The Realization Effect: Risk-Taking After Realized Versus Paper Losses

By Alex Imas *

Understanding how prior outcomes affect risk attitudes is critical for the study of choice under uncertainty. A large literature documents the significant influence of prior losses on risk attitudes. The findings appear contradictory: some studies find greater risk-taking after a loss, whereas others show the opposite – that people take on less risk. I reconcile these seemingly inconsistent findings by distinguishing between realized versus paper losses. Using new and existing data, I replicate prior findings and demonstrate that following a realized loss, individuals avoid risk; if the same loss is not realized, a paper loss, individuals take on greater risk.

"A person who has not made peace with his losses is likely to accept gambles that would be unacceptable to him otherwise."

– Kahneman and Tversky (1979)

"Losses that come on the heels of prior losses may be more painful than average...after a prior loss, the person becomes more loss averse."

– Barberis, Huang and Santos (2001)

Understanding how prior outcomes affect risk attitudes is critical for the study of choice under uncertainty in many economically important contexts. If the value of a stock falls below the purchase price, does this affect the investo's subsequent behavior, and if so, does he seek a riskier position or switch to a safer one? If a casino gambler loses money at the roulette table, does he get discouraged and quit gambling or chase his losses?

Standard expected utility theory assumes that prior outcomes influence risk-taking only if there is a substantial change in wealth (Savage, 1954); an individual's risk preferences should be stable with respect to small and moderate losses (Rabin, 2000). However, empirical evidence suggests that risk attitudes are often dependent on an individual's history of gains and losses. Prior losses have been

* Imas: Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh PA, aimas@andrew.cmu.edu. I thank Jim Andreoni, Linda Babcock, Nicholas Barberis, Aislinn Bohren, Gary Charness, Ayelet Gneezy, Uri Gneezy, Sebastian Kodritsch, Michael Kuhn, Thomas Langer, George Loewenstein, Ted O'Donoghue, Yuval Rottenstreich, Martin Weber, and Florian Zimmermann for invaluable input and useful comments, and Michael Rosenberg for excellent research assistance. I also thank numerous seminar participants for helpful feedback. The author declares that he has no relevant or material financial interests that relate to the research described in this paper.

found to affect subsequent risk-taking in a variety of settings, including choices over lotteries in laboratory experiments (Thaler and Johnson, 1990), trading decisions of experienced market-makers (Coval and Shumway, 2005; Liu et al., 2010) and investors (Kaustia and Knupfer, 2008; Andersen, Hanspal and Nielsen, 2015). However, research exploring dynamic effects has produced contradictory results: some studies find that individuals become more risk-seeking following losses (Andrade and Iyer, 2009; Langer and Weber, 2008), while other studies have found the opposite, that they become more risk-averse (Shiv et al., 2005; Liu et al., 2010).

This contradiction in empirical findings is mirrored in theoretical work on dynamic choice under uncertainty. For example, in the models of Barberis, Huang and Santos (2001) and Dillenberger and Rozen (2014), individuals respond to losses by taking on less risk, while in Shefrin and Statman (1985) and Weber and Camerer (1998) losses lead to more risk-taking. The extent of this inconsistency is encapsulated by the two opening quotations, both from papers that explore how prior losses affect risk attitudes. The first statement suggests that individuals take on *more* risk after losses, whereas the second posits the opposite, that losses lead to *less* risk-taking.

In this paper I propose a reconciliation of these seemingly inconsistent findings based on the distinction between losses that are *realized*, when money or another medium of value is transferred between accounts (e.g. selling a losing stock or cashing out and parting with the money after a loss), and those that are not realized, paper losses (e.g. holding a losing stock or not cashing out after a loss). Analyzing existing data and presenting new experimental evidence, I show that after a paper loss individuals become more likely to chase their losses and take on greater risk, while after a realized loss they take on less risk.

The first new study replicates prior findings and demonstrates the differential effect of realization in the same experiment: *Ceteris paribus*, the same loss is followed by less risk-taking if it is realized and by more risk-taking if it is not. The second experiment extends these results in two different populations, further tests the mechanisms behind realization and rules out alternative explanations. The third experiment demonstrates the differential effect of realized versus paper losses in an environment akin to a casino, where risk-taking is detrimental to wealth, and shows that giving individuals flexibility in realizing their asset positions leads to greater loss chasing and lower earnings than when realization is imposed exogenously. This study also demonstrates that the increase in risk-taking following paper losses is a deviation from individuals' ex-ante risk-taking strategies – after a paper loss they take on more risk than they had initially planned – and shows that realization mitigates this dynamic inconsistency.

Section II demonstrates that the main result of differential risk taking after realized versus paper losses follows from Cumulative Prospect Theory (Tversky and Kahneman, 1992) and the additional assumption on how prior losses are 'bracketed' with prospective risky choices (Read, Loewenstein and Rabin, 1999;

Rabin and Weizsäcker, 2009): namely, paper losses are integrated and evaluated jointly with prospects in the same bracket while realized losses are not. After experiencing a paper loss, an individual offered a risky prospect integrates the loss with the prospect and evaluates the two jointly. Due to loss aversion, the individual becomes *more* willing to accept a lottery if it offers a possibility of erasing the previous negative outcome. However, realization closes the bracket of prior outcomes; the individual internalizes the loss and updates the reference point, evaluating subsequent prospects relative to a new bracket. Since, without integration, risk taking no longer offers the possibility of recovering from the prior loss, individuals are predicted to take on less risk after a realized loss than after a paper loss. Additional predictions of the framework are also derived including how realization mitigates deviations from ex-ante risk-taking plans, i.e. dynamic inconsistency.

The proposed distinction in how individuals respond to realized versus paper losses has important implications for commitment and monitoring, particularly in contexts in which loss chasing has potentially negative consequences for wealth (e.g. casino gambling, losing stocks due to momentum). Barberis (2012) shows that Cumulative Prospect Theory generates a dynamic inconsistency in preferences: after experiencing losses, individuals take on more risk than they had originally planned. This may result in greater risk-taking than the individual, and potentially their employer, would have initially deemed optimal. Ebert and Strack (2014) argue that such dynamic inconsistency leads to unrealistic behavioral predictions – that individuals will chase their losses until they are bankrupt. Section II and the Appendix demonstrate that realization mitigates such dynamic inconsistency: since realized losses are not integrated with prospective risky choices, individuals will be more likely to follow their ex-ante optimal risk-taking strategies after realized losses than after paper losses. Section III provides direct empirical support for this prediction and demonstrates that flexibility in realizing one's position leads to greater deviations from the ex-ante optimal strategy than if realization was imposed exogenously.

In what follows, Section I reviews prior work on how losses affect risk attitudes and provides empirical support for the differential effect of realized versus paper losses on risk-taking. Section II shows how Cumulative Prospect Theory along with assumptions on bracketing can explain the discrepant findings regarding the impact of prior losses on risk-taking. Section III presents investment studies that provide direct support for the predictions. Section IV concludes with a discussion of implications for optimal monitoring and contracts. The Appendix provides a more detailed theoretical analysis and shows that the results are robust to a variety of behavioral specifications.

I. An Empirical Contradiction

In this section, I analyze existing empirical evidence on how prior losses affect risk attitudes and demonstrate a contradiction in the literature: some studies find that people take on more risk after a loss while other studies find the opposite – that individuals avoid risk after a loss. I then show that distinguishing between realized and paper losses reconciles these contradictory results. The section begins with studies conducted in the lab, and then discusses non-experimental data.

Langer and Weber (2008) adapted the investment game of Gneezy and Potters (1997) to study risk-taking in a dynamic context. In the game, individuals were asked to make a series of investment decisions from an initial endowment. Over the course of 30 rounds, participants could either invest part of their endowment in a risky, positive expected-value asset, or they could keep it. After each round, a randomization device determined whether the investment in the asset would be lost or multiplied. Importantly, if the investment was lost, the participant learned this information but did not part with it; there was no transfer of the money lost. All earnings were realized at the end of the experiment. If the participant ended the study with less money than the endowment, he would pay the difference to the experimenter; if he ended the study with more money than the endowment, the difference would be paid to him. In turn, participants in this study experienced *paper* gains and losses after each round prior to the end of the experiment.

I obtained and analyzed the individual-level data from the authors to examine risk-taking following a gain or a loss. As can be seen in Figure 1(a), overall investment in risk increased as the rounds progressed. I separated the data by investment after a loss and investment after a gain, and calculated the change in investment conditional on each outcome. Running a non-parametric Mann-Whitney test on the results reveals that the increase in risk-taking was driven by participants' responses to losses: individuals took on significantly *more* risk after a loss relative to the previous round ($p < .001$) and after a prior gain ($p < .001$).

Shiv et al. (2005) used a similar investment game to explore the effects of prior gains and losses on risk attitudes. As in Langer and Weber (2008), individuals started the experiment with an endowment and made a series of investment decisions over the course of 20 rounds. In each round, participants made the choice of either investing part of the endowment in a risky, positive expected-value asset, or keeping it. After each round, a randomization device determined whether the investment would be lost or multiplied. Unlike in Langer and Weber (2008), however, earnings were transferred after *every single round*. After each round, if the participant learned the investment was lost, he took that part of his endowment and transferred it to the experimenter; if the investment was multiplied, the participant received the positive difference. As such, participants in Shiv et al. (2005) experienced *realized* gains and losses after each round.

Figure 1(b) shows investment in risk over time from Shiv et al. (2005). Unlike in Langer and Weber (2008), overall risk-taking decreased as the rounds progress. A similar analysis to the one above reveals that changes in investment are again driven by a differential response to losses, but in the opposite direction. Relative to investment after a prior gain, participants responded to a loss by taking on *less* risk ($p < .01$) than before the loss, and did not change their investment behavior in response to gains.

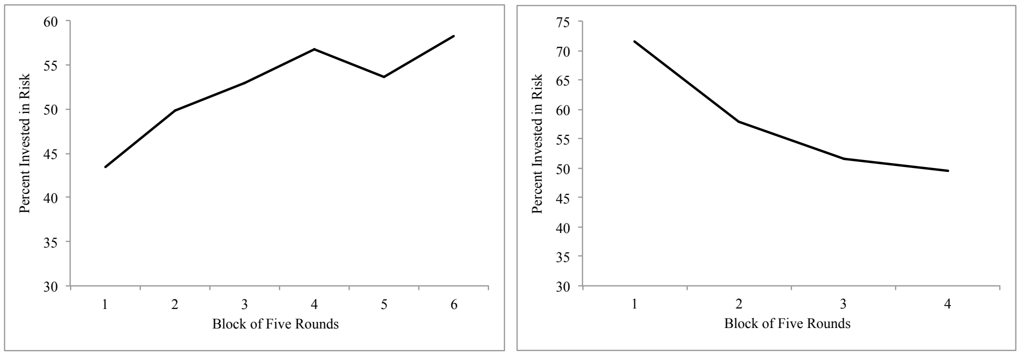

(a) *Langer and Weber* (2008) (b) *Shiv et al.* (2005)

FIGURE 1. Comparison of Risk-taking after Realized and Paper Losses

A review of the literature reveals that the distinction between realized and paper losses reconciles the contradictory findings. Andrade and Iyer (2009) gave participants an endowment of $10 and asked them to make two rounds of betting decisions on a gamble akin to a roulette wheel. Participants could bet up to $5 in each around. As in Langer and Weber (2008), earnings were realized at the end of the study, such that first round outcomes were paper gains or losses. The authors found that participants took on significantly *more* risk after a loss (Experiment 2). Using a similar design of sequential investment decisions with paper outcomes, Barkan and Busemeyer (1999) found similar results: individuals took on more risk after a loss than before the loss. In contrast, Shiv, Loewenstein and Bechara (2006) found that individuals took on less risk after a loss than before it; in their experiment, outcomes were realized after every round.[1]

Non-experimental studies examining the dynamics of risk attitudes present a similar contradiction. Coval and Shumway (2005) and Liu et al. (2010) studied the risk-taking behavior of professional traders. Both papers found that morning losses significantly affected risk-taking in the afternoon. Such behavior is not consistent with expected utility theory since, given the horizon of one day, wealth effects from prior losses were negligible, agency concerns were neutral, and other incentive effects can be ruled out. However, Coval and Shumway (2005) found that morning losses led traders to take on more risk in the afternoon, while Liu et al. (2010) found the opposite – morning losses led to less risk-taking in the afternoon. When discussing discrepancies between the two papers, Liu et al. (2010) note a

[1]It is important to note that in all studies examined, taking on risk can erase a prior loss if the gamble or investment is successful. As posited in Section II, this is a critical motivation for increased risk-taking after a paper loss. For example, Thaler and Johnson (1990) and Heath (1995) do not find support for loss chasing if the gamble does not allow the individual to offset the prior loss.

difference in the composition of paper and realized losses. A significantly larger portion of morning losses were realized in Liu et al. (2010) than in Coval and Shumway (2005). In light of the findings presented in this paper, the difference in the ratio of realized to paper losses may explain the contrast in the results.

Weber and Zuchel (2005) examined how the effect of prior outcomes on risk attitudes is influenced by whether participants made risky choices by betting on a lottery or deciding how many shares of an asset to buy. However, unlike the other studies in this section, people were asked how much risk they wanted to take *before* finding out the outcome of their prior choice; they were asked to make a contingent strategy of how much money to risk if the outcome of their prior choice was a gain and if it was a loss. Participants' plans were carried out after the outcomes of prior choices were determined. The authors found that contingent on a prior loss, participants stated they would like to decrease the amount of money risked relative to before the loss both in the lottery and investment frames. These results are consistent with the findings reported in Section III: individuals' planned risk-taking contingent on a prior loss was significantly lower than actual risk-taking after a (paper) loss.

Although the studies above provide suggestive evidence for the differential effect of realized versus paper losses on risk-taking, it is not causal. Besides differences in the nature of losses, there were several other discrepancies that may have contributed to the contrasting results (participants were drawn from different subject pools or populations, the risky assets had different distributions, etc). The next section demonstrates how Cumulative Prospect Theory (CPT) and an assumption on bracketing predicts the differential effect of prior realized versus paper losses on risk-taking, while Section III provides causal evidence for the framework's predictions.

II. Realization Effect

In one of the first papers to differentiate between realized and paper losses, Shefrin and Statman (1985) model an investor who engages in *mental accounting* (Thaler, 1985). A mental account defines which earnings and prospects are evaluated jointly within the same bracket (Read, Loewenstein and Rabin, 1999; Choi, Laibson and Madrian, 2009). Earnings and prospects associated with different mental accounts are 'bracketed' separately and evaluated independently. Shefrin and Statman (1985) posit that when an individual buys a stock or makes an investment, a mental account is opened. If the asset is sold for lower (higher) than the initial purchase price, the individual codes this as a realized loss (gain) and closes the associated mental account. On the other hand, if the asset goes down (up) in value without being sold, the individual codes this as a paper loss (gain) and the mental account remains open. Individuals experience a pronounced drop in utility when a mental account is closed at a loss relative to the same loss while the mental account is open; the authors argue that a realized loss is more painful than a paper one. Barberis and Xiong (2012) formalize the distinction between paper and realized outcomes as defined by Shefrin and Statman (1985) in a model of realization utility (see Ingersoll and Jin (2013) for an extension of this framework).[2]

In this paper, I define *realization* as an event in which money or another medium of value is transferred between accounts. These accounts could be real, (e.g. brokerage, savings), or mental accounts. For example, consider an individual who accepts a positively skewed gamble (where the potential upside is greater than the downside), loses, and is then offered the same gamble again. The loss is realized if he parts with the money after the first gamble – if the cash leaves his account and is transferred to another. If he does not part with the money lost before making a decision about whether to accept or reject the second gamble, then it is defined as a paper loss. Note that this definition subsumes the definition of Shefrin and Statman (1985): selling a losing stock transfers an amount smaller than the original purchase price to another account, e.g. the individual's brokerage account (realized loss), whereas holding a losing stock does not involve a transfer (paper loss).[3]

The studies presented in the next section test four predictions: (1) risk-taking will be greater after a paper loss than after a realized loss; (2) risk-taking will be greater after a paper loss than before the paper loss; (3) risk-taking will be lower after a realized loss than before the realized loss; and (4) people will deviate from their ex-ante risk-taking plans to take on more risk after a paper loss and realization mitigates this dynamic inconsistency. I show how these predictions can be derived from CPT (Tversky and Kahneman, 1992) applied to a sequential three-period decision problem in which the individual is offered two opportunities to accept or reject a positively-skewed lottery. I outline the conditions under which the predictions hold and discuss how potential modifications to the theory suggested in the literature may affect them. The Appendix presents a more detailed analysis of the decision problem, including extensions and robustness checks.

A critical issue in addressing the consequences of realization is whether an individual considers each lottery one at a time or forms a plan of how much risk to take in each period conditional on the prior outcome. In the former (myopic) case, the individual does not consider the second lottery when making the first choice; in the latter (non-myopic) case, he forms an ex-ante, fully contingent strategy of whether to accept or reject the lottery at each opportunity. The first two predictions follow in both cases from a straightforward application of CPT. Prediction 3 holds in the non-myopic case under a broad set of conditions; in

[2]A recent study by Frydman et al. (2014) provides neural support for realization utility: activation in brain regions associated with utility shocks was more significant after a realized outcome than after a paper outcome.

[3]The definition offered here is not meant to capture all instances which determine when a mental account is closed or opened. This would involve a formal theory of mental accounting that outlines the necessary and sufficient conditions for these processes, and is outside the scope of the current paper. Rather, it is argued that realization as a transfer is a sufficient condition for predicting when a loss will lead to less risk-taking versus more. See Section IV for further discussion.

the myopic case, the prediction follows from one of several possible modifications of CPT, including the assumption of sensitization (Barberis, Huang and Santos, 2001; Thaler and Johnson, 1990) applied to realized losses. Given that Prediction 4 involves departures from a plan the individual makes with respect to sequential risk-taking, it only applies to the non-myopic case where the individual forms such a plan, and holds generally in that case.

The key assumption driving these predictions is that a prior realized loss is not integrated and evaluated jointly with a prospective risky choice. After a paper loss, the individual may feel that there is still hope that the upside of a lottery will allow him to avoid the pain of parting with the amount lost. When evaluating prospects, paper losses are integrated with the potential payoffs, and, as such, a lottery whose upside allows the individual to erase the prior losses becomes more attractive since rejecting it results in a sure outcome of a realized loss. The realization of a loss serves as a natural point for an individual to internalize the negative outcome, update the reference point and close the associated mental account. Since the prior loss is not integrated, accepting risk no longer affords the possibility of avoiding a negative transfer. In turn, the individual is less likely to accept the lottery after a realized loss than a paper one.

To set up the basic framework, allow preferences to satisfy the standard assumptions of CPT (Tversky and Kahneman, 1992). Let $(x^1, p^1; ...x^n, p^n)$ represent a lottery with n possible outcomes, where x^i is the outcome with objective probability p^i, $x^i > x^j$ iff $i > j$, and $\sum_{i=1}^n p^i = 1$. The decision maker (DM) evaluates the lottery as $\sum_{i=1}^n \pi^i \ V(x^i|r)$ relative to reference point $r \in \mathbb{R}$, where π^i and $V(x^i|r)$ are the decision weight and output of the value function V, respectively, used to evaluate outcome x^i. Let $V(x^i|r)$ satisfy the standard assumptions of a prospect theory value function, such that

$$V(x^i|r) = \begin{cases} v(x^i - r) & if \ x^i \geq r \\ -\lambda v(-(x^i - r)) & if \ x^i < r \end{cases} \tag{1}$$

where $V(r|r) = 0$, v is concave, and $\lambda > 1$ implies loss aversion. Note that this implies that the value function V is concave for gains $(x^i \geq r)$ and convex for losses $(x^i < r)$. The probability weighting function $w : [0, 1] \to [0, 1]$ transforms objective probabilities into decision weights π, such that $\pi_i = w(p^i + ... + p^n) - w(p^{i+1} + ... + p^n)$ for $x^i \geq r$ and $\pi_i = w(p^1 + ... + p^i) - w(p^1 + ... + p^{i-1})$ for $x^i < r$, with $w(0) = 0$ and $w(1) = 1$. The weighting function w is assumed to satisfy the reflection property, assigning the same weight to a given gain-probability as to a given loss-probability (Prelec, 1998; Tversky and Kahneman, 1992). Unlike in the original prospect theory of Kahneman and Tversky (1979), in CPT probability weighting is rank-dependent. Weighting functions proposed in the literature are typically non-linear, S-shaped transformations of objective probabilities e.g. $w(p) = exp(-(-ln \ p)^\alpha)$ from Prelec (1998), reflecting the observed overweighting of small probabilities and the underweighting of large probabilities in empirical studies.

Let there be three periods. The DM is offered a mixed, positively skewed lottery $L = (x^g, p; x^l, 1-p)$ in the first of two periods, where $p < .5$, $x^g > 0 > x^l$ and $x^g > |x^l|$. If the DM accepts L in the first period, he learns the outcome and proceeds to the next period to make the same choice over L again; if he rejects, no further lotteries are offered. An outcome of a choice is realized (or not) in the beginning of the next period, and all outcomes are realized in the third period. As in Shefrin and Statman (1985) and Barberis and Xiong (2012), the DM derives ex-post utility only from realized outcomes.[4]

Myopic Risk-Taking

If the DM evaluates each lottery choice one at a time, without considering the second lottery when making the first choice, in each period he evaluates a prospect with two possible outcomes. As such, CPT coincides with the original version of prospect theory of Kahneman and Tversky (1979), where $\pi^g = w(p)$ and $\pi^l = w(1-p)$. Without assuming a particular functional form, let $w(p)+w(1-p) \leq 1$, a condition met by all weighting functions proposed in the literature, e.g. Prelec (1998), and following Barberis (2012) and Tversky and Kahneman (1992), take the reference point r to be the status quo ($r = 0$).

In the first period, the DM evaluates accepting the lottery or rejecting it and retaining the status quo. The DM accepts the lottery if

$$0 < w(p)v(x^g) - \lambda w(1-p)v(-x^l). \tag{2}$$

Although it seems counterintuitive given the assumption of loss aversion, individuals with prospect theory preferences may accept actuarially fair (or even unfair) mixed lotteries if they are positively skewed due to probability weighting, which leads them to overweigh low probability outcomes.[5] To analyze the effects of a prior loss on subsequent risk-taking, assume that the DM chooses to accept the first lottery.

Suppose that the DM suffers a loss that is not realized – a paper loss – and integrates it with the prospect when offered the second lottery. Since all outcomes will be realized at the beginning of the final period, now the DM compares the valuation of accepting the second lottery, $w(p)v(x^g+x^l)-\lambda w(1-p)v(-2x^l)$, which allows him to potentially avoid realizing the prior loss, to rejecting the lottery and realizing the loss with certainty, $-\lambda v(-x^l)$. The DM accepts the lottery if

$$0 < w(p)v(x^g + x^l) - \lambda w(1-p)v(-2x^l) + \lambda v(-x^l). \tag{3}$$

[4]The results also hold if the utility derived from a realized outcome is greater than from a paper one.

[5]Barberis (2012) shows that individuals with CPT preferences will even accept actuarially fair lotteries with even odds as long as they are offered enough such lotteries in a sequence to formulate an ex-ante strategy that generates a positively skewed lottery over final, accumulated earnings. The next subsection and the Appendix discuss this in more detail.

Suppose that the DM accepted the first lottery and suffers a realized loss. In the second period, he is offered the same lottery again. The prior loss is not integrated with the prospect and the DM accepts the lottery if (2) is met.

Since the same condition (2) specifies if the DM accepts the lottery both before a loss and after a realized loss, if Prediction 2 holds, Prediction 1 holds as well. For Prediction 2 to hold, it is necessary to demonstrate that if the DM accepts L before a loss, even when indifferent, he would always be willing to accept L after a paper loss; specifically, that the DM's valuation of accepting the lottery (relative to rejecting it) is greater after a paper loss than before the loss:

$$\lambda > \frac{w(p)(v(x^g) - v(x^g + x^l))}{w(1-p)(v(-x^l) - v(-2x^l)) + v(-x^l)}. \tag{4}$$

Condition (4) holds for any level of loss aversion, $\lambda > 1$ (see Appendix for proof).

For Prediction 3 to hold in the myopic case, such that the DM takes on less risk after a realized loss not only relative to a paper loss but relative to before the loss, more structure is needed. A number of different mechanisms discussed in the literature lead to Prediction 3 when applied to realized losses, including, *sensitization* (Barberis, Huang and Santos, 2001; Thaler and Johnson, 1990), a diminished capacity for dealing with bad "news" (Koszegi and Rabin, 2009; Pagel, 2012; Linville and Fischer, 1991), the increased salience of the potential downside of risk (Bordalo, Gennaioli and Shleifer, 2012), or a change in mood (Loewenstein, 1996). Since (2) specifies the DM's willingness to accept the lottery both before and after a realized loss, it is straightforward to show that any of these factors that produce a greater distaste for losses after a realized loss lead to Prediction 3.[6]

Non-Myopic Risk-Taking

If the second lottery is taken into account in the first period, with the DM formulating an ex-ante strategy that is a fully contingent plan of risk-taking, the decision problem changes significantly. Barberis (2012) examines such a decision problem with paper outcomes, demonstrating the difficulty of finding an analytical solution given non-linear probability weighting; this section and the Appendix build on his proposed framework and use similar simulation techniques to outline the conditions under which the predictions hold. Note that Prediction 1 – the differential effect of realized versus paper losses on risk-taking – holds under the same, general conditions in both the myopic and non-myopic case.

Prior to the first choice, an example of a contingent plan could be to accept the lottery in each period regardless of the prior outcome; another could be to accept the first lottery, then contingent on winning, to accept the second lottery and to reject it otherwise. The latter strategy is termed a "loss-exit" plan. As in

[6]The Appendix provides further discussion and a formal illustration of sensitization applied to realized losses.

prior work on CPT in dynamic settings (Barberis, 2012; Ebert and Strack, 2014), assume that the DM makes a choice believing he will follow the optimal strategy and is naïve about any dynamic inconsistency in his preferences, i.e. subsequent deviations from the strategy.[7] Here, consider the case where realization in the second period is not anticipated. The Appendix provides an analysis of the case where realization is anticipated and shows that the predictions hold in both cases.

In the first period, the DM evaluates the utility over the distribution of *accumulated* earnings induced by each available strategy; these accumulated earnings represent potential outcomes that will be realized in the last period. Strategies may generate lotteries over accumulated earnings with more than two outcomes; for example, the "loss-exit" plan generates the lottery $L_{loss-exit} = (2x^g, p^2;\ x^g + x^l, p(1-p);\ x^l, 1-p)$. The DM compares the valuation of each strategy and chooses to accept or reject the first lottery according to the ex-ante optimal strategy that maximizes utility. The DM's valuation of the "loss-exit" plan, $w(p^2)v(2x^g) + [w(p) - w(p^2)]v(x^g + x^l) - \lambda w(1-p)v(-x^l)$, is compared to the valuation of every other available strategy including rejecting the first gamble and retaining the status quo. Note that the "loss-exit" plan maximizes the possible upside while minimizing the downside and generates a more positively skewed lottery than the one-shot lottery L; probability weighting makes this plan particularly attractive. As shown in Barberis (2012) and the Appendix, if the DM accepts the first lottery, the optimal strategy is the "loss-exit" plan under the vast majority of preference parameter values including structural estimates from prior work (e.g. Tversky and Kahneman, 1992).

To analyze the effects of a prior loss on risk-taking, suppose the DM accepts the first lottery with the aim of following the "loss-exit" plan. If the DM suffers a paper loss, he accepts the second lottery if (3) is met. Alternatively, if the loss is realized, the lottery is accepted if (2) is met. As in the myopic case, Prediction 1 follows if (4) holds, which it does for any level of loss aversion, $\lambda > 1$.

For Prediction 2 to hold, the DM's valuation of accepting the second lottery (relative to rejecting it) after a paper loss must be higher than the valuation of accepting the first lottery as part of the "loss-exit" plan, expression (5) below; for Prediction 3, the relative valuation of the lottery after a realized loss must be lower than that of the "loss-exit" plan, expression (6) below.

$$\lambda > \frac{w(p^2)[v(2x^g) - v(x^g + x^l)]}{w(1-p)[v(-x^l) - v(-2x^l)] + v(-x^l)}. \tag{5}$$

$$0 < w(p^2)v(2x^g) + [w(p) - w(p^2)]v(x^g + x^l) - w(p)v(x^g). \tag{6}$$

Both (5) and (6) are met under a broad set of conditions (see Appendix, section

[7] See also Loewenstein, O'Donoghue and Rabin (2003), O'Donoghue and Rabin (1999) and Della Vigna and Malmendier (2006) for discussion on naiveté in other domains.

1.2). Specifically, both hold for the benchmark conditions used in simulations of the decision problem in Barberis (2012) – the structural estimates of Tversky and Kahneman (1992) – as well as the range of estimates found in other work (Abdellaoui, 2000; Gonzalez and Wu, 1999; Tanaka, Camerer and Nguyen, 2010). Note that unlike the myopic case, additional assumptions such as sensitization are not required to show that Prediction 3 holds broadly.

Since the DM accepted the first lottery with the aim of following the "loss-exit" plan, being more willing to take on risk after a paper loss constitutes a deviation from the DM's ex-ante optimal strategy. This dynamic inconsistency is due to probability weighting and the difference in skewness of options available to the DM in the first period versus the second. Intuitively, prior to accepting the lottery, the "loss-exit" plan generates a skewed distribution over accumulated earnings. Probability weighting makes the plan more attractive than the one-shot lottery by overweighting the low probability outcome of two wins in a row and underweighting the more likely outcome of experiencing a loss. After a paper loss, the available strategies generate a less skewed distribution over accumulated earnings than prevailed in the first period; the DM compares accepting the second lottery to experiencing the sure loss if the lottery is rejected. Since taking the risk allows the DM to avoid the negative realization, loss aversion leads him to accept greater risk than originally planned (Prediction 2). As with a paper loss, after a realized loss the strategy of accepting the lottery generates a less positively skewed distribution than in the previous period, making it less attractive. However, since the realized loss is not integrated, the DM is no longer motivated to accept risk in order to recover from the prior outcome and hence is less willing to take it (Prediction 3).

For Prediction 4, dynamic inconsistency is observed when the DM accepts the first lottery with the aim of following the "loss-exit" plan, but after a loss, values accepting the second lottery more than rejecting it. After a paper loss, the DM deviates by accepting the second lottery if expression (3) holds; after a realized loss, he deviates if expression (2) holds. However, for any level of loss aversion $\lambda > 1$ expression (4) holds, implying that the DM is less willing to accept the second lottery after a realized loss than a paper one. In turn, fewer deviations of greater risk-taking should be observed when losses are realized than when they are not. The next section tests the predictions directly.

III. Investment Experiments

As discussed in Section I, prior work using the investment game of Gneezy and Potters (1997) has found that a loss leads to more risk-taking (Langer and Weber, 2008) and less risk-taking (Shiv et al., 2005) depending on whether it was transferred or not. To test whether paper versus realized losses lead to the differential effects on subsequent risk attitudes, I adopted the same game to both replicate these findings within the same experiment and further examine how the nature of a prior loss affected individuals' risk-taking in a sequence of investment decisions. The decision maker receives an endowment, $\$E$, and makes investment decisions over a series of rounds. In each round, he can choose how much of an amount, X, he would like to invest in a risky option and how much to keep. The amount invested in the risky option, x, yields a dividend of kx ($k > 1$) with probability p and is lost with probability $1 - p$. The money not invested ($X - x$) is kept by the decision maker. The expected payoff in each round is:

$$p \cdot (X - x + kx) + (1 - p) \cdot (X - x).$$

After the choice of x is made, the outcome of the risky option is determined and revealed to the decision maker. The decision maker then moves on to the next round where he is presented with the same choice.

The amount invested x provides a robust metric for capturing treatment effects and differences in attitudes toward risk. Similar paradigms have been used to test for myopic loss aversion in students (Gneezy and Potters, 1997) and professional traders (Haigh and List, 2005), to demonstrate decreased (and increased) risk-taking following a prior loss (Shiv et al., 2005; Langer and Weber, 2008), to examine gender differences in risk attitudes (Charness and Gneezy, 2012) and to show the effect of ambiguity aversion and illusion of control on portfolio choice (Charness and Gneezy, 2010).

A. *Study 1*

Realized and Paper Losses

The first study aimed to identify the differential effect of realized versus paper losses on risk-taking by replicating prior results within the same experiment. Undergraduates ($N = 128$) from a university-wide subject pool were recruited to participate in an experiment on decision-making.[8] All were given a \$5 show up fee at the end of the experiment. Participants were randomly assigned to individual computer stations and given a set of instructions that was read aloud. Each person was endowed with an account of \$8 in an envelope in the beginning of the study and asked to count it: the envelope contained 7 one-dollar bills and 4 quarters.

Participants were told that they would make 4 rounds of investment decisions. In each round the participant would decide how much of \$2 to invest in a lottery (in increments of quarters). With a 1/6 chance the lottery would succeed and pay dividends $k = 7$ times the amount invested x; with a chance of 5/6 the lottery would fail and the money invested would be lost. In each round, participants were randomly assigned one "success number" between 1 and 6. This number was displayed on their computer screen in the beginning of each round. Participants would then enter the amount they would like to invest x. Note that in this case, p (1/6) and k (7) were chosen such that $p \cdot k > 1$, making the expected value of investing higher than the expected value of not investing.

[8]See Appendix for all experimental instructions.

Once this was done, the experimenter rolled a six-sided die in the front of the room. Participants were welcomed to examine the die to make sure it was fair. If the outcome of the die roll matched the participant's success number, the lottery would succeed and they would earn $7x$ plus the amount they did not invest $(2-x)$. If the outcome was any other number, the lottery would fail and participants would earn the amount they did not invest. After learning the outcome of the die roll, participants would move on to the next round, be assigned a new "success number" and make the same decision again. All die roll outcomes were written on a board in front of the room to keep information constant between treatments.

To test for the differential effect of realized versus paper losses, participants were randomized into either the Realized or Paper treatment. In the Realized treatment, at the end of the 3^{rd} round participants had their wealth positions realized: if they had lost money by the end of the 3^{rd} round, they took this amount out of their envelope and transferred it out of their endowed account by handing it to the experimenter. If they had won, this amount was given to them. After realizing their earnings, participants made one last investment decision in the 4^{th} round and were paid according to the outcome.

In the Paper treatment, participants did not realize their earnings at the end of the 3^{rd} round. They continued on to the 4^{th} round and were paid at the end of the experiment. Time between rounds was normalized across treatments such that the break between each round was on average how long it took to realize earnings. As such, those in the Realized treatment did not have a longer break between the 3^{rd} and 4^{th} rounds than those in the Paper treatment.

A second Paper Social (Paper S) treatment was run as a robustness check to ensure that the intervention and interruption of the experimenter in the Realized treatment did not drive the results. The procedure in Paper S was the same as in the Realized treatment – at the end of the 3^{rd} round the experimenter came up to each participant and verbally informed them how much money they had won or lost relative to the original endowment – but no money was transferred between the participant and the experimenter. Namely, although participants in both treatments were similarly interrupted and informed of their earning, no transfer occurred in the Paper S treatment.[9]

Note that given a sequence of decisions and outcomes, the wealth positions and information were the same for participants in all three treatments. However, those who had lost money from their \$8 endowment by the end of the 3^{rd} round in the Realized treatment parted with it by having this amount transferred out of their envelope. This served as the manipulation of exogenously inducing realization. In contrast, those who had a similar loss by the end of the 3^{rd} round in the Paper and Paper S treatment could still potentially avoid parting with their endowment and experiencing a negative realization by taking on more risk in the 4^{th} round.

In this setup, the predictions are that those who were losing at the end of

[9]The Paper S treatment was run as a robustness check after running the Paper and Realized treatments, using the same lab and subject pool.

the 3^{rd} round in the Paper and Paper S treatments should take on *more* risk in the 4^{th} round compared to those who were similarly losing in the Realized treatment. Relative to 3^{rd} round investment, participants who had lost in the two Paper treatments should *increase* their 4^{th} round investment while those in the Realized treatment should *decrease* their 4^{th} round investment.

Results To analyze the effect of prior losses on risk-taking, I examine the change in investment between rounds 3 and 4 for participants who lost the lottery in rounds 1, 2 and 3 – the lottery failed all three times.[10] Figure 2 shows the investment changes by treatment. Looking at the effect of losses between treatments, consistent with the first prediction the change in investment between the 3^{rd} and 4^{th} rounds was significantly different between the Paper and Realized treatments ($t(51) = 3.19, p < .01$). Importantly, this effect did not seem to be driven by the additional interruption or interaction with the experimenter: the same difference in risk-taking was significant between the Paper S and Realized treatments ($t(63) = 3.25, p < .01$). To quantify the effect of realizing a loss on risk-taking, an OLS regression of Investment Change on a treatment dummy (Realized = 1; Paper = 0) for those who lost by the end of the 3^{rd} round revealed that realizing one's loss leads to a significant decrease in risk-taking relative to not realizing a loss of the same size ($\beta = -.38, p < .01$). Changes in investment did not differ between treatments for any other round.

Examining the effect of prior losses within treatment, consistent with the second prediction, individuals who had lost in the Paper treatment increased their investment in the lottery by \$0.23, taking on significantly *more* risk than the null hypothesis of zero change in investment ($t(26) = 2.28, p = .03$).[11] However, in line with the third prediction, those who had similarly lost in the Realized treatment decreased their investment by \$0.15, taking on significantly *less* risk than specified by the null hypothesis ($t(25) = 2.42, p = .02$). As in the Paper treatment, those who had lost in the Paper S treatment increased their investment in the lottery by \$0.16, also taking on significantly more risk than the null ($t(38) = 2.41, p = .02$).[12]

B. *Study 2*

The second study further examined the mechanism of realization as a transfer between accounts and tested the robustness of the results. First, it examined whether the decrease in risk-taking after parting with the money lost was a general phenomenon following a transfer out of one's account or whether the physical separation was required. Particularly, the experiment aimed to test whether the

[10] Across the studies, investments did not significantly differ by treatment in rounds 1 through 3. See Appendix for total investment and changes in investment by round and treatment.

[11] One-sample t-test with null hypothesis of zero investment change.

[12] Since the experiment was designed to test the effect of prior losses, probabilities were chosen such that most participants had lost by the end of the 3^{rd} round. There were too few observations to conduct any meaningful analyses on prior gains.

FIGURE 2. Investment Change ($) after a Realized versus Paper Loss

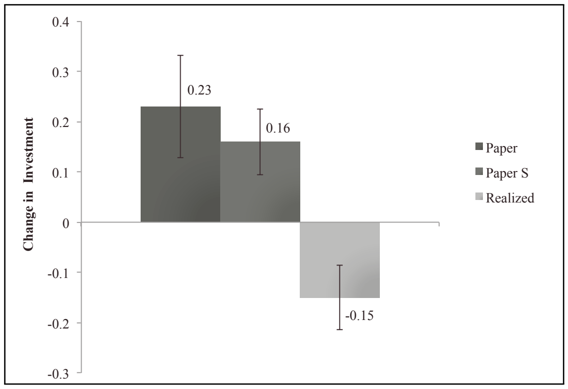

framework outlined in Section II applied more generally to contexts where money is transferred between accounts rather than specifically to situations involving physical cash transactions. In addition, supplementary to the Paper S treatment, the study tested whether the decrease in risk-taking in the first experiment was due to a transfer of money or an artifact of an interruption effect. Lastly, the second study replicated the first in a more representative population that varied in age and income.

Realization as a Transfer

Individuals ($N = 246$) from a participant pool open to the general population were recruited for an experiment on decision-making. The average age of participants was 26.9 and the average income was $18,000. All were given a $5 show up fee at the end of the experiment. Procedures were largely the same as those in the first experiment: participants were randomly assigned a "success number" from 1 to 6 at the beginning of each round and the lottery yielded 7 times the amount invested $x with a 1/6 probability, and lost the amount invested with probability 5/6.

Participants were randomly matched into one of 4 treatments. The Realized and Paper S treatments were identical to those in the first study. The Interrupt treatment was similar to the Paper S treatment except that after participants learned the outcomes of their third round investments, each was given a filler task for 5 minutes. The filler task consisted of solving anagrams of moderate difficulty for no additional monetary incentives. At the end of the 5 minutes, as in the Paper S treatment the experimenter informed each participant about their earnings up until that point but did not realize their position. Participants then made a final investment decision in the 4^{th} round.

The Transfer treatment was designed to test if the physical exchange of money is necessary for the effect of realization or whether a transfer between accounts is sufficient. Instructions were similar to the Realized treatment except that participants were told they were endowed with \$8 from the experimenter but were not actually given the cash in an envelope. At the end of the 3^{rd} round, earnings up until that point were presented. Each participant was then informed that to finalize the earning position for the three rounds, any amount lost would be withdrawn from their endowed account and transferred back to the experimenter; alternatively, any amount won would be deposited into their endowed account. Participants were asked to type in the word "Closed" into a text box on the computer screen to finalize their earnings position for the 3 rounds. Each then moved on to make a final 4^{th} round investment decision where, as in all other treatments, up to \$2 could be invested in the lottery.

The predictions state that those who had a loss by the end of the 3^{rd} round in the Paper S and Interrupt treatments, where the outcomes are not realized, should change their investment by taking on *more* risk; those who similarly lost in the Realized and Transfer treatments, where the outcomes are realized, should respond by taking on *less* risk.

Results Figure 3(a) shows the investment changes by treatment. Examining the effect of a prior loss between treatments, the results of the first study were replicated: investment change after a loss at the end of the 3^{rd} round was significantly different between the Realized and Paper S treatments ($t(57) = 2.91, p < .01$). An interruption did not seem to drive this effect, risk-taking after a loss was significantly different between the Interrupt and Realized treatments as well ($t(60) = 3.02, p < .01$). Additionally, a physical transfer of money was not necessary for the effect of realization: risk-taking after a loss in the Transfer treatment differed significantly from risk-taking after a similar loss in the Paper S treatment ($t(56) = 3.32, p < .01$). Examining the size of the effect, those who lost in the Realized treatment responded by taking on significantly less than those who similarly lost in the Paper S treatment ($\beta = -.47, p < .01$). A loss in the Transfer treatment had an analogous effect, with participants responding by taking on significantly less risk than those who similarly lost in the Paper S treatment ($\beta = -.34, p < .01$) and the Interrupt treatment ($\beta = -.35, p < .01$). Changes in investment did not differ between treatments for any other round.

Looking at the effect of prior losses within treatment, in line with the second prediction, those who had lost by the end of the 3^{rd} round in the Paper S treatment increased their investment by \$0.19, taking on significantly more risk than the null hypothesis of zero investment change ($t(27) = 2.64, p = .01$). Consistent with the third prediction, those who had similarly lost in the Realized treatment decreased their investment by \$0.28, taking on significantly less risk than the null ($t(30) = 2.00, p = .05$). As in the Paper S treatment, those who had lost in the Interrupt treatment increased their investment by \$0.20 ($t(29) = 2.65, p = .01$). The Transfer treatment produced a similar decrease in risk-taking after a loss as in the Realized treatment, with participants decreasing their investment by \$0.14 ($t(30) = 2.04, p = .05$).

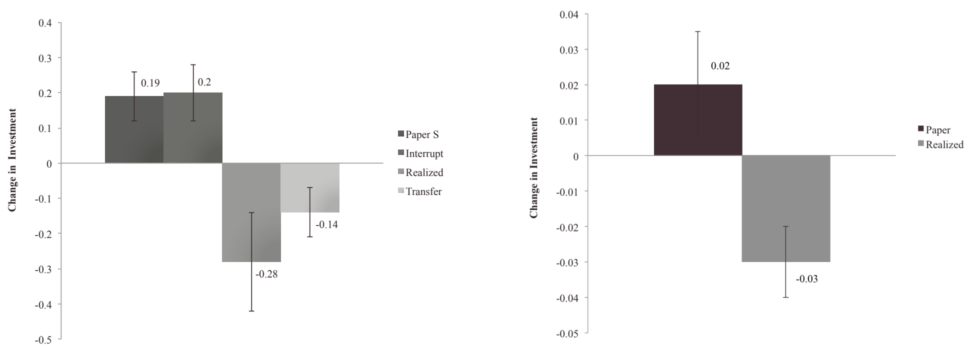

(a) *Lab Experiment* (b) *Mechanical Turk Experiment*

FIGURE 3. Investment Change ($) after a Realized versus Paper Loss

Robustness Check

To check the robustness of the results, a separate experiment was run using the online labor market on the Amazon Mechanical Turk platform. Individuals (N=151) were recruited from Mechanical Turk for an experiment on decision-making. In keeping with payments typically seen on the platform, participants were given a flat show up fee of $0.30. Each was endowed with $1.00 and told that they would make 4 rounds of investment decisions. In every round, participants decided how much of $0.25 (in increments of $.01) to invest in the same lottery as above. After being assigned one "success number" between 1 and 6 and making their investment decision, the participant rolled a virtual die that returned an outcome using a random number generator. Earnings were calculated as before: the investment was multiplied by 7 if the outcome matched the assigned success number; if it did not match, the investment was lost. A new success number was assigned for each round.

Participants were randomly assigned to one of two treatments. The Realized treatment was similar to the Transfer treatment above. At the end of the 3^{rd} round, participants were told their earnings, and that to finalize their position up until that point, any amount lost would be withdrawn from their endowed account and transferred to the experimenter; any amount won would be transferred to their endowed account. As in the Transfer treatment, participants were asked to type in the word "Closed" into a text box to finalize their earnings position for the 3 rounds and moved on to the 4^{th} round to make a fourth investment decision. There, they could again invest up to $0.25 in the lottery.

The Paper treatment was similar to the Paper S treatment above. At the end of the 3^{rd} round participants were interrupted and led to a screen displaying their earnings up until that point. They were not informed of any transfer and continued on to the 4^{th} round.

As shown in Figure 3(b), the results were consistent with the predictions. Looking between treatments, changes in investment for those who had lost were significantly different between the Realized and Paper treatments ($t(85) = 3.32, p = .001$). Within the Paper S treatment, individuals who experienced a loss at the end of the 3^{rd} round increased their investment by \$0.02, taking on more risk than the null of no investment change ($t(45) = 2.44$, $p = .02$), while those in the Realized treatment decreased their investment by \$0.03 ($t(40) = 2.29$, $p = .03$).

Note that despite the differences in the size of the stakes involved, responses to prior losses as a percentage change relative to investment in the previous round were similar in magnitude across the two experiments in Study 2. Those who had lost in the lab study ($X = \$2$) increased their investment by 22% in the Paper S treatment and decreased it by 28% in the Realized treatment. Participants who had lost in the online study ($X = \$0.25$) increased their investment by 24% in the Paper treatment and decreased it by 25% in the Realized treatment.

C. Study 3

The third study explored whether giving individuals flexibility in when to realize their positions could lead to lower earnings overall relative to those whose positions were exogenously realized. Particularly, the experiment was designed to mimic environments where taking on risk and chasing losses leads to lower expected returns than keeping one's money (e.g. casinos, race-tracks), and the choice to realize one's position is endogenous.[13] The second prediction in Section II states that individuals should avoid realizing their losses and instead choose to take on greater risk than in the previous round. Hence, in this context, flexibility in realization is predicted to lead to lower expected earnings. The second part of the study demonstrates that this behavior is due to a dynamic inconsistency: prior to making the first choice, individuals plan to take on less risk after a loss. After the loss actually occurs, however, they deviate from their strategy to take on more risk. On the other hand, imposing realization exogenously should mitigate such deviations.

Realization and Flexibility

Undergraduates (N=120) from a university-wide subject pool were recruited to participate in an experiment on decision-making. All were given a \$5 show up fee at the end of the experiment. The lottery was set to yield 2.5 times the amount invested \$x with a 1/3 probability, and to lose with probability 2/3. Since $p \cdot k < 1$, the expected value of investing in the lottery is slightly lower than the expected value of not investing, similar to gambling on a roulette wheel. Procedures were

[13]Although it seems counterintuitive that individuals with prospect theory preferences, i.e. first order risk aversion, would gamble in such environments, Barberis (2012) and Ebert and Strack (2014) demonstrate how a dynamic version of prospect theory predicts that for a wide range of parameters, individuals will accept positively skewed (or even fair, unskewed) gambles such as the ones in this study even if the overall expected value of any single gamble is negative.

largely the same as those in the first experiment, except that now participants were randomly given two different "success numbers" from 1 to 6 at the beginning of each round to reflect the higher probability of the lottery succeeding.

In addition to the Paper and Realized loss treatments, a third Flexible treatment was added to test whether flexibility in realization indeed reduced expected earnings relative to when realization was exogenously imposed due to individuals choosing to delay realization and taking on greater risk after a loss. In the Flexible treatment, individuals were asked at the end of the 3^{rd} round whether they would like to realize their earnings similar to those in the Realized treatment, or to move on to the 4^{th} round. If they chose to realize their positions, the procedure was identical to the Realized treatment; if they chose to move on, the procedure was identical to the Paper treatment.

Both those in the Paper and Flexibility treatments are predicted to take on more risk after a loss than those in the Realized treatment. Those who had a paper loss by the end of the 3^{rd} round should change their investment by taking on more risk, while those who had a realized loss should decrease their investment and take on less risk.

Results Figure 4(a) shows the change in investment after a loss by treatment. In line with the first prediction, investment changes after a loss between the 3^{rd} and 4^{th} rounds were significantly different between the Realized and Flexible treatments ($t(48) = 3.68, p < .001$) and the Realized and Paper treatments ($t(50) = 3.44, p = .001$), but not between the Paper and Flexible treatments ($t(48) = .37, p = .71$). Participants who lost in the Realized treatment decreased their investment relative to both the Paper treatment ($\beta = -.57, p < .001$) and the Flexible treatment ($\beta = -.60, p < .001$). Changes in investment did not differ between treatments for any other round.

Consistent with the second prediction, those who had lost by the end of the 3^{rd} round in the Paper treatment increased their investment in the lottery by \$0.29, taking on significantly *more* risk than the null ($t(25) = 2.75, p = .01$). Similarly, those in the Flexible treatment who had lost by the end of the 3^{rd} round *also* took on significantly more risk, increasing their investment in the lottery by \$0.33 ($t(23) = 3.14, p < .01$). Importantly, participants in the Flexible treatment had lost by the end of the 3^{rd} round chose to realized their positions only 13% of the time. This was not driven by participants' desire not to be interrupted: those who had won realized their positions 44% of the time, a significant difference ($t(38) = 2.33, p = .025$). In contrast, and in line with the third prediction, those who had lost by the end of the 3^{rd} round in the Realized treatment took on *less* risk and decreased their subsequent investment by \$0.27 ($t(25) = 2.19, p = .038$).

Giving individuals flexibility in realizing their positions led to greater losses than when realization was imposed exogenously. Since investing in the lottery had a lower expected return compared to not investing, participants in the Paper and Flexible treatments stood to earn less in the 4^{th} round than those for whom realization was imposed exogenously. Indeed, expected 4^{th} round earnings in the

Realized treatment were significantly higher than in both the Paper and Flexible treatments ($t(78) = 2.68, p < .01$ and $t(79) = 3.49, p < .001$, respectively).

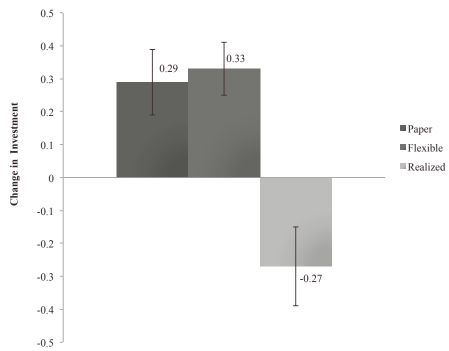

(a) *Realization and Flexibility*

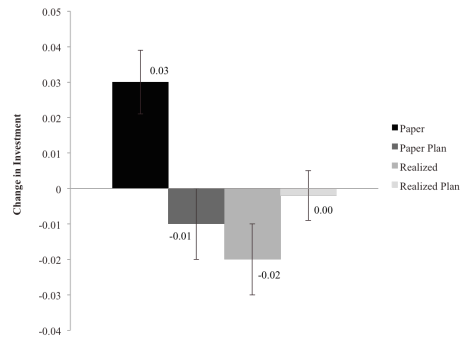

(b) *Strategy and Deviation*

FIGURE 4. Investment Change ($) after a Realized versus Paper Loss

Strategy and Deviation

A separate experiment explored the fourth prediction on whether realized versus paper losses have a differential affect on deviations from ex-ante investment strategies. Individuals (N=150) were recruited from Mechanical Turk in the same manner as in the second study, for the same baseline incentives. The protocols for choosing an investment and the resolution of uncertainty were also the same. Participants were assigned one "success" number between 1 and 6 in each round. As above, the lottery was analogous to a casino-like bet such that an investment was multiplied by 5 if the outcome of a virtual die roll matched the success number; if it did not match, the investment was lost. New success numbers were assigned in each round.

Participants were randomly matched into either the Paper treatment or the Realized treatment which were analogous to the treatments in the second experiment of Study 2. After reading the instructions but before making their investment decisions, each was asked to enter their fully contingent investment strategy over the four rounds. Participants began by entering how much they planned to invest in the first round, then how much they planned to invest in the second round conditional on winning in the first round as well as their planned investment conditional on losing in the first round, and so on. In the fourth round participants entered their planned investment for 8 possible contingencies. A picture of the outcome tree with nodes illustrating all possible contingencies was provided. Each participant then continued to make the actual investment decisions.

When making their actual investment decisions, those who accumulated losses at the end of the 3^{rd} round in the Paper treatment are predicted to have a significant, positive change in investment between the 3^{rd} and 4^{th} rounds, while those in the Realized treatment should have a negative change in investment. Comparing behavior to reported strategies allows for a direct test of the framework outlined in Section II against alternate explanations for the effects of realization such as learning (Erev and Roth, 1998). Individuals report their risk-taking strategy prior to experiencing any realization; they make actual choices with the same level of experience regarding realization in the Paper treatment but not in the Realized treatment. Therefore, if realization affects behavior by triggering learning, then planned and actual behavior should differ in the Realized treatment but not in the Paper treatment. In contrast, if realization operates through the mechanism outlined in Section II, then greater deviations from planned risk-taking after a loss should be observed in the Paper treatment than in the Realized treatment.

Results Examining actual investment decisions between treatments, changes in investment for those who had accumulated losses by the end of the 3^{rd} round were significantly different between the Paper and Realized treatments ($t(90) = 3.12$, $p < .01$). Participants in the Paper treatment *increased* their investment by \$0.03, taking on more risk than the null ($t(44) = 2.30$, $p = .03$). In contrast, participants who had similarly lost in the Realized treatment *decreased* their investment by \$0.02 ($t(46) = 2.11$, $p = .04$).

Figure 4(b) shows that the increase in risk-taking following a loss in the Paper treatment was a deviation from the participants' risk-taking plans. Comparing actual behavior to the ex-ante strategy after a loss within the *same individual* reveals a significant inconsistency between planned and actual investment after a loss: on average, participants in the Paper treatment invested \$0.04 *more* than they had planned to after a loss at the end of the 3^{rd} round ($t(89) = 2.51$, $p = .01$). Participants in the Realized treatment invested \$0.01 less than their strategy, but this difference was not significant ($t(93) = 1.34$, $p = .18$). Regressing a treatment dummy on the difference between the planned and actual change in investment after a loss provides support for the fourth prediction: following a loss, individuals in the Paper treatment exhibited a significantly larger positive deviation from their ex-ante strategy than those who had similarly lost in the Realized treatment ($\beta = .06, p < .01$).

IV. Discussion and Conclusion

The findings of this paper offer a unifying principle – realization – that reconciles two rich strands of literature on the dynamics of risk attitudes (Barberis, Huang and Santos, 2001; Benartzi and Thaler, 1995; Gneezy and Potters, 1997). Three studies are presented that demonstrate the differential effect of paper versus realized losses on risk-taking. Relative to the period before, individuals take on more risk after a paper loss and less if the loss is realized. Moreover, I show that the increase in risk-taking following a paper loss is a product of dynamic inconsistency in preferences – individuals deviate from their planned risk-taking strategies to take on more risk after a paper loss, and that realization mitigates these deviations.

These results have implications for recent work on CPT in dynamic contexts (Barberis, 2012; Ebert and Strack, 2014). For example, Ebert and Strack (2014) argue that the dynamic inconsistency implied by CPT leads individuals to gamble until "until the bitter end," i.e. until they are bankrupt, which speaks against CPT as a descriptive theory of dynamic choice. However, the differential effect of realized versus paper losses on deviations from individuals' ex-ante plans suggests that dynamic inconsistency should be observed less frequently than would be otherwise implied if realization was not considered. The results presented in this paper provide evidence for the predictions of CPT in dynamic settings, lending support for the theory as a descriptive model of choice when realization and its effects on bracketing are taken into account.

The interplay between realization and dynamic inconsistency in choice under uncertainty has significant implications for the role of monitoring in investor behavior. The results of the third study demonstrate that individuals whose investments were unsuccessful were reluctant to realize their losses, preferring to instead take on more risk before their positions were finally realized, and that this increased in risk-taking was a deviation from their ex-ante risk-taking plans. Such trading behavior can spiral out of control and lead to significant losses, which is consistent with the literature documenting overly aggressive trading (Barber et al., 2006) and a pronounced disposition effect displayed by individual investors (Odean, 1998). These effects can be particularly detrimental in contexts where greater risk-taking after a loss leads to lower expected returns than available alternatives (e.g. as a result of momentum). Since realization brings actual behavior after a loss closer to planned behavior, individuals sophisticated about their dynamic inconsistency should display a demand for realization after a loss as a commitment device against detrimental loss chasing. For example, an individual can automatically set his asset positions to be reported to a third party who can exogenously influence the realization of his positions. The results are also related to the prescriptions proposed by Weber and Zuchel (2005) regarding the benefits of binding precommitments in investment strategies.

For institutional traders, given the relationship between compensation and trading performance, the reporting of asset positions to the overseeing risk manager can be taken as a natural point of realization – analogous to the closing of the respective mental account. Anecdotal evidence suggests that some of the largest losses suffered by financial institutions occurred as a result of traders hiding prior losses while taking on excessive risk in an attempt to cover them.[14] A firm's monitoring strategy could utilize realization of traders' asset positions while lowering the incentives for them to hide loses. For example, Camerer and Loewenstein (2004) describe an investment banker whose firm forced traders to periodically

[14]See Jerome Kerviel's 4.9 billion Euro loss for Societe Generale, Kweku Adoboli's 2.3 billion dollar loss for UBS and Nick Leeson's 1.3 billion dollar loss for Barrings, which wiped out the firm.

switch positions (the portfolio of assets that the trader bought and is blamed or credited for) with the position of another trader. Such a policy would be an effective tool to curb loss chasing, particularly if position switches were performed at times not announced to the traders ex-ante.

The differential effect of realized versus paper losses on risk-taking also contributes to the emerging literature on how non-standard factors such as emotions (Caplin and Leahy, 2001, 2004; Caplin, 2003; Koszegi, 2006) and other psychological factors (Karlsson, Loewenstein and Seppi, 2009; Loewenstein, 1996; Loewenstein et al., 2001) affect risk attitudes. This paper tests whether realization as a transfer is a sufficient condition for predicting when a loss will lead to less risk-taking rather than more. However, other factors involved with or following a loss, e.g. very long passages of time, may lead to updating of the reference point and closure of the associated mental account such that the individual does not respond by taking on more risk. Future research should examine other factors that affect when mental accounts are opened or closed in order to determine the necessary and sufficient conditions for the theory.

REFERENCES

Abdellaoui, Mohammed. 2000. "Parameter-Free Elicitation of Utility and Probability Weighting Functions." *Management Science*, 46(11): 1497–1512.

Andersen, Steffen, Tobin Hanspal, and Kasper Meisner Nielsen. 2015. "Once Bitten, Twice Shy: Do Personal Experiences or Wealth Changes Affect Risk Taking?" *mimeo.*

Andrade, Eduardo B., and Ganesh Iyer. 2009. "Planned Versus Actual Betting in Sequential Gambles." *Journal of Marketing Research*, 46(3): 372–383.

Barber, Brad M., Yi-Tsung Lee, Yu-Jane Liu, and Terrance Odean. 2006. "Just How Much Do Individual Investors Lose by Trading?" *Review of Financial Studies*, 22(2): 609–632.

Barberis, Nicholas. 2012. "A Model of Casino Gambling." *Management Science*, 58(1): 35–51.

Barberis, Nicholas, and Wei Xiong. 2012. "Realization Utility." *Journal of Financial Economics*, 104(2): 251–271.

Barberis, Nicholas, Ming Huang, and Tano Santos. 2001. "Prospect Theory and Asset Prices." *The Quarterly Journal of Economics*, 116(1): 1–53.

Barkan, Rachel, and Jerome R. Busemeyer. 1999. "Changing Plans Dynamic Inconsistency and the Effect of Experience on the Reference Point." *Psychonomic Bulletin & Review*, 6(4): 547–554.

Benartzi, Shlomo, and Richard H. Thaler. 1995. "Myopic Loss Aversion and the Equity Premium Puzzle." *Quarterly Journal of Economics*, 110(1): 73–92.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2012. "Salience Theory of Choice Under Risk." *The Quarterly Journal of Economics*, 127(3): 1243–1285.

Camerer, Colin F., and George Loewenstein. 2004. "Behavioral Economics: Past, Present, Future." In *Advances in Behavioral Economics.* , ed. Colin F. Camerer, George Loewenstein and Matthew Rabin, Chapter 1, 3–51. New York:Princeton University Press.

Caplin, Andrew. 2003. "Fear as a Policy Instrument." In *Time and Decision: Economic and Psychological Perspectives on Intertemporal Choice.* , ed. George F. Loewenstein, Daniel Read and Roy F. Baumeister, 441–458. New York, NY:Russell Sage Foundation.

Caplin, Andrew, and John Leahy. 2001. "Psychological Expected Utility Theory and Anticipatory Feelings." *Quarterly Journal of Economics*, 116(1): 55–79.

Caplin, Andrew, and John Leahy. 2004. "The Supply of Information by a Concerned Expert." *The Economic Journal*, 114(497): 487–505.

Charness, Gary, and Uri Gneezy. 2010. "Portfolio Choice and Risk Attitudes: An Experiment." *Economic Inquiry*, 48(1): 133–146.

Charness, Gary, and Uri Gneezy. 2012. "Strong Evidence for Gender Differences in Risk Taking." *Journal of Economic Behavior & Organization*, 83(1): 50–58.

Choi, James J., David Laibson, and Brigitte C. Madrian. 2009. "Mental accounting in portfolio choice: Evidence from a flypaper effect." *American Economic Review*, 99(5): 2085–2095.

Coval, Joshua D, and Tyler Shumway. 2005. "Do Behavioral Biases Affect Prices?" *Journal of Finance*, 60(1): 1–34.

Della Vigna, Stefano, and Ulrike Malmendier. 2006. "Paying Not to Go to the Gym." *American Economic Review*, 96(3): 694–719.

Dillenberger, David, and Kareen Rozen. 2014. "History-Dependent Risk Attitude." *Journal of Economic Theory*, 157: 445–477.

Ebert, Sebastian, and Philipp Strack. 2014. "Until the Bitter End: On Prospect Theory in a Dynamic Context." *American Economic Review*, forthcoming.

Erev, Ido, and Alvin E. Roth. 1998. “Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria.” *American Economic Review*, 88(4): 848–881.

Frydman, Cary, Nicholas Barberis, Colin Camerer, Peter Bossaerts, and Antonio Rangel. 2014. “Using Neural Data to Test a Theory of Investor Behavior: An Application to Realization Utility.” *Journal of Finance, forthcoming*, 69: 907–946.

Gneezy, Uri, and Jan Potters. 1997. “An Experiment on Risk Taking and Evaluation Periods.” *The Quarterly Journal of Economics*, 112(2): 631–645.

Gonzalez, Richard, and George Wu. 1999. “On the Shape of the Probability Weighting Function.” *Management Science*, 166: 129–166.

Haigh, Michael S., and John A. List. 2005. “Do Professional Traders Exhibit Myopic Loss Aversion? An Experimental Analysis.” *The Journal of Finance*, 60(1): 523–534.

Heath, Chip. 1995. “Escalation and De-escalation of Commitment in Response to Sunk Costs: The Role of Budgeting in Mental Accounting.” *Organizational Behavior and Human Decision Processes*, 62(1): 38–54.

Ingersoll, Jonathan E., and Lawrence J. Jin. 2013. “Realization Utility with Reference-Dependent Preferences.” *Review of Financial Studies*, 26(3): 723–767.

Kahneman, Daniel, and Amos Tversky. 1979. “Prospect Theory: An Analysis of Decision under Risk.” *Econometrica*, 47(2): 263–292.

Karlsson, Niklas, George Loewenstein, and Duane Seppi. 2009. “The ostrich effect: Selective attention to information.” *Journal of Risk and Uncertainty*, 38(2): 95–115.

Kaustia, Markku, and Samuli Knupfer. 2008. “Do Investors Overweight Personal Experience? Evidence from IPO Subscriptions.” *The Journal of Finance*, 63(6): 2679–2702.

Koszegi, Botond. 2006. “Emotional Agency.” *Quarterly Journal of Economics*, 121(1): 121–155.

Koszegi, Botond, and Matthew Rabin. 2009. “Reference-Dependent Consumption Plans.” *American Economic Review*, 99(3): 909–936.

Langer, Thomas, and Martin Weber. 2008. “Does Commitment or Feedback Influence Myopic Loss Aversion?” *Journal of Economic Behavior & Organization*, 67(3-4): 810–819.

Linville, P W, and G W Fischer. 1991. "Preferences for separating or combining events." *Journal of Personality and Social Psychology*, 60(1): 5–23.

Liu, Yu-Jane, Chih-Ling Tsai, Ming-Chun Wang, and Ning Zhu. 2010. "Prior Consequences and Subsequent Risk Taking: New Field Evidence from the Taiwan Futures Exchange." *Management Science*, 56(4): 606–620.

Loewenstein, George. 1996. "Out of Control : Visceral Influences on Behavior." *Organizational Behavior and Human Decision Processes*, 65(3): 272–292.

Loewenstein, George F., Elke U. Weber, Christopher K. Hsee, and Ned Welch. 2001. "Risk as Feelings." *Psychological Bulletin*, 127(2): 267–286.

Loewenstein, George, Ted O'Donoghue, and Matthew Rabin. 2003. "Projection Bias in Predicting Future Utility." *The Quarterly Journal of Economics*, 118(4): 1209–1248.

Odean, Terrance. 1998. "Are Investors Reluctant to Realize Their Losses?" *The Journal of Finance*, LIII(5): 1775–1798.

O'Donoghue, Ted, and Matthew Rabin. 1999. "Doing It Now or Later Doing It Now or Later." *American Economic Review*, 89(1): 103–124.

Pagel, Michaela. 2012. "Expectations-Based Reference-Dependent Life-Cycle Consumption." *mimeo.*

Prelec, Drazen. 1998. "The Probability Weighting Function." *Econometrica*, 66(3): 497–527.

Rabin, Matthew. 2000. "Risk Aversion and Expected-Utility Theory: A Calibration Theorem." *Econometrica*, 68(5): 1281–1292.

Rabin, Matthew, and Georg Weizsäcker. 2009. "Narrow Bracketing and Dominated Choices." *American Economic Review*, 99(4): 1508–1543.

Read, Daniel, George Loewenstein, and Matthew Rabin. 1999. "Choice Bracketing." *Journal of Risk and Uncertainty*, 19: 171–197.

Savage, Leonard J. 1954. *The Foundations of Statistics.* New York:Wiley.

Shefrin, Hersh, and Meir Statman. 1985. "The Disposition to Sell Winners Too Early and Ride Losers Too Long: Theory and Evidence." *Journal of Finance*, 40(3): 777–790.

Shiv, Baba, George Loewenstein, and Antoine Bechara. 2006. "The Dark Side of Emotion in Decision-Making: When Individuals with Decreased Emotional Reactions Make More Advantageous Decisions." *Cognitive Brain Research*, 23(1): 85–92.

Shiv, Baba, George Loewenstein, Antoine Bechara, Hanna Damasio, and Antonio R Damasio. 2005. "Investment behavior and the negative side of emotion." *Psychological science*, 16(6): 435–9.

Tanaka, Tomomi, Colin F. Camerer, and Quang Nguyen. 2010. "Risk and Time Preferences: Linking Experimental and Household Survey Data from Vietnam." *American Economic Review*, 100(1): 557–571.

Thaler, Richard. 1985. "Mental Accounting and Consumer Choice." *Marketing science*, 4(3): 199–214.

Thaler, Richard H, and Eric Johnson. 1990. "Gambling with the House Money and Trying to Break Even: The Effects of Prior Outcomes on Risky Choice." *Management science*, 36(6): 643–660.

Tversky, Amos, and Daniel Kahneman. 1992. "Advances in prospect theory: Cumulative representation of uncertainty." *Journal of Risk and Uncertainty*, 5(4): 297–323.

Weber, Martin, and Colin F. Camerer. 1998. "The Disposition Effect in Securities Trading: An Experimental Analysis." *Journal of Economic Behavior & Organization*, 33: 167–184.

Weber, Martin, and Heiko Zuchel. 2005. "How Do Prior Outcomes Affect Risky Choice? Further Evidence on the House-Money Effect and Escalation of Commitment." *Decision Analysis*, 2: 30–43.