EFFECT REVERSALS OVER ONE-SHOT VERSUS RECURRENT DECISIONS: THE CASE OF REGRET

ALEX IMAS, DIEGO LAMÉ, AND ALISTAIR J. WILSON

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 one-shot settings, with "regret lotteries" that always reveal outcomes, both upon entry and as a counterfactual upon non-entry, being priced higher than comparable standard lotteries that only reveal outcomes for entrants. These findings have motivated a literature advocating for the use of regret lotteries as a policy tool to boost the power of incentives for desirable behaviors such as excercise and drug adherence, which often require repeated decisions. Importantly, however, differences in learning opportunities and the interaction between anticipated and realized regret make the predictions of regret theory in repeated settings far from clear. Through a series of experiments, we show that indeed in one-shot settings regret lotteries are valued higher than standard lotteries. In contrast, for sequential decisions this pattern reverses: in repeated settings regret lotteries are valued significantly lower than standard lotteries and their certainty equivalents—even in the very first period. In light of 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 repeated settings.

Regret theory (Bell, 1982; Loomes and Sugden, 1982) proposes a simple modification to expected-utility theory through which individuals derive utility not only from choice outcomes, but also from counterfactual knowledge of outcomes from alternative choices. Specifically, the decision-maker factors in her chosen action under each possible realization, just as in expected utility theory (EUT), and also derives utility from realizations under counterfactual choices. Regret theory 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 of the theory for decision making is that, relative to optimal choices under EUT, individuals may select options that seek to minimize 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

---

Date: February, 2017.
Imas: Carnegie Mellon University, Social and Decision Sciences, 5000 Forbes Street, Pittsburgh, PA; aimas@andrew.cmu.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, Lise Vesterlund and Stephanie Wang.

al., 2008b). Regret lotteries differ from standard lottery incentives by ensuring that the decision-maker is aware that she will learn the lottery's realization regardless of her actual decision. In a typical regret lottery, the entire population of possible entrants are first assigned a winning state (a randomly generated number, a zip code). They then make choices of whether to enter the lottery or not by making a purchase or embracing the incentivized activity (e.g., exercising, taking prescribed drug). Regardless of entry decisions, the lottery is resolved and the entire population is informed of the realized winning state. If the realized state matches the individual's *and* she chose to entering the lottery, then a corresponding prize is awarded. If either she did not enter the lottery 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. Given a regret-averse decision-maker, the knowledge that feedback will be provided yields a behavioral subsidy to the incentive, where anticipated regret from not entering and missing out on the prize can push those on the margins to comply with the incentivized activity.

A prominent example from the field is the Dutch postcode lottery, designed to incentivize the collection of public revenues. To realize the lottery's outcome, a winning postcode is drawn from the population 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 (and neighbors) 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 (an assigned winning state) is not provided to individuals who do not purchase 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. However, both the anticipated and realized aspects of regret change in repeated settings; in turn, the predictions on its efficacy as a behavioral boost to incentive become less clear when considering sequential decisions. Unlike with standard lotteries that provide no counterfactual information, individuals facing a sequence of regret lotteries can learn about the offered incentives without having to enter the lottery. This difference in anticipated learning can lead to *less* people entering regret lotteries than standard ones. Additionally, in contrast to one-shot settings, sequential decisions are subject to both anticipated and realized regret. Does the regret from entering the lottery and losing build up over time? Can this agglomeration be anticipated?

Due to these potential differences, it is critical to examine the impact of regret on decision-making both in one-shot and repeated settings. In this paper, we directly study the effects of counterfactual information in both settings through a series of controlled laboratory experiments which compare subjects' valuations for identical lotteries with and without counterfactual feedback. While our results do corroborate a higher valuation of regret lotteries in one-shot settings, when we examine valuations in a sequences of statistically independent decisions the effect is *reversed*. These results suggest that regret lotteries are an inferior incentive tool when compared to standard lotteries for repeated, sequential decision making.[1]

Our paper adds a new facet to an economic literature focusing on the potential of regret lotteries as an effective incentive for better policy design(Madrian, 2014), improving health outcomes (Kessler and Zhang, 2014), increasing motivation in firms (Babcock et al., 2012), 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 that does support regret lotteries as a superior policy tool relative to standard lotteries—or to fixed payments of the expected value—come primarily from experiments examining one-shot decisions (Loomes and Sugden, 1987; see Zeelenberg, 1999 for a review).

While some studies 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 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 that regret lotteries offer a positive inducement for both one-shot and sequential decisions, they suggest that regret lotteries may be an inferior incentive device in repeated settings. Policy-makers with a fixed incentive budget may have better available options incentivizing sequential decisions.

Additionally, our findings speak to a more general point of caution in extrapolating results from a one-shot behavioral effect to repeated settings. Beyond just an attenuation of the effect size, our paper instead shows that the behavioral comparative-static *direction* can be reversed when moving from single choices to sequential decisions.

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 anticipated regret, there has been substantially less focus on how the presence of counterfactual information affects outcomes through realizations and differing opportunities for learning.[2] While realized regret or rejoicing will not be an issue when the goal is to incentivize one-time decisions (such as opt-ins or initial allocation decisions for a portfolio) it will be when incentivizing recurrent decisions (such as exercise or refraining from smoking). Some participants have to lose, and no single participant can

[1]Where 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.

[2]In contrast, the focus within the learning literature on regret does not consider anticipated effects (see for instance Erev and Barron, 2005)

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 gambles 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, the availability of counterfactual feedback in a regret lottery can reinforce the decisions of participants who chose not to enter the lottery initially. Most participants who chose not to enter will receive frequent feedback reinforcing their initial decision.[4]

Importantly, in a repeated setting even anticipatory effects become a double-edge sword due to the additional opportunities for learning from the counterfactual information. 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 having to enter. Anticipation of learning in a repeated setting with counterfactual information could therefore reduce engagement relative to a standard lottery even before the first outcome is realized. Indeed, we find evidence that the anticipated learning channel does lead to less engagement with regret lotteries in our repeated environment. Subjects are less willing to enter the regret lotteries than standard lotteries even at the very first decision before any realizations have occurred. This suggests that in contrast to one-shot settings, anticipatory factors can actually *decrease* the appeal of regret lotteries.

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.

2. Experimental design

Our experiment consists of a 2×2 between-subject design examining 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) The repetition and feedback of the decision:

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

[4] In addition, other behavioral dynamic forces may serve to attenuate the effects which push in the other direction. For example, a gambler's fallacy may mitigate the discouraging effects of realized regret on the choice not to enter.

(A) A *simultaneous* one-time binding decision over thirty independent, identically distributed (*iid*) lotteries carried out after the decision.

(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 publically drawing three numbered balls $\{a, b, c\}$ without replacement from a physical bingo cage at the from the room, which contains 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 $\$250$.

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

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

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.[5]

The design over the two treatment dimensions is summarized in Table 1.[6] For the lottery-feedback dimension (**I**) 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 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 (**II**) 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).

[5]We also conduct an alternative simultaneous standard/regret treatment pair where we elicit values over a single lottery (rather than thirty *iid* draws), where the prize amounts are multiplied by 4 to increase the stakes. Results from these treatments are qualitatively the same as in our main treatments, and are reported in the Appendix.

[6]The Appendix provides a more-detailed discussion of the precise experimental details, as well as providing representative instructions.

	Standard	**Regret**
Simultaneous	• Random ticket on entry • Single Decision • 30 Subjects	• Fixed, printed ticket • Single Decision • 30 Subjects
Sequential	• Random ticket on entry • Thirty sequential decisions • 30 Subjects	• Fixed, printed ticket • Thirty sequential decisions • 30 Subjects

TABLE 1. Core Design

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 subjects to enter the lottery at greater rates. Subjects 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 simultaneous treatments will on average give up larger offers to enter the regret lottery than the standard version.

However, as discussed in the introduction the predictions of regret theory are far from clear once anticipation and realization interact. It is possible that anticipated regret outweighs the effect of ex-post realizations, and the regret lottery will continue to have greater value than the standard lottery in the sequential decisions. Conceptually, though intuitively unlikely, it is possible that anticipation and realization ramp up the effects of regret, increasing the level effect of regret. However, if realization effects are dominant—or indeed the anticipation of learning the realizations without entering—the higher valuation of regret lotteries can be reduced. Moreover, if these effects are large enough, the valuations of regret lotteries can actually be lower than those of standard lotteries.[7] The main motivation of our experiment is to understand and document the comparative statics associated with moving from a one-shot decision to a repeated setting.

3. RESULTS

Figure 1 illustrates the aggregate results from each treatment, where the figure provides the averages for subject's valuations V_t^i (for each subject i in each round t), which we express relative to the expected value as a percentage.[8] For statistical inference to complement the figure we use Table 2, which regresses subject-level averages on two treatment-dummies reflecting the regret

[7]We should note here that our priors going into this project were that ex-post realizations and the high probability of not winning any prize in the lottery (82.7 percent) would retard the efficacy of regret as an incentive over time. We did not *anticipate* that the downward pressure of anticipatory learning would be the dominant force here.

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

FIGURE 1. Average Valuations Relative to Actuarial Value

effects, as well as three level controls. The two main regret effects are captured by the *Regret* coefficient, while the interaction *Seq*$\times$ *Regret* provides the estimated difference-in-difference for the results.

In the simultaneous treatments, our experiments indicate that subjects in the standard lottery value it at approximately half its expected value (52.5 percent). However, in the same simultaneous choice environment subjects price up the regret lottery to 64 percent of its expected value. Matching the prior regret literature, the data indicates that when making a one-time decision participants' anticipated of regret drives up their lottery valuations. The positive effect from regret in the one-shot setting is marginally significant with a two-sided test.[9] The one-shot regret effect is estimated at 13.0 percentage points (of the expected value), reflecting a pricing increase of 23.8 percent over the standard lottery for the average subject. The conclusion for the one-shot treatments is summarized as follows:

Result 1 (One-Shot)**.** *Regret lotteries are effective in motivating entry in one-shot decision settings relative to standard lotteries.*[10]

In contrast to the finding for the one-shot environment, when looking at subjects' average valuations across the 30 sequential decisions, regret effects are entirely *reversed*. While standard

[9]We additionally ran one-shot sessions where subjects provided valuations for a single lottery (rather than a sequence of thirty), where we multiplied all prize amounts by 4 ($10/$100/$1,000). The regression indicates a marginally significant regret effect of 14.2 percentage points—measured relative to the quadrupled expected value of $2.43. The relative regret effects are therefore remarkably similar for one-shot decisions, both for a simultaneous decision over 30 realizations and for a single realization. Pooling the two one-shot decision environments, the positive effect for regret is significant ($p = 0.024$), where we report these results in Tables A1 and A2 in the appendix.

[10]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 the point of view of its use as an incentive device with a fixed budget, the lottery is inferior to paying the expected value with certainty. 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 $p = 0.000$), but particularly so for women (total effect of 58.0 percent combined with regret, $p = 0.000$).

Var.	Coeff.	Std. Err	p-Value
Regret Effects:			
Regret	0.130	0.075	0.084
Seq×Regret	-0.322	0.106	0.005
Level Controls:			
Sequential	0.350	0.085	0.002
Female	-0.188	0.056	0.001
Constant	0.638	0.062	0.000
N	120		

TABLE 2. Regression results

lotteries are valued at 85 percent of the expected value, average valuations for the regret lottery is 15 percentage points lower.[11]

The regressions for subject-level average valuations in Table 2 indicates a large and substantial difference-in-difference across our the two treatment pairs, -32.2 percentage points. Factoring in the positive estimated effect for regret in the simultaneous treatment, the net regret effect in the sequential treatments is -17.3 percent (significantly different from zero with $p = 0.023$). Where regret boost the standard lottery's incentive power in the simultaneous setting by about a quarter, in the sequential environment it decreases the incentive power by a fifth.[12] The conclusion for risk in sequential decision making is therefore:

Result 2 (Sequential Implementation)**.** *Regret lotteries induce less entry than standard lotteries in the repeated environment.*

Though regret lotteries can be an effective incentive boost for one-time decisions, our data indicates they *reduce* the incentive efficacy in repeated settings.[13]

In Imas et al. (in preparation) we further examine within-subject learning across time in our sequential valuation environment, and similar environments where we vary the dynamic feedback. Though we show that in that analysis that the previous round's *realizations* does have strong effects on subject's round-to-round valuations, from the aggregate point of view which we take here for

[11]In the appendix we provide details on the dynamics across the sequence of decisions, where we repeat the average-value regressions in Table 2 for the values from the first and last rounds (Tables A3 and A4, respectively). The main finding there is that the net regret effect for sequential decisions is significantly negative from the very first valuation, and remains so over the course of the entire experiment.

[12]The overall level effects for valuations indicate that subjects are substantial more favorable towards the lottery in the sequential treatments. As our experimental design is constructed around differences we focus on these effects rarther than the levels. However, we do plan to examine the level effects in future work. In particular, the risk aversion demonstrated in the simultaneous treatment valuations is strange, where repetition for thirty rounds should *decrease* the effective risks.

[13]Moreover, in both cases lotteries are inefficient relative to offering the expected value of the lottery with certainty.

assessing regret lotteries there is little effect on our overall conclusions. The negative regret difference in sequential emerges from the very first decision, and persists across the decision sequence.

The more-involved task of identifying and pulling apart all of the factors driving differential regret effects in simultaneous and sequential settings is beyond this short paper's scope, and is left for future research. Instead, our paper focuses on demonstrating the possibility for reversal of behavioral comparative-statics as we move from one-time to sequential decision-making. We make this point not as a claim that all behavioral incentive available to us might fail in generalizing across these settings, but in making clear that some caution is warranted.

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 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.

Imas, Alex O., Diego Lamé, and Alistair J. Wilson, "Learning about a Lotteries Value: Realized Regrets," in preparation.

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

Var.	Coeff.	Std. Err	p-Value
	Regret Effects		
Regret	0.142	0.083	0.089
Seq×Regret	-0.322	0.118	0.007
	Level Controls		
Sequential	-.272	0.085	0.002
Female	-.242	0.061	0.000
Constant	0.660	0.060	0.000
N	120		

Table A1. Regression Results (Averages) replacing Simultaneous with One-Shot

Var.	Coeff.	Std. Err	p-Value
Regret	0.137	0.060	0.024
Seq×Regret	-0.314	0.104	0.003
	Level Controls		
Sequential	.358	0.080	0.000
One-Shot	0.636	0.060	0.000
Female	-0.228	0.051	0.000
Constant	0.660	0.060	0.000
N	180		

Table A2. Regression Results Pooling Simultaneous and One-Shot

Var.	Coeff.	Std. Err	p-Value
	Regret Effects		
Regret	0.129	0.080	0.110
Seq×Regret	-0.376	0.114	0.001
	Level Controls		
Sequential	0.461	0.080	0.000
Female	-0.178	0.060	0.004
Constant	0.632	0.069	0.000
Subjects	120		

Table A3. Regression Results (First Round)

A.1. **Single Lottery Treatment.**

Var.	**Coeff.**	**Std. Err**	**p-Value**
	Regret Effects		
Regret	0.135	0.100	0.178
Seq×Regret	-0.431	0.141	0.003
	Level Controls		
Sequential	0.596	0.100	0.000
Female	-0.270	0.074	0.004
Constant	0.687	0.084	0.000
Subjects	120		

TABLE A4. Regression Results (Last Round)

A.2. **Regret with Choice.** In our implementation of the regret lottery, we fix subjects' lottery tickets. 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.[14]

A.3. **Valuations over time.** Where the main paper summarizes the data across all four treatments and purposefully keeps the presentation condensed, we here examine the valuations across time in the repeated treatments. Figure **??** 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.

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 A5 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.[15] 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.

[14]Using 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.

[15]We 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.

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 $t \in \{2, 3, \ldots, 30\}$. 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 A5. Within-subject valuation changes

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 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.[16]

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.[17] 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.

[16]In 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.
[17]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 subjects (uniformly for both), determining the outcome by drawing a ball from the cage.

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.[18]

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.[19] 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 offer they would turn down to enter the lottery, incentivized through a BDM over the \$0.00 to \$5.00 interval.

For the main 30-round 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 ($P = \$2.50$).[20] 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. 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.[21] 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.[22]

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

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

[20]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.

[21]Subjects 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.

[22]The 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.