott} + \beta_5 * \text{Choice} + \beta_5 * \text{Social})+\text{controls}.$$

The right-hand-side variables are dummies for:

SEQ: Sequential environment with feedback between choices.
STD.LOT: A standard lottery where the ticket is only generated on entry.
True-1S: Treatment indicator for *True one-shot* treatment pair.
Choice: Treatment indicator for *Regret Choice* treatment.
Social: Treatment indicator for *Regret Social* treatment.

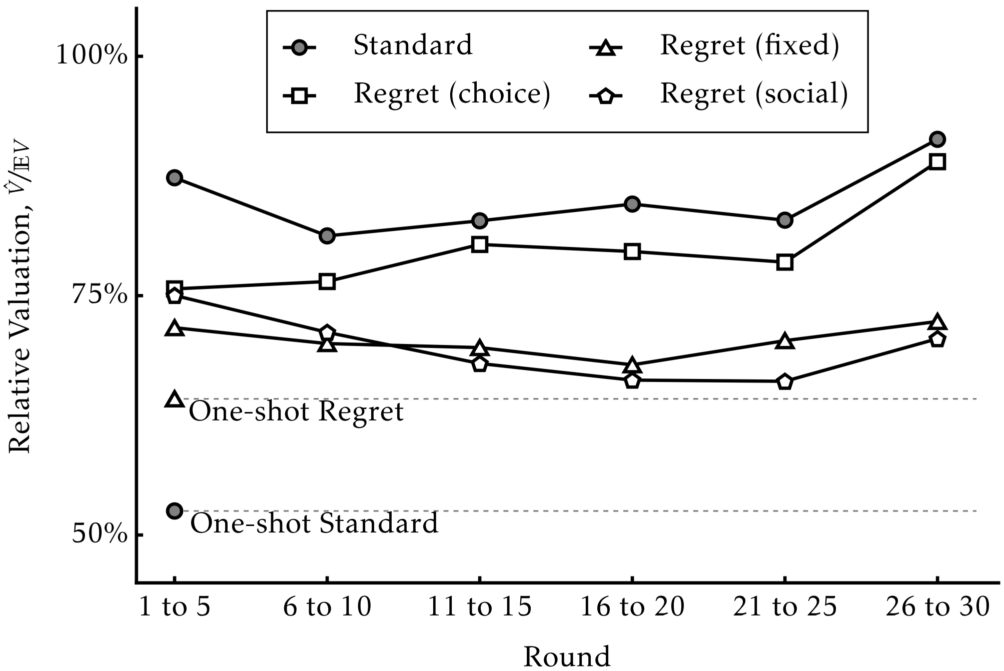

FIGURE A1. Sequential Robustness Treatments

A.2. **Experimental Details.** Our experiments were conducted at the Pittsburgh Experimental Economics Laboratory with participants recruited from the general undergraduate population of the University of Pittsburgh. Each experimental session comprised 10 participants, where no participant participated in more than one session.

In all treatments participants received a $5 fixed fee for their participation, plus their earnings from the decisions in the experiment. Common to all experiments, the first part of the instructions outlined a multiple price-list where participants make 20 decisions between a 50-50 gamble over $10 and nothing.[15] The fixed task in the price list was the lottery, and in each of the twenty-one questions the other option was a fixed certain amount increasing from $0.00 to $10.00, in $0.50 increments.

While the initial task elicits a risk and consistency measure, the main purpose of this first task was to familiarize participants with the valuation method for a lottery. After participants completed the price-list, we used the price-list to motivate the BDM procedure elicitation in the main experimental treatments.[16]

In all treatments, the instructions carefully outline the prize lottery and how realization of the outcome is determined. In explaining how an entry ticket translates into the three different prizes, we explicitly provide the odds of each winning event.[17]

For the main 30-round treatments, participants are told there will be a sequence of 30 rounds, where in every round prizes of $2.50, $25 and $250 are possible ($P = \$2.50$).[18] The maximum offer they would turn down to enter the lottery was elicited in every round with a BDM over the $0.00 to $1.00 interval. At the end of each round, the bingo cage was spun several times by the experimenter and three balls were drawn in turn. The numbers on the three balls were then publicly announced and entered into the monitor computer. Participants' screens informed them of their current earnings for the round.

[15] This was framed with the bingo cage at the front of the room, where we allowed them to choose 'Odd' or 'Even'. At the end of the session, after the main data collection, we chose one of the twenty-one price-list tasks for one of the ten participants (uniformly for both), determining the outcome by drawing a ball from the cage.

[16] We thank P.J. Healy for this suggestion in implementing instructions for the BDM.

[17] Representative instructions and screenshots are included as a supplement to this paper.

[18] We chose a different prize amounts for the one-shot treatment to make sure the lottery was well incentivized, as our focus is the standard/regret comparison in each feedback environment. In those treatments participants are informed there will be a single drawing with possible prizes of $10, $100 and $1,000 ($P = \$10$). Participants then indicated the valuations through a BDM over the $0.00 to $5.00 interval.

The lottery-type treatments were implemented through the manner in which tickets were issued. In the main regret treatments, participants' entry tickets were pre-assigned—printed on a piece of paper and placed on their desks as they arrived.[19,20] The assigned entry tickets were held constant across the entire session. Since the bingo cage drawing was public for each round, participants in the regret-lottery treatments always learned the outcome of the lottery—whether they would have won or lost—even if their decision meant they did not enter the lottery, i.e. the counterfactual.

In the standard-lottery treatments, participants were instead told they would only be assigned a ticket if they entered the lottery. After indicating their valuation, a random ticket was generated (a uniform draw across all possible tickets) and displayed on participants' screens during the lottery draw only if they had entered. Participants that did not enter therefore had no way to know whether they would have won or not.[21]

[19]Participants were randomly assigned to desks as they came in to the session through a draw from the bingo cage (ten balls, one for each desk). The pre-assigned entry tickets were uniform random draws from the set of all possible tickets.

[20]In one robustness treatment (Regret Choice) participants had to choose their own lottery ticket before the valuation, to check that having a fixed ticket was not causing the effect. In a second robustness treatment (Regret Social) the printed lottery tickets were shared across five participants, and were displayed publicly so that the other group of five could also see their neighbors numbers.

[21]While the sequence of uniform draws that determined the BDM's constant offers in each round were pre-drawn at the session level and matched, because the lottery outcomes were from a physical cage there was no way of matching these realizations across treatment.