

Imas, Alex; Lamé, Diego; Wilson, Alistair J.

Working Paper

Regret in One-Shot and Recurrent Decisions: A Cautionary Tale

CESifo Working Paper, No. 5939

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Imas, Alex; Lamé, Diego; Wilson, Alistair J. (2016) : Regret in One-Shot and Recurrent Decisions: A Cautionary Tale, CESifo Working Paper, No. 5939, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/144974>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Regret in One-Shot and Recurrent Decisions: A Cautionary Tale

Alex Imas
Diego Lamé
Alistair J. Wilson

CESIFO WORKING PAPER NO. 5939

CATEGORY 13: BEHAVIOURAL ECONOMICS

JUNE 2016

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: www.CESifo-group.org/wp

ISSN 2364-1428

Regret in One-Shot and Recurrent Decisions: A Cautionary Tale

Abstract

Under regret theory, decision-makers derive utility both from the outcome of their chosen action and the counterfactual. Evidence for anticipatory regret aversion has been found in oneshot settings, with "regret lotteries" that always reveal outcomes, as a counterfactual on non-entry, being priced higher than comparable standard lotteries that only realize outcomes for entrants. However, as anticipation and realization of regret necessarily interact in dynamic settings, the predictions of regret theory for repeated decisions are far from clear. Indeed while our one-shot experimental data corroborate the previous findings, data from a sequence of decisions show the reverse, with regret lotteries priced lower than standard lotteries and their certainty equivalents. Given the recent literature on the use of regret lotteries as incentives, our results suggest that while these lotteries can be effective for motivating one-time decisions, their benefits as a repeated incentive is less than clear. More generally, the paper illustrates the issues that can arise when extrapolating behavioral effects from one-shot to recurrent settings.

Keywords: regret aversion, separability, incentives.

Alex Imas
Carnegie Mellon University
Social and Decision Sciences
5000 Forbes Street
USA – Pittsburgh, PA
aimas@andrew.cmu.edu

Diego Lamé
University of Pittsburgh
Department of Economics
230 Bouquet Street
USA – Pittsburgh, PA
djw63@pitt.edu

Alistair J. Wilson
University of Pittsburgh
Department of Economics
230 Bouquet Street
USA – Pittsburgh, PA
alistair@pitt.edu

May, 2016

Our thanks to the following for their feedback: Linda Babcock, Julie Downs, Judd Kessler, PJ Healy, David Huffman, George Loewenstein, Graham Loomes, Lise Vesterlund and Stephanie Wang.

Regret theory (Bell, 1982; Loomes and Sugden, 1982) proposes a simple modification to expected-utility theory through which individuals derive utility not only from the outcome of their choices, but also from their knowledge of the counterfactual. Specifically, the decision-maker factors in both the utility of the chosen action under each possible realization, but also the utility of realizations under counterfactual choices. Regret aversion makes two assumptions: (a) people experience feelings of regret (rejoicing) when the outcome of their chosen action does worse (better) than the counterfactual outcome; and (b) these feelings are anticipated ex-ante. The consequence is that the decision maker may choose to avoid choices that have some possibility for ex-post regret.

A recent body of work has sought to utilize regret as a non-pecuniary boost to the power of deployed incentives through the use of *regret lotteries* (Zeelenberg and Pieters, 2007; Volpp et al., 2008b). Regret lotteries differ from a standard lottery in that a decision-maker knows she will learn the lottery realization regardless of her actual decision. In a typical regret lottery, all possible entrants are assigned a winning state (e.g. a randomly generated number, a zip code) and entered into the lottery either by making a purchase or embracing the incentivized activity (e.g. exercising). Regardless of the individual's entry decision, the lottery is resolved and the entire population is informed of the realized state. If the realized state matches an individual's winning state *and* she entered the lottery, then a corresponding prize is awarded. If either she did not enter the lottery or her particular winning state was not realized, then there is no prize. The lottery seeks to exploit regret aversion through the provision of counterfactual feedback regardless of the entry decision. Given regret-averse decision-makers, the knowledge that feedback will be provided yields a behavioral subsidy to the incentives. Entering the lottery becomes more desirable, with the anticipated regret from entering and not winning small in relation to the anticipated regret from not entering and missing out on a large cash prize.

A prominent example in the field is the Dutch postcode lottery, designed to incentivize the collection of public revenue. To realize the lottery's outcome, a winning postcode is drawn from the set of all Dutch postcodes. Each individual living within the particular postcode drawn that bought a ticket wins a large cash prize. Those in the drawn neighborhood that did not buy tickets do not get prizes but do observe that their postcode had won. Zeelenberg and Pieters (2004) argue that regret aversion should make this lottery program more successful than a national lottery in which the counterfactual (an assigned winning state) is not announced for those who chose not to purchase a ticket.

However, the argument for the increased incentive power of regret lotteries draws on the theoretical and empirical literature on anticipated regret, which focuses on one-shot decisions. Since anticipated and realized regret necessarily interact in dynamic settings, the predictions become less clear when considering repeated decisions. Are preferences over a series of lotteries different than

over a single one-time lottery? Does the regret from entering the lottery and losing build up over time? Can this agglomeration be anticipated?

In this paper we directly examine the effects of counterfactual information in one-shot and repeated settings through a series of controlled laboratory experiments. We compare subjects' valuations for statistically identical lotteries, with and without counterfactual feedback. While our results do corroborate regret-aversion in a one-shot setting, when we examine valuations in a repeated environment (a sequence of statistically independent decisions) the effect is reversed. In the repeated setting regret lotteries are an inferior incentive tool when compared to standard lotteries without counterfactual feedback.

Several recent review articles highlight the potential for regret lotteries as a policy tool (Madrian, 2014), improving health outcomes (Kessler and Zhang, 2014), the structure of incentives in firms (Babcock et al., 2012), and as a tool in development programs (Datta and Mullainathan, 2014). In much of the discussion regret lotteries are advocated as a tool for incentivizing recurrent, ongoing choice. However, the evidence that supports regret lotteries as a superior policy tool to more standard lotteries—or to fixed payments of the expected value—come from experiments with 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 any 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 to incentivize adherence to a prescription-drug regimen. Participants are randomized into either a regret-lottery treatment or a control with no incentive. Those in the treatment were assigned a "success number" for the duration of the study and informed of the outcome of a daily drawing. If the outcome of the draw matched the assigned number *and* the participant took a pill that day, she would win a prize; if the outcome matched the number but she had not taken a pill, no prize would be awarded. The authors report greater adherence to the drug regimen in the lottery treatment than in the no incentives control (see Volpp et al., 2008b; Kimmel et al., 2012; Haisley et al., 2012, for applications to other settings).

Our paper directly compares regret lotteries to the same lottery without counterfactual information, both in one-shot and repeated settings. While our results do not dispute regret lotteries as a positive inducement, they do suggest that policy-makers with a fixed incentive budget may have better options when incentivizing recurrent and ongoing behavior. As a more general point, our paper provides a prototype for a behavioral effect in a one-shot-setting that does not extrapolate to a series of repeated statistically independent decisions. Beyond demonstrating a non-separability across the decision sequence that attenuates the effect size, our paper instead shows that the *direction* of the effect's comparative static can be reversed.

Why would the effects of regret reverse in repeated decisions? While the aspect of regret theory that motivates counterfactual information has focused on the effects of anticipation, there has been substantially less focus on how the presence of counterfactual information affects outcomes through realizations.¹ While realized regret or rejoicing will not be an issue when the goal is to incentivize one-time decisions such as opt-ins or making an initial allocation decision for a portfolio, it will be when incentivizing repeated decisions such as exercise or refraining from smoking.² Some participants have to lose, and no single participant can consistently win the lottery again and again. Realized regret becomes a certainty, and the extent to which the effects of anticipated or realized regret may dominate becomes far less clear.

With the large-prize–low-probability lotteries that are typically implemented for regret lotteries the vast majority of participants will lose, and those that chose the incentivized activity will experience realized regret. Feelings of regret specifically have been shown to cause individuals to switch away from actions that produced the regret (Ku, 2008), while positive emotions such as rejoicing at having made the ex-post optimal choice tend to reinforce that same action in subsequent decisions (Keltner and Lerner, 2010).³ Of particular interest are those participants on the margins, where the counterfactual information produced a change in behavior. Over time, successive losses may build to induce a switch away from the incentivized activity. Unlike a standard lottery, lotteries with counterfactual feedback will reinforce the decisions of participants who chose not to enter the lottery and saw the counterfactual loss. Most participants choosing not to engage will receive frequent feedback reinforcing that their response was ex-post optimal.⁴

Importantly, in a repeated setting with counterfactual information even the anticipation of regret becomes a double-edge sword. Participants can adopt a wait-and-see approach if they are uncertain about the offered incentives, anticipating an opportunity to learn about the lottery without actually entering. Anticipation of learning without entering in the repeated setting could reduce engagement relative to a standard lottery even before the first outcome is realized, and counterfactual realizations will largely reinforce further non-engagement.

The paper’s organization is as follows. Section 2 outlines the experimental design. Section 3 presents the results from both one-shot and repeated settings. The paper’s organization is as follows. Section 4 concludes.

¹In contrast, the focus within the learning literature on regret is entirely backward looking through realizations (see for instance Erev and Barron, 2005)

²Where our paper examines decision-making in *repeated* but statistically independent decision making, Strack and Viefers (2015) examine a dynamic decision to divest from a risky asset with a persistent state. In this setting they find that counterfactual information helps, correcting for excessive risk aversion.

³See Marchiori and Warglien (2008) and Hart (2005) for the use of regret as a process to predict subsequent choices.

⁴In addition, other behavioral forces such as the gambler’s fallacy may serve to attenuate the effects from realized regret on *not entering* that push in the other direction.

2. EXPERIMENTAL DESIGN

Our experiment consists of a 2×2 between-subject design examining valuations across: a) the type of lottery, either a *regret lottery* or a *standard lottery* without counterfactual information for non-entry; and b) the feedback environment, either a one-shot or a repeated decision. Participants in each treatment round face a decision problem where they must choose between entering a lottery and taking a certain amount of money. The decision can be interpreted as being analogous to a choice between engaging with an incentivized activity or not, but with a foregone opportunity cost as represented by the constant amount.

The lottery incentive in all treatments is implemented through an assigned ticket: three distinct numbers $\{A, B, C\}$ between 1 and 50. Lottery realizations are determined through the drawing of three balls $\{a, b, c\}$ without replacement from a physical bingo cage, which initially contains 50 bingo-balls, numbered from 1 to 50. Prizes for the lottery are determined by the number of matched balls on the assigned ticket, the cardinality of $\{A, B, C\} \cap \{a, b, c\}$. Matching one ball yields a prize of P , matching two balls yields a prize of $10 \cdot P$, while matching all three balls yields a prize of $100 \cdot P$.

Given 50 balls in the bingo cage, there are 19,600 possible outcomes and the expected value (EV) of the lottery is given by

$$EV(P) = \left(\frac{3,243}{19,600} \cdot 1 + \frac{141}{19,600} \cdot 10 + \frac{1}{19,600} \cdot 100 \right) P.$$

Our decision task elicits participant's valuations for entering the lottery by asking for the maximum amount of money the participant would turn down to enter the lottery. Truthful reporting for this value 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 maximum, then the participant enters the lottery; if the offered amount is greater, then the participant takes the offer instead.

The basic design over the two treatment dimensions is summarized in Table 1.⁵ For the lottery-feedback dimension, we vary whether participants receive their lottery entry tickets before the decision to enter (printed out on their desks), or only after choosing to enter (randomized by the computer). The former represents a *regret lottery* and the latter a *standard lottery*. Participants in the regret-lottery treatment know that they will observe the lottery draw and the realized outcome regardless of whether they enter the lottery or not. In contrast, in the standard-lottery treatment participants find out their assigned numbers only upon entry—those who accept the offer instead of entering cannot find out whether their ticket would have won or not. For the feedback dimension, we vary whether participants make a single one-shot decision for a single lottery, or make repeated decisions over 30 rounds with lottery realizations after each decision.

⁵The appendix provides a more-detailed discussion of the precise experimental details and choices made.

	One-Shot	Repeated
Standard	$P = \$10.00$ Random ticket on entry 30 Subjects	$P = \$2.50$ Random new ticket on entry 30 Subjects
Regret	$P = \$10.00$ Fixed, printed ticket 30 Subjects	$P = \$2.50$ Fixed, printed ticket (all rounds) 30 Subjects

TABLE 1. Core Design

The literature on regret theory motivates our predictions in the one-shot setting. The standard lottery has the exact same prize structure as the regret lottery but removes the possibility of anticipated regret from not entering. A regret-averse agent can therefore take the constant offer knowing they will not be confronted with a counterfactual realization. In contrast, regret aversion provides a motive for individuals to enter the lottery in our regret treatment, with subjects anticipating the potential regret they may experience if they do not enter but win a large prize (up to \$1000 in the one-shot decision). Based on regret theory, we therefore predict that participants in the one-shot treatments are willing to give up larger offers to enter the regret lottery than the standard lottery. Separately, for a lottery to be an effective incentive (under a fixed prize budget), we also expect that the valuations will exceed the lottery’s actuarial value.

However, as discussed in the introduction, it is unclear whether the predictions of regret theory hold once anticipated and realized regret interact. If anticipated regret outweighs the effects of ex-post realizations, then the regret lottery will still be preferred to the standard lottery in our repeated treatments. If the ex-post effects dominate, then the regret lottery will be inferior to the standard lottery. Anticipated learning in the regret lotteries can also be examined in the repeated setting by restricting attention to first round outcomes.

Finally, in a repeated setting the standard risk-aversion concepts can become more muted. Given a large sequence of realizations, the law of large numbers will reduce the relative variation over final wealth for a participant who always enters. Repeated lotteries may therefore be less effective in exploiting risk-seeking behavior driven by factors such as probability weighting to induce entry. We will examine this idea by comparing subjects’ valuations to the lotteries’ expected values.

3. RESULTS

Figure 1 illustrates the overall results from our experiments. Our data consists of the elicited valuations V_t^i for each subject i in each round t . However, to facilitate direct comparisons between valuations provided for our one-shot and repeated treatments (given the different prize scaling

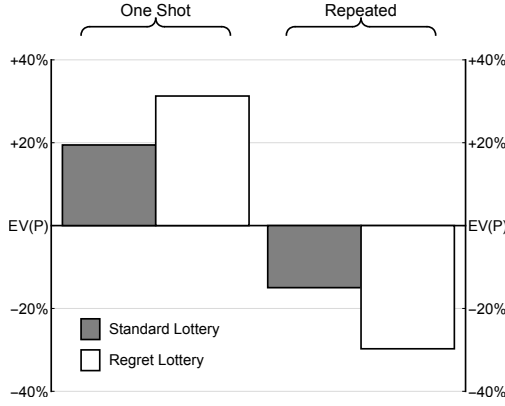


FIGURE 1. Average Valuations

	Coeff	Std. Err	p-Value
Regret	0.142	0.083	0.089
Repeated	-0.272	0.085	0.002
Repeated+Regret	-0.322	0.118	0.007
Female	-0.242	0.061	0.000
Constant	0.300	0.064	0.000
N	120		

TABLE 2. Regression results

P) we rescale and normalize, indicating subjects' percentage mark-ups on the lottery's expected value.⁶

In the one-shot treatments, we find subjects in the standard-lottery treatment (who only receive an entry ticket after giving up the offer) value the lottery at approximately a 20 percent premium over the lottery's expected value (19.5 percent, a \$0.47 premium). Subjects in the regret lottery treatment (who have an entry ticket before deciding) price the lottery at a premium of more than 30 percent (31.3 percent, a \$0.76 premium). The results therefore indicate that in the one-shot treatment the average participant has a higher valuation for the regret lottery than the standard lottery (see Table 2).

Result 1 (One-Shot Implementation). *Regret lotteries are more effective in motivating entry in one-shot settings relative to standard lotteries. Average valuations for both lotteries are greater than the lottery's expected value.*

In contrast to the one-shot environment, when looking at subjects' average valuations across the 30 rounds in the repeated treatment, the results are reversed. First, valuations are lower than the expected value for both types of lotteries. Moreover, the comparative-static relationship between the standard- and regret-lottery treatments is the opposite of that in the one-shot environment. Standard lotteries produce a 15 percent discount relative to the expected value, while the regret lotteries double this reduction to a 30 percent discount (\$0.09 and \$0.18 reductions per round, respectively).

Result 2 (Repeated Implementation). *Regret lotteries induce less entry than standard lotteries in the repeated environment. Average valuations for both lotteries are less than the lottery's expected value.*

⁶That is in each round of a treatment the expected value is $EV(P)$, and our data will focus on the rescaled valuation $\hat{V}_t^i = \frac{V_t^i - EV(P)}{EV(P)}$.

Looking at data at the subject level (averaging across the 30 rounds in the repeated treatment), we regress each subject’s relative premium/discount on treatment dummies for the regret lottery, feedback environment, and their interaction, as well as a gender control. The results of this regression are given in Table 2.

The regression results replicate the behavioral pattern in the aggregate results. While in the one-shot treatment the regret lottery is valued higher than the standard lottery, repetition drives the valuations below both the standard lottery and the expected value.⁷ The net effect of the regret lottery in the repeated treatment (regret treatment dummy plus the interaction) is an 18 percent reduction in valuation that is significant at the 95 percent confidence level ($p = 0.031$). In turn, our data suggest that regret lotteries *reduce* the effectiveness of the incentives in repeated settings, with lotteries in general seemingly less effective in motivating entry relative to offering the expected value with certainty.⁸

In the appendix we further examine the response across time in our repeated sessions, where we show that the preference for standard over regret lotteries emerges in the very first decision and persists across the entire session. In addition, at the subject level we show that the previous round’s realizations do affect valuations going forward. The task of identifying which particular factors are driving the differential effect of regret in one-shot versus repeated settings is beyond the scope of our note, and is left for future research. Note that because our paper focuses on demonstrating the reversal of a comparative static on behavior when going from a one-shot setting to a repeated one, our design does not allow us to identify which particular factors are driving the reversal. While this task is beyond the scope of our note, it is clearly an interesting question, and is left to future research.

4. DISCUSSION AND CONCLUSION

Our results indicate that the possibility of anticipated regret can increase the efficacy of a lottery incentive in a one-shot setting. However, in a repeated setting the result is reversed, with the prospect of counterfactual information driving down valuations of the lottery. The reversal holds both in relation to the standard lottery, as well as to the lottery’s expected value in each round.

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 settings where incentivized decision will be made repeatedly. If a lottery is to be used, rather than a repeated implementation for the same decision-maker, it may be more effective to instead design the incentive to be closer

⁷Results are qualitatively similar if we look at the subject’s first or last valuation instead of the average. In the appendix, we provide robustness checks to the assumptions involved in rescaling the prize magnitudes across repeated and one-shot settings.

⁸Table 2 includes a gender control to control for the effect that women are in general less willing to enter the lottery. While we add this control to reduce variability, the size of the effect points to greater possibilities from tailoring incentives: women in our sessions prefer the lottery’s EV in both the repeated and static treatments.

to a one-shot setting. For example, a regret lottery paid one time at the end of a prolonged period (with required engagement across the entire period) may be more effective than both types repeated lotteries (i.e. either with or without counterfactual feedback).

Lastly, our findings caution against the untested extensions of behavioral decision-making phenomena from one-shot settings 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 statistically independent, as in our setting. In some settings this non-separability might exacerbate behavioral effects, in others it may attenuate them. Our paper provides an example of a worst case from the point of view of policy, where the direction of an effect is entirely reversed with repetition. Our results illustrate the need for understanding how factors in repeated decision-making interact and affect behavioral phenomena identified in one-shot settings.

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.
- 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.
- 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.
- 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. Regret with Choice. In our implementation of the regret lottery, we fix subjects lottery entries. This would be similar to how an employer might fix the employees number in a lottery, or match lottery tickets with particular meaning (a child’s birthdate, an anniversary, etc.). But this particular implementation might mean a “gambler’s fallacy” could plausibly contribute to Result 2. For example, suppose that a subject with the fixed ticket $< 10, 18, 24 >$ saw that last round the numbers chosen from the cage were $< 10, 16, 22 >$. Because the 10 ball was recently chosen, they may mistakenly believe it will be less likely to be drawn in the next round. Beyond just matches to their number, they may think that because all three numbers drawn last time were all lower than 25, or were all even, that their particular ticket (with these same characteristics) will be less likely to win next round.

To examine this idea, we conducted an additional treatment with 30 more subjects, which we call our *Regret Choice* treatment. Before entering their valuation for the lottery, we ask them to choose their ticket $< X, Y, Z >$. Through the act of choosing their numbers before realization of entry subjects are still able to anticipate regret of not entering if their numbers are drawn from the cage. But, to counteract the gambler’s fallacy, their choice can now adjust to choose whichever three numbers they think most likely to come up. Everything else was held constant from the Regret treatment.

Our results for the repeated *Regret Choice* treatment are not substantially different from repeated *Regret*. Relative to the actuarial value subjects on average discount their values by 20 percent. Though this is less discounted than the 30 percent observed in the *Regret* treatment, it is more-discounted than the 15 percent discount in *No Regret*.

Though we cannot reject the hypothesis that *No Regret* and *Regret Choice* produce the same effect with a two-sided alternative, we can reject the one-sided alternative hypothesis that motivated the research: neither form of regret lottery is superior to the standard lottery.⁹

A.2. Removing Information. In all of the experiments detailed above we provide subjects with both a description of the lottery procedure, and the explicit probabilities of each winning outcome. While information on probabilities is certainly calculable from a description of the procedure, and may be provided in some form to participants, in many settings this information will be hard to access. For example, participants in state-run lotteries will typically understand the procedure for draws, but formal statements of the probabilities are not prominent at the point of sale, and are available if sought out or asked for.

⁹Using a regression similar to Table 2, but focused only on the repeated treatment averages, we find a significant effect for *Regret* ($p = 0.007$) and *Regret Choice* ($p = 0.085$) against the one-sided alternative that the relative effect of regret is positive.



FIGURE A1. Average valuations relative to midpoint

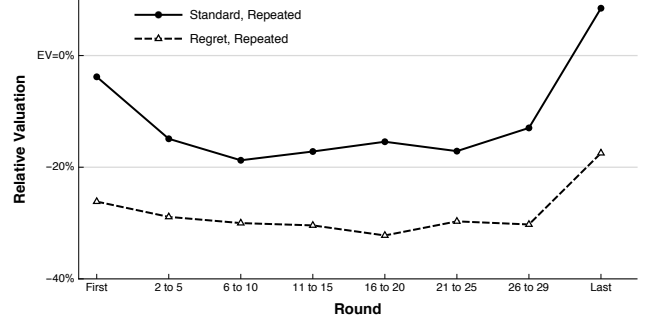


FIGURE A2. Average valuations across time

To examine the effect of providing probability information to subjects we re-ran our repeated *No Reret* and *Regret* treatments. Everything in these replications was identical to before, except that we now remove provided information on the explicit probabilities of each prize amount.

Results from the 60 subjects across these two treatments indicate no significant differences between the valuations for regret and no-regret lotteries, though the relative effect of the regret lotteries is still negative (which are discounted by 4 percent more than the standard lotteries).

Why then are the results for regret lotteries neither as negative, nor as sharp, as those with information on probabilities? The reason seems to be a much greater variation of values across subjects. The between-subject standard deviation for values in our treatments is 21.1 percent (standard error of 1.8 percent) of the actuarial value (\$0.13 per round). However, in our treatments without information this increases by approximately two-thirds to 35 percent (2.9 percent) of the actuarial value (\$0.21 per round).¹⁰

A.3. Valuation over time. Where the main paper summarizes the data across all four treatments, purposefully keeping the presentation condensed, we here examine the valuations across time in the repeated treatments. Figure A2 presents the relative valuations over time for the two types of lotteries.

The first striking result is that the valuation of the regret lottery is significantly lower than the standard lottery across almost every round, including the very first. In the first round, the valuation of the standard lottery is at a slight 4 percent discount relative to the expected value of \$0.61. In contrast, the regret lottery is valued significantly lower at \$0.45—a 26 percent discount relative to expected value ($p = 0.001$). Repeating this procedure for each round reveals that the regret lottery is directionally valued less than the standard lottery for all rounds, significantly less at a 90 (95) percent confidence level in 24 (17) of the 30 rounds.

¹⁰Standard errors on the between-subject standard deviation are obtained by 1,000 bootstraps of a random-effects regression

Covariate	Standard			Regret		
	Coeff.	Std.Err	p-Value	Coeff.	Std.Err	p-Value
Enter & Won	0.132	0.026	0.000	0.060	0.026	0.023
Accept Offer, \emptyset	0.179	0.032	0.000	–	–	–
Accept Offer, CF Win	–	–	–	0.133	0.034	0.000
Accept Offer, CF Loss	–	–	–	0.106	0.031	0.001
Trend	-0.105	0.017	0.000	-0.071	0.022	0.002
Last Round ($t = 30$)	0.156	0.050	0.003	0.118	0.066	0.080
N	870			870		

Note: Dependent variable is $\hat{V}_t^i - \hat{V}_{t-1}^i$ for subject i in period $t \in \{2, 3, \dots, 30\}$. Reported standard errors are clustered at the subject level, omitted category for constants in the regression is Entering and Losing. Variables in the control/treatment columns are interactions with the relevant treatment.

TABLE A3. Within-subject valuation changes

We now turn to examining how valuations evolve over time in more detail. Looking at the variation across rounds at the subject-level, it becomes apparent that there is substantial learning based upon previous outcomes. To demonstrate this, we examine the change in each subject's valuation $V_t^i - V_{t-1}^i$ between rounds and how this change responds to the observed realizations in the previous round.

Table A3 reports results from a single regression for how the observable outcome in the previous round, a trend variable and a last-round dummy affect the change in valuation in each treatment. For the standard lottery, there are three possible prior-round outcomes: (a) enter and win; (b) enter and lose; and (c) accept the offer.¹¹ In contrast, the accept offer event in the regret lotteries is further divided into accept-offer and observe a counterfactual win, and accept offer and observe a counterfactual loss, resulting in four different outcomes in total.

Examining the regression coefficients and their significance, it is clear that past realizations do affect changes in valuation. Entering the lottery and losing is the omitted category for each lottery type and the effect of this outcome is captured by the Trend term. The coefficient is negative and significant for both lotteries, implying that entering and losing has a negative effect on valuation in the two treatments. In contrast, the coefficient for accepting the offer with no counterfactual observation is positive and significant, reflecting a relative increases in valuation.

While the choice to accept the offer without counterfactual information leads to a net valuation increase of 7.4 percent of the lottery's EV, the net response for regret lotteries after accepting the offer and observing a counterfactual loss is only 3.5 percent. Since counterfactual losses are by the

¹¹We will call a win matching one or more numbers, where the vast majority of these are one-number matches. No subject matched all three numbers on an entered ticket across all of our experiments, while only one subject in a single round matched three numbers on a counterfactual.

far the most common outcomes, the overall effect is that a decision to disengage from the regret lottery is significantly more persistent in the regret treatment. As such, the gap in valuation between the regret and standard lotteries formed in the first round—which may be driven by subjects anticipating that they will learn about the lottery through the provided counterfactual information even if they do not enter—is not closed.¹²

A.4. Experimental Details. Our experiments were conducted at the Pittsburgh Experimental Economics Laboratory between March and December 2015, with subjects recruited from the general undergraduate population of the University of Pittsburgh. We ran a total of 21 sessions with 10 participants in each session, where no subject participated in more than one session.

In all treatments subjects 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 subjects make 20 decisions between a 50-50 gamble over \$10 and nothing.¹³ 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 does provide an elicited measure of risk-aversion and consistency (whether the price-list response was monotone), the main purpose of this first task was to familiarize subjects 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.¹⁴

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.¹⁵ For our one-shot treatments, subjects are told that there will be a single drawing with possible prizes of \$10, \$100 and \$1,000 ($P = \10). Subjects then indicated the maximum certain value they would turn down to enter the lottery, incentivized through a BDM over the \$0.00 to \$5.00 interval. For the repeated-implementation treatments, subjects are told there will be a sequence of 30 rounds, where in every round prizes of \$2.50, \$25 and \$250 are possible ($P = \2.50).¹⁶ 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

¹²In the Appendix we present additional tests of the robustness of our results such as manipulating the level of lottery information provided and giving subjects the choice to pick their regret lottery numbers before the offer decision.

¹³This 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 subjects (uniformly for both), determining the outcome by drawing a ball from the cage.

¹⁴We thank P.J. Healy for this suggestion in implementing instructions for the BDM.

¹⁵Representative instructions and screenshots are included as a supplement to this paper.

¹⁶We 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.

drawn in turn. The numbers on the three balls were then publicly announced and entered into the monitor computer. Subjects screens informed them of their current earnings for the round.

The lottery-type treatments were implemented through the manner in which tickets were issued. In the regret treatments, subjects' entry tickets were pre-assigned—printed on a piece of paper and placed on their desks as they arrived.¹⁷ The assigned entry tickets were held constant across the entire session. Since the bingo cage drawing was public for each round, subjects 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, subjects 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 subjects' screens during the lottery draw only if they had entered. Those subjects that did not enter had no way to know whether they would have won or not.¹⁸

¹⁷Subjects were randomly assigned to desk numbers as they came in 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.

¹⁸The sequence of uniform draws that determined the BDM's constant offers in each round were pre-drawn at the session level, and so across the standard and regret treatments we can match subjects with the same series of constant offers.