Dynamic Inconsistency in Risky Choice: Evidence from the Lab and Field

RAWLEY HEIMER, ZWETELINA ILIEWA, ALEX IMAS, and MARTIN WEBER*

ABSTRACT

Many economically important settings, from financial markets to consumer choice, involve sequential decisions under risk. Data from these dynamic settings runs counter from findings in one-shot settings: people are anomalously risk-averse in the latter, while taking on even negative expected-value risk in the former. We use two pre-registered experiments ($N = 940$) and a unique brokerage dataset of traders' investment plans and subsequent decisions ($N = 190,000$) to shed light on this discrepancy. Participants in our experiments made a sequence of choices over whether or not to accept a symmetric fair gamble with feedback provided after every decision. We use both between-subject and within-subject design to compare participants' ex-ante strategies conditional on future outcomes with their actual behavior. A large majority of participants plan to follow 'loss-exit' strategies - to continue taking risk after gains and to stop after losses. Actual behavior exhibited the *reverse* pattern: participants cut their gains early and chased their losses. We find an analogous dynamic inconsistency in the investment plans and subsequent decisions of traders in our unique brokerage dataset. We formally demonstrate that this behavioral pattern identifies the dynamic predictions of Cumulative Prospect Theory. A significant demand for commitment devices points to at least partial sophistication about the dynamic inconsistency. We use simulations to quantify that the welfare costs of naïveté for a representative agent are over *one hundred and ten percent* of the stakes in a one-round investment. Moreover, participants' wide-spread demand for non-binding commitment, which is ineffective in mitigating dynamic inconsistency, highlights a second form of naïveté with respect to the effectiveness of such 'soft' commitment. Our results have implications for evaluating unintended effects of recently introduced regulation that mandates soft commitment.

---

* Rawley Heimer is at the Carroll School of Management, Boston College (e-mail: heimer@bc.edu). Zwetelina Iliewa is at the Max Planck Institute for Research on Collective Goods (e-mail: iliewa@coll.mpg.de). Alex Imas is at the Carnegie Mellon University (e-mail: aimas@andrew.cmu.edu). Martin Weber is at the Department of Finance, University of Mannheim, and the Centre for European Policy Research (e-mail: weber@bank.bwl.uni-mannheim.de). We are grateful to Peter Bossaerts, Markus Dertwinkel-Kalt, Benjamin Enke, Johannes Haushofer, Camelia Kuhnen, Matthias Sutter, and Florian Zimmermann for helpful suggestions. We also thank participants at the Workshop on Behavioral and Experimental Economics and the Annual Meeting of the Society for Experimental Finance. Pascal Kieren and Yanghua Shi provided excellent research assistance. All remaining errors are our own.

I. Introduction

People making risky decisions often face sequences of independent choices. Stock market investors, after observing the performance of their stocks, can increase their position, sell their shares, or continue to hold the same assets. Consumers face similar dynamic choices: for example, a family that is used to eating at home can try a new restaurant knowing they can return if they like it or go back to eating in.

Behavior in such dynamic environments—where people decide to take on risk knowing they can stop anytime—appears to run counter to findings from one-shot settings. People tend to avoid uncertainty when it is presented in isolation, overwhelmingly rejecting positive-expected-value and displaying a seemingly anomalous preference for safer options (Kahneman and Tversky (1979)). At the same time, people appear risk-seeking when risk is offered as part of a sequence, even if the cumulative gamble has a *negative expected value*. For example, while casino and other forms of commercial gambling leave the average player with less money than when she started, the $240 billion dollar industry is booming, collecting record revenues year after year.[1] Similarly, investments in the retail foreign exchange market (FOREX) yield negative returns for the typical trader, yet the daily turnover is roughly equivalent to the entire NYSE family of stock exchanges (King and Rime (2010)). Studies in the lab have confirmed these patterns, with participants expressing a willingness to take on fair (Andrade and Iyer (2009)) and negative expected-value gambles (Imas (2016)) when they are presented as part of a sequence.[2] Understanding what drives the increase in risk-taking in dynamic versus one-shot deci-

[1] American Gambling Association annual report, 2018.

[2] Note that this discrepancy between one-shot and sequential risk-taking is fundamentally distinct from risk-taking as a function of the evaluation period. Prior work has shown that people take on more risk when feedback on outcomes is provided less frequently. This phenomenon, termed myopic loss aversion (MLA; Gneezy and Potters (1997); Benartzi and Thaler (1995)), cannot explain the outlined differences in risk-taking in sequential versus one-shot environments because 1) feedback is provided after every choice and 2) an MLA agent would reject fair or negative expected-value risk regardless of feedback frequency.

sions can potentially shed light on other phenomena associated with sequential choice, such as the costly active trading of individual stocks (Barber and Odean (2000)) and the tendency to 'double-down' on failed strategies (Heath (1995)).

This paper explores how people choose to take risks in dynamic environments versus one-shot settings, and how they re-evaluate their ex-ante strategies after experiencing gains and losses. We begin the investigation by conducting two pre-registered experimental studies ($N = 940$). In the experiments, participants are offered the choice to accept or reject a sequence of fair symmetric gambles while being provided feedback after every decision and having the choice to stop anytime. We use a mixture of between-subject and within-subject design to (i) test whether participants are more willing to take risk if offered a sequence of fair gambles than a gamble in isolation and (ii) understand the mechanism behind the discrepancy in risk-taking by comparing their ex-ante strategies before accepting risk to their actual behavior once outcomes are realized. The latter allows us to identify dynamic inconsistency in sequential risk-taking and quantify sophistication through demand for commitment. Participants are assigned to different treatments that vary the number of gambles they are presented with and whether we elicit their strategies prior to the initial choice. In some treatments, these strategies binding; in others, participants are reminded of their strategies as they make their sequential decisions but are free to deviate. We refer to the former treatments as offering 'hard commitment' and the latter as 'soft commitment.' Finally, in a separate set of treatments participants make decisions without stating their ex-ante strategies.

Our experimental findings can be summarized as follows. First, we show that participants are significantly more likely to accept risk when it is part of a larger sequence of gambles than in isolation. Comparing ex-ante strategies to behavior allows us to identify the mechanism for this discrepancy. The vast majority – more than 80% – of strategies can be classified as "loss-exit" plans; in contrast, only 7% of strategies can be classified as "gain-exit." Actual choices follow the *reverse* pattern: participants are significantly more likely to stop after winning than after losing. Differences of the entry rates be-

tween treatments suggest that a portion of our participants are sophisticated about this dynamic inconsistency. People are significantly more likely to begin taking risk when provided with a commitment opportunity, though a substantial proportion appears to be naïvé. In contrast, there is no evidence for sophistication when comparing soft and hard commitment opportunities. Despite most people deviating from their strategies when loss or gain limits are reached, they are equally likely to take on risk when commitment is non-binding as when it is binding. As explained below and demonstrated formally in the Appendix, this pattern of findings are most consistent with the dynamic predictions of Cumulative Prospect Theory (CPT thereafter, Tversky and Kahneman (1992)). Models that incorporate diminishing sensitivity without a gain-versus-loss asymmetry (Expected Utility Theory) or probability weighting (Rank Dependent Utility; Quiggin (1982); Yaari (1987)) cannot rationalize the observed behavioral patterns. Having provided reduced-form evidence for the CPT framework, we use simulations to quantify the welfare loss resulting from dynamic inconsistency. Naïvé agents who accept the sequential gamble with the illusion that they will stick to their loss-exit plan incur a utility cost from participation. For a representative CPT agent in our setting, the predicted cost of naïveté is equivalent to losing more than *one hundred and ten percent* of the one-round stake with certainty.

Our experimental paradigm can be used to formally test recent theoretical work that rationalizes people's willingness to take fair or negative expected-value risk in dynamic environments (Barberis (2012); Ebert and Strack (2015)). Barberis (2012) shows that dynamic CPT predicts that same person will reject a single fair gamble while accepting the same gamble as part of a sequence of choices. Identification of this framework comes from comparing people's ex-ante risk-taking strategies to their ex-post decisions. Specifically, non-linear probability weighting generates a dynamic inconsistency between ex-ante strategies and behavior: people begin to take on risk with the plan of stopping after losses and continuing on after gains. Actual decisions are predicted to follow the reverse pattern,

with a greater willingness to take on risk after losses than after gains.[3] To illustrate the intuition for this prediction, consider a decision-maker facing a sequence of fair gambles. She makes a choice to accept or reject the first gamble as part of a strategy. A strategy that does not depend on prior outcomes generates the same level of skew over final wealth as the individual gamble; on the other hand, a "loss-exit" plan—exiting earlier after losses than after gains—generates a positive-skewed lottery over final wealth. For example, a "loss-exit" strategy over two 50:50 gambles, each with an upside of G and downside of $-G$ generates the the following lottery over final wealth, $(2G, 1/4; 0, 1/4; -G, 1/2)$. This lottery has substantially more positive skew than the single gamble, or a non-contingent strategy of accepting both gambles regardless of outcome, $(2G, 1/4; 0, 1/2; -2G, 1/4)$. CPT preferences lead the decision-maker to overweight smaller probabilities; in turn, the positive skew generated by the "loss-exit" strategy makes the sequence of gambles more attractive than the single gamble with no skew. Diminishing sensitivity of the value function—which generates risk-seeking in the loss domain and risk-aversion in the gain domain—lead the decision-maker to systematically deviate from this plan after accepting the first gamble: she rejects the second gamble after experiencing a gain and accepts it after experiencing a loss. The framework also makes a prediction for the role of commitment in this setting. A decision-maker sophisticated about her dynamic inconsistency will be more likely to accept risk if she can commit to the "loss-exit" strategy than if no such commitment opportunities exist.

Having demonstrated that people systematically deviate from their risk-taking strategies in a controlled lab environment, we provide field evidence for this dynamic inconsistency in financial markets. To do so, we employ a unique dataset from a large online brokerage with round 190,000 traders from over 150 countries. This dataset records traders' investment plan by means of mandatory stop-loss and take-profit orders at the time of the opening of each position, as well as any subsequent revision of these orders

[3] The theoretical predictions regarding actual behavior after gains depend on whether the sequence is assumed to be finite (Barberis (2012)) or infinite (Ebert and Strack (2015)). A higher tendency to cut losses early is predicted for settings with finite number of rounds, such as our experimental paradigm.

for the entire duration of the position. The combination of initial limit orders and subsequent revisions allows us to directly characterize the intended risk-taking strategies of traders, as well as how strategies evolve as positions accumulate paper gains and losses. In this setting, the initial limit orders can be viewed as traders' ex-ante strategies. For example, by placing a stop-loss order of 10% and a take-profit order of 20%, the trader reveals a willingness to lose 10% for the chance of gaining 20%. Data on revisions allows us to observe how traders' risk-taking strategies evolve as they experience paper gains and losses on open positions.

Our findings from the field can be summarized as follows. Similar to the experimental setting, we find that the majority of limit orders can be classified as "loss-exit" plans. That is, the average newly opened position will include a stop-loss order that is smaller in absolute percentage value than the position's take-profit order. However, after the active position accumulates paper gains or losses, the trader tends to reverse the "loss-exit" plan. Traders experiencing paper losses are most likely to further lower their stop-loss orders. For example, when the position is currently at a 5% loss, a trader who sets a 10% stop-loss will modify the stop-loss order by lowering it to 20%. On the other hand, when traders experience a paper gain, they are most likely to manually close the trade before the take-profit order is triggered. For example, when the trader has set a 20% take-profit order and the position hits a 10% paper gain, she is most likely to manually close the trade at 10%.

Together, the experimental findings complement results from the field. The analogous pattern of results in both settings increase confidence in their generalizability and portability. The observed behavioral patterns highlight that the option to stop taking on risk in sequential choice after receiving feedback is a crucially appealing feature of these environments—individuals begin to take on risk that they would avoid in isolation *because* they can condition future choices on past outcomes. However, dynamic inconsistency in ex-post behavior can potentially lead to welfare losses due to individuals cutting their gains early and chasing losses once outcomes are actually realized. Importantly,

the costs of naïveté about this dynamic inconsistency also depend on the availability of soft commitment devices, which are the most wide-spread type of commitment devices in the real world.[4] Our experimental results show that the presence of soft commitment opportunities led a substantial fraction of individuals who would have avoided risk absent commitment opportunities to accept it. This 'illusion of commitment' is costly, as these same individuals end up systematically deviating from their non-binding strategy. Providing 'hard commitment' opportunities that do not allow ex-post revision would not only increase naïvé agents' welfare, but also the welfare of sophisticated agents who correctly anticipate their dynamic inconsistency and do not accept sequential risks as a result. Hard commitment is valuable for these sophisticated agents because it enables them to accept risk as part of a utility-maximizing "loss-exit" strategy. As we show in Section 5, the opportunity for hard commitment is equivalent in terms of welfare to a certain *gain* of up to one hundred and twenty percent of the one-round stake for some sophisticated agents who exhibit stronger non-linear probability weighting than the representative agent.

Our results shed light on prior findings in sequential risky choice and suggest scope for policy in these settings. For example, the propensity to realize gains and hold on to losses (Odean (1998)), and the underperformance of investors due over-trading in response to prior outcomes Barber and Odean (2000) may run counter to their ex-ante strategies when they begain to trade. Heimer and Imas (2019) document the outcomes of a policy that restricted the amount of leverage available to retail FOREX traders. Counter to the predictions of the standard model, but consistent with dynamic inconsistency, decreasing traders' choice sets in this manner actually *improved* their outcomes. As discussed further in Section VII, our results on soft commitment are particularly relevant for recent policy proposals aiming to stem traders' losses through non-binding interventions (e.g. the Markets in Financial Instruments Directive). Behavior under soft commitment suggests

[4] For example, both loss and gain limits in financial markets can be adjusted at will. Regulation aimed at limiting the scope for unintended losses, such as the regulation on "depreciation reporting", which is a part of the recently revised financial instruments regulation in European markets (MiFID II), essentially urges investors to think about a loss-exit strategy while leaving the loss-limit non-binding.

that this type of policy may backfire by encouraging investors to take on more risk than they otherwise would without effectively preventing them from chasing losses.

Our findings contribute to the literature on dynamic inconsistency between planned and actual behavior. A large literature explores systematic deviations from ex-ante strategies in intertemporal choice (Frederick, Loewenstein, and O'Donoghue (2002)). Sophistication about dynamic inconsistency and demand for commitment have been studied both theoretically (O'Donoghue and Rabin (1999)) and empirically (DellaVigna and Malmendier (2006)). The proposed mechanism for time inconsistency – hyperbolic discounting – is conceptually distinct from the driver of dynamic inconsistency in risky choice and cannot explain the behavioral patterns described here. In contrast to the significant body of work on the former, few papers have explored the latter; and of those that do, the majority are theoretical rather than empirical.

The rest of the paper proceeds as follows. Section II describes the experimental design. Section III outlines the theoretical predictions of dynamic CPT in our setting. Section IV presents experimental results and discusses them. Section V describes the implications of our findings for welfare. Section VI describes the field setting and presents complimentary behavioral findings. Section VII concludes. Finally, Appendix A formally compares the predictions of various models of dynamic choice under uncertainty.

II. Experimental Design

We conduct two preregistered experimental studies with a total of 940 subjects to examine the dynamics of decision-making under risk.[5] The studies are design to (i) test for a dynamic inconsistency in risky choice that arises when subjects have the opportunity to invest in the gamble for many rounds and stop endogenously at any point of time,

[5] Both experiments were preregistered: https://aspredicted.org/blind.php?x=x954rp and https://aspredicted.org/blind.php?x=tn4dt4.

and (ii) estimate subjects' sophistication regarding the dynamic inconsistency and the effectiveness of different commitment devices.

Subjects are presented with a fair symmetric gamble with an expected value of zero. In particular, if the subject accepts the gamble, depending on the outcome of a random draw her investment may either double or be lost, each with a probability of 50%. To ensure that subjects understand the gamble, they first draw ten samples from a stratified sample (for the benefits of sampling for the understanding of probabilities see, e.g., Kaufmann, Weber, and Haisley (2013) and Hogarth and Soyer (2015)).

The main task is to decide whether or not to invest 10 cents in the lottery. There are 3 (4) between-subject treatments in Experiment 1 (Experiment 2), as illustrated in Figure 1. In each experiment there is a one-shot treatment and several multi-round treatments. Experiment 1 introduces two multi-round treatments: a sequential treatment, in which subjects receive feedback after every decision and subsequently make their next choice, and a hard commitment treatment, in which subjects commit to an ex-ante strategy. Experiment 1 thus allows us to test for the dynamic inconsistency and subjects' sophistication about it. Experiment 2 replicates Experiment 1 and includes a soft plan treatment which tests for the effectiveness of a realistic, non-binding commitment device and subjects' sophistication about it.

Subjects in the one-shot treatment receive an endowment of 10 cents and decide only once whether or not to invest it in the fair gamble (binary choice). In the multi-round treatments (i.e. sequential, hard commitment and soft commitment), the main task consists of a sequence of binary risky choices over a maximum of 26 rounds. We select the maximum number of rounds to be the minimum number of rounds for which a representative prospect theory agent would accept the multi-round gamble as shown by Barberis (2012).[6] The subject receives an initial endowment of $2.60 at the beginning of

[6] Previous experimental studies by Andrade and Iyer (2009) and Barkan and Busemeyer (1999) analyze dynamics of risk-taking over only two rounds.

the experiment and decides whether to invest 10 cents in each round or to keep it (binary choice). Once the subject stops investing, she cannot re-enter and the main part of the experiment is over. In the sequential treatment, subjects make decisions round by round receiving feedback in between. At the beginning, subjects decide whether they want to invest in the first round. If the subject decides to enter the lottery, the outcome for the first round is revealed and she decides whether or not to continue investing in the second round, and so forth. Participants are informed about the total gains or losses they have accumulated since the beginning alongside the outcome of the last round.

In the commitment treatments, we ask subjects for an exit strategy before they decide whether or not to start investing in the first round. The stopping rule is elicited by asking subjects to indicate their gain limit (i.e. minimum gain they will be happy to walk away with rather than continuing to gamble) and loss limit (maximum loss they are willing to accept). The decision to enter is elicited subsequently. In the hard commitment treatment (Experiment 1 and 2), subjects are informed that they will automatically stop investing if either one of their limits is triggered, hence their limits serve as a binding commitment device. In the soft commitment treatment (in Experiment 2), subjects are informed that they will be notified immediately as soon as either one of their limits is triggered and they will then have the chance to decide how to proceed; hence, their limits are a non-binding commitment device.[7]

We choose a fair gamble to distinguish sequential risk taking with endogenous stopping from risk taking across multiple rounds with differential timing of feedback, such as in the studies examining myopic loss aversion. For example, Gneezy and Potters (1997) show that subjects are more often willing to accept a positive expected-value gamble across multiple rounds with cumulative feedback presented at the very end than plays of the gamble over the same number of rounds, but with feedback provided after each

[7] Fischbacher, Hoffmann, and Schudy (2017) use gain and loss limits as a treatment against disposition effect. In contrast to the previous study, the limits in our experiments are not optional to prevent any selection bias and cannot be revised in every round to ensure that they accurately measure subjects' ex ante plan.

outcome is realized. Since the expected value of each gamble is positive, presenting cumulative feedback leads to a greater likelihood that the participant sees a positive outcome than if feedback is presented after every round. The authors explain this result through loss aversion and narrow bracketing: a loss-averse individual would prefer to avoid seeing negative information, so the gambles with cumulative feedback are more attractive. This mechanism cannot explain a greater willingness to accept a *fair* gamble as part of a sequence than in isolation. First, in our setting, feedback is provided after every round. Additionally, even if only cumulative feedback was provided, investing for a fixed number of rounds with endogenous exit generates greater variance in outcomes without any risk premium compared to a single play of the gamble.

Note that, by design, the only way for subjects to achieve a skewed outcome distribution is through an outcome-dependent dynamic strategy. This is because gains and losses in the individual rounds are linearly compounded, which is why investing longer would not result in a positively skewed outcome distribution.[8]

[Insert Figure 1 about here]

Both experiments were conducted online on MTurk via TurkPrime. Experiment 1 was conducted between February 18th and 21st, 2019, and Experiment 2 took place between June 28th and July 3th, 2019. Both experiments were conducted in four batches to isolate potential day-of-week effects. Participants in Experiment 1 were excluded from the subject pool for Experiment 2. Both experiments consist of an entry-level questionnaire, that elicits demographic characteristics (age, gender, highest level of education and study field) as well as self-reported level of statistical skills, a main task, and an exit-level questionnaire that elicits further control variables for the purpose of distinguishing between different hypotheses. The participation rate in the main task in all treatments of

[8] However, outcome-independent strategies would affect the variance and the kurtosis of the outcome distribution. Investing longer unconditionally is connected to larger variance and smaller kurtosis.

both Experiments 1 and 2 is above 95%, hence there is little room for selective attrition. Table I shows that the average participant is between 34 and 35 years old male with a college or bachelor degree in computer sciences or business. Except for the participants in Experiment 2 reporting higher perception of their own statistical skills, there are no significant differences in the characteristics of the participants in Experiments 1 and 2.

[Insert Table I about here]

III. Theoretical Predictions

In this section, we provide the intuition for the dynamic predictions of Cumulative Prospect Theory (CPT) in our setting. Appendix A presents the formal derivation and compares them to the predictions of other models of dynamic choice under uncertainty, such as Rank Dependent Utility and Expected Utility Theory.

How does CPT rationalize risk-taking when there is no risk premium in the single gamble? When considering sequential choices, the decision-maker evaluates which strategy is most attractive in expectation. Strategies are evaluated as a function of the initial choice and the future choices she is planning to make in response to realized outcomes.[9] For example, a trader may follow a strategy where she initially invests in a stock, reinvests if its price increases by 10 percent, and closes the position if its price falls by 5 percent; she may decide differently—e.g., not invest at all—if the subsequent investment opportunities were not available. This type of strategy—stopping early after losses and continuing to take on risk after gains—is referred to as a "loss-exit" plan (Barberis, 2012). Compared to a "gain-exit" strategy or any outcome-independent strategy, a "loss-exit"

[9] Optimization in dynamic decision-making of this sort is typically referred to as *dynamic programming* (Dixit, Sherrerd et al. (1990)). Dynamic programming identifies the strategy which yields the highest value not just in terms of the expected outcomes of the initial choice, but rather the expected outcomes of future choices that are contingent on the outcomes of prior realizations.

plan generates greater positive skewness in the probabilities over final outcomes; "loss-exit" plans can even create positive skew when each individual gamble has no skew at all. For example, take a fair gamble with even odds that is offered twice. A "loss-exit" plan would yield a 1/4 probability of a large gain, 1/4 of breaking even, and a 1/2 probability of losing; notably, this strategy generates a small probability of a big win—an outcome that does not exist in the case of a one-shot gamble.

Even though the strategy does not change the expected value of the outcome distribution, which remains zero, it does transform the skewness of the outcome distribution to make it more attractive. There is substantial evidence that, holding other things constant, people prefer positively-skewed risk (Dertwinkel-Kalt and Köster, forthcoming; Eraker and Ready, 2015). This phenomenon has typically been explained through CPT because non-linear probability weighting leads to the overweighting of small likelihoods (Barberis and Huang, 2008). In turn, accepting the gamble as part of a "loss-exit" plan with the option—but no obligation—to take on a second gamble is more attractive to an individual with CPT preferences than accepting the gamble in isolation: the "loss-exit" plan generates a low likelihood of a large gain while limiting the potential downside.

Importantly, CPT also generates a dynamic inconsistency between a person's ex-ante strategy and actual behavior once she starts taking on risk. Standard models of risky decision-making assume that people stick to their strategies regardless of the realized outcomes: the trader who plans to sell after a loss and double down after a gain will stick to this strategy. In contrast, a decision maker with CPT preferences who accepts the first bet as part of a "loss-exit" strategy systematically deviates from this plan by stopping too early after winning and continuing on too late after losing (relative to her plan). To illustrate the intuition, consider what happens in the above example after seeing a gain. The plan that involves accepting the next gamble does not generate as much positive skewness as the one which included the previous gamble—the decision-maker is now looking at a prospect with even odds. Since there is no low-likelihood outcome to overweight, she will deviate from her plan and reject the second gamble. After losing,

the decision-maker faces the choice of accepting the loss with certainty or potentially compensating for it by taking on the second gamble. CPT preferences predict that she will chase her losses and accept more risk. As we formally demonstrate in Appendix A, both deviations are driven in part by the assumption of diminishing sensitivity in the CPT value function, such that the decision-maker is risk-averse in the gain domain and risk-seeking in the loss domain. There, we also show that this prediction is unique to CPT: such dynamic inconsistency is not predicted by Expected Utility Theory or other models of Rank Dependent Utility, which also features non-linear probability weighting.

The framework also makes predictions on how awareness of this dynamic inconsistency will affect behavior. Sophisticated individuals will only accept the first gamble if provided with an opportunity to commit to their strategy; naíve individuals will be equally likely to accept the first gamble in a sequence whether or not they have a commitment opportunity. Standard models in decision theory predicts no such demand for commitment: in a market setting, a trader who is dynamically consistent will take on the same amount of risk regardless of whether she can commit to this plan through, for example, a stop-loss or take-profit order, or not.

We now proceed to present our empirical findings and compare them to these dynamic predictions.

IV. Experimental Results and Discussion

A. Accepting a Fair Gamble (Entry Decision)

In the following, we examine subjects' initial decisions to accept a fair gamble, i.e. the *entry decision*. First, we aim to learn whether subjects are more likely to take fair risk if they have the option to take it more than once. This would indicate that subjects are either risk-loving or they perceive beneficial dynamic strategies that increase

the perceived value of the expected outcome distribution. Second, we analyze whether subjects value the availability of a commitment device, which would be indicative of a dynamic inconsistency and, at the same time, subjects' sophistication about it. We distinguish between binding and non-binding commitment devices to estimate the scope of subjects' sophistication.

First, our results indicate that subjects are more likely to accept the fair gamble if they can take risk over many rounds compared to taking it only once. This result follows from Figure 4 and Table II, Panel A. Figure 4 displays the proportion of subjects who accept the gamble in the first round in each between-subject treatment. It shows that the entry rate is at smallest in the one-shot treatment in both experiments.[10] Table II, Panel A, additionally displays the marginal fixed effects (mfx) of Probit regressions of the binary entry decision. The main independent variables are dummy variables for each of the multi-round treatments (i.e. sequential, hard commitment and soft commitment treatments). The mfx measures the difference between the probability to accept the gamble in the respective multi-round treatment compared to the one-shot treatment, which is the reference treatment in Panel A. Since the treatments are randomly assigned and the participation rate is almost 100% as shown in Section II, the treatments can be viewed to be orthogonal with respect to any personal characteristics. We display all regressions with and without demographic control variables to account for potential imbalances between various treatments.[11] Standard errors are clustered at the subject level to account for potential heteroscedasticity.

Our results demonstrate substantially higher entry rates in sequential treatments compared to the one-shot treatment. Differences between sequential and one-shot treatments are all significant at the 1% level in both experiments. In should be noted that, by design, taking the fair gamble many times does not result in a higher expected value nor

[10] The proportion of subjects who accept the one-shot gamble may be due to preferences for risk or entertainment or experimental demand effects.

[11] A list of all variables included in the paper is provided in Appendix B.

a lower probability of experiencing loss. Thus, the higher tendency to accept risk for multiple rounds cannot be explained by loss aversion as in the case of positive-expected-value gambles (see, Gneezy and Potters (1997)). It is, however, consistent with the CPT. As we show in Section III, skewness preferences in an EUT framework alone cannot explain our result of a higher tendency to accept a multi-round gamble because at any level of risk-aversion, agents would not accept the one-shot nor the multi-round gamble. RDU can only explain this result under the assumption that subjects' wealth level outside of the experiment is low. Assuming realistic levels of wealth would result in risk-loving behavior on the part of RDU agents; hence they would accept the gamble independent of whether it is one-shot or multi-round.

[Insert Figure 4 about here]

[Insert Table II about here]

Second, we find that subjects are more likely to accept the multi-round fair gamble if they could commit to an ex-ante strategy. This result follows from Figure 4 and Table II, Panel B. The latter presents the marginal fixed effects (mfx) of Probit regressions of the binary entry decision excluding the one-shot treatment; hence, the reference treatment is the sequential treatment, in which subjects decide round-by-round without any commitment device. Analogous to Panel A, standard errors are clustered at the subject level and we control for demographic characteristics as indicated. We find that subjects are significantly more likely to accept the multi-round gamble if they can commit to an ex-ante strategy with a binding commitment device.

One alternative explanation is that the complexity of the main task differs between the commitment treatments and the sequential treatment. In particular, since the commitment treatments restrict the strategy-choice set to strategies that can be described by a pair of limits, it may be the case that the perceived task-complexity is lower than the

task-complexity of the sequential treatment, and may explain why subjects in the commitment treatments have a higher tendency to enter. To test this claim, in Experiment 2 we elicit the perceived complexity of the main task using the four-item score of Maynard and Hakel (1997). In contrast to the complexity explanation, we find that subjects in the hard commitment treatment perceive higher complexity of the task than subject in the sequential treatment; this difference is significant at the 1% level. A related alternative explanation is that the higher acceptance rate in the commitment treatments is due to decision aversion as the number of decisions to be made in those treatments (i.e., 3) is on average smaller than the number of decisions in the sequential treatment. However, if decision aversion was driving the results, the treatment with the highest acceptance rate should have been the one-shot treatment, which is not the case.

The difference in entry rates between the commitment and sequential treatments suggests that participants are at least partly sophisticated about their dynamic inconsistency. To measure the scope of this sophistication, we distinguish between a binding and a non-binding commitment device. Figure 4 and Table II, Panel B, show that a non-binding commitment device also significantly increases the likelihood that participants enter and accept risk. The increased entry rate is comparable to the effect of a binding commitment device; the difference between a binding and a non-binding commitment device is not significant at the 10% level. Importantly, non-binding commitment devices are closer to commitment devices used in the real world. For instance, a gambler who has committed to limiting her losses by withdrawing a certain amount of cash may use credit or withdraw additional cash to deviate from his non-binding commitment. A FOREX trader, who has set certain stop-loss and take-profit limits can always revise them to prevent them from being executed. In Europe, advisors and portfolio managers are required by the recently introduced MiFID II regulation to inform their clients immediately in case their portfolio depreciates by more than 10% from the beginning of the quarter (i.e. depreciation reporting). Such reporting rules, just like any reminders, can be viewed as non-binding commitment devices. They highlight a particular dynamic strategy, namely

to take action in case of a loss, that the investor may ex-ante intend to follow but fail to execute ex post. Whether such real-world commitment devices have unintended, adverse effects on subjects' dynamic inconsistency depends on whether investors overly trust them ex ante. Our results on entry decisions suggest that subjects presume that non-binding commitment devices are just as effective against dynamic inconsistency as binding commitment devices. In the next section, we show that this assumption is largely incorrect as a large majority of the subjects override their ex-ante loss limits when they are non-binding.

B. Dynamics of Risk Taking: Ex-Ante Exit Plan versus Ex-Post Behavior

In this section we examine what dynamic strategies participants choose ex ante. We then study whether they follow these strategies in the absence of a commitment device, i.e. whether they are dynamically consistent or not. The manner of deviation allows us to distinguish between different theoretical explanations. In the subsections that follow, we analyze the ex-ante strategies and the ex-post behavior both in a between-subject setting (i.e., by comparing the sequential and the hard plan treatment in Experiments 1 and 2) and in a within-subject setting (i.e., ex-ante versus ex-post behavior in the soft plan treatment in Experiment 2).

B.1. Ex-Ante Plan

Figure 5 illustrates the cumulative distribution of gain and loss limits in our sample. Panel A covers all commitment treatments in both experiments, whereas Panel B displays only the non-binding limits in the soft commitment treatment in Experiment 2. First, the figure shows that consistent with CPT, the most likely loss limit is 10 cents, which means that subjects plan to exit the gamble as soon as they have lost one round more than they have won. In contrast, the most likely gain limit is a redundant limit of 260 cents.

The distribution of gain limits first-order stochastically dominates the distribution of loss limits. Table III further shows that the average subject who accepts the gamble sets a gain limit that is 3.81 times higher than her loss limit. That makes an average difference of 80.26 cents—more than 30% of the total endowment.Subjects can be classified according to their ex-ante strategy as "loss-exit" (i.e., loss limit closer to the reference level than gain limit), "gain-exit" (i.e. gain limit closer to the reference limit than loss limit) and "symmetric" (both limits are equidistant). The table shows that a large majority of 80.8% of subjects begin taking risk with a "loss-exit" strategy, whereas only 7%choose a "gain-exit" strategy.

[Insert Figure 5 about here]

[Insert Table III about here]

The prevalence of the loss-exit strategy leads to a positive skewness (0.31) of the expected final outcome distribution, as shown in Table III. To calculate this number, we run 100,000 independent outcome paths for each subject in our sample and determine the individual expected final outcome distributions, which we then aggregate to form the expected final outcome distribution of the representative subject. Note that a positive skewness, by design, can only result from an outcome-dependent dynamic strategy, as we use a symmetric gamble and the outcomes are linearly compounded across rounds. The skewness of the sample distribution of final outcomes in the commitment treatments is close to the skewness of the simulated outcome distribution - between 0.247 and 0.471. In Table IV we report the skewness of the outcome distribution among all subjects in the commitment treatments (see Panel A), only amongst those who enter (see Panel B), and the hypothetical outcome distribution if subjects in the soft plan treatment would stick to their limits (see Panel C). In all cases, the sample skewness is positive and statistically significant at the 5% level. The table further shows that the sample skewness in the sequential treatment is close to zero and not statistically significant at

the 10% level, unlike the sample skewness in the commitment treatments. Furthermore, we find that the actual outcome distribution of the subjects in the soft commitment treatment does not exhibit any statistically significant skewness at the 10% level, unlike their hypothetical outcome distribution. Both results indicate that the ex-post behavior deviates substantially from participants' prevailing "loss-exit" plans. In the following subsection, we shed light on exactly how subjects behave ex-post.

[Insert Table IV about here]

B.2. Ex-Post Behavior

We now proceed to analyze the ex-post behavior and compare it to the ex-ante plan. In particular, we aim to learn whether subjects' ex-post exit decisions are outcome-specific (i.e., different after gains versus losses) and whether the direction of the asymmetry between gains and losses is consistent with their ex-ante "loss-exit" plan.

First, we analyze the probability of exiting after gains versus losses. Figure 6 shows that subjects in the sequential treatment are more likely to end up with a cumulative gain than a cumulative loss even though the gamble is symmetric (see Panel A and Panel B). This is partly due to a substantially higher tendency (i.e, by 20 percentage points) to stop taking risks after the outcome of the first round is announced and it is a gain rather than a loss (see Panel C). The latter is direct evidence of a causal effect of losses on the probability to take more risks because the outcome after the first round is random and thus orthogonal to any subsequent choices. Critically, this pattern is opposite to subjects' ex-ante strategies. Table V, columns (1) and (2), compare the probability of realizing a cumulative gain versus a cumulative loss. The table displays the marginal effects of Probit regressions with and without demographic controls and with cluster-robust standard errors. We find that the probability of realizing a cumulative gain is 8.9 - 9.8 percentage points higher in the sequential treatment than in the plan treatment.

With respect to the causal effect based on the probability to exit after the first round, Table V, columns (3) and (4), display the results of OLS regressions with the main effects of a sequential-treatment dummy and a gain-dummy, as well as their interaction. It shows that the higher probability of realizing first-round gains than first-round losses is significantly different than subjects' ex-ante planned behavior in response to first round outcomes. The difference, as measured by the interaction term, is statistically significant at the 1% level.

[Insert Figure 6 about here]

[Insert Table V about here]

A higher probability of realizing gains, by design, does not mean that subjects in the sequential treatment earn more money on average. As the gamble is fair and the individual draws are iid, subjects can only achieve a higher probability of gains if they cut their gains early and chase their losses longer, thus earning small gains and enduring large losses. Figure 7 shows that this is indeed the case in our data. It compares the absolute cumulative gains and the absolute cumulative losses in the sequential treatment and the commitment treatments (Panel A). It further shows how long (i.e., percentage of rounds) subjects need to stay in a particular domain before they stop taking further risks as a function of whether they are in the gain versus loss domain. For both measures, the difference between gains and losses is of opposite sign in the sequential treatment compared to the commitment treatments. In particular, in the commitment treatments, subjects leave more room for large gains and restrict the scope of their losses (see Panel A). To achieve that, they are more patient in the gain domain to wait for larger gains to accumulate than they are in the loss domain, where they cut their losses sooner (see Panel B). Table VI shows that the difference-in-difference in both measures is statistically significant at the 1% level, indicating an ex-post deviation from subjects' ex-ante strategy.

[Insert Figure 7 about here]

[Insert Table VI about here]

Critically, the ex-post deviation is also present in the soft commitment treatment, where subjects are allowed to decide how to proceed after being notified that a limit has been triggered. From the 57 subjects whose limits were triggered, 80.7% decide to continue investing and thus deviate from their ex-ante plan, as shown in Figure 8. The majority of 70.2% of the triggered limits are loss limits. Moreover, the majority of subjects who deviate do so from the beginning until the end, thus repeatedly deciding against their ex-ante plan.

[Insert Figure 8 about here]

The presented patterns of deviations allow us to distinguish between different theoretical explanations for the observed dynamic inconsistency. Taken together, our findings are consistent with the predictions of CPT. The theory predicts that agents with an ex-ante "loss-exit" plan (majority of our sample) will exhibit "gain-exit" behavior ex post. This prediction is unique to CPT, as RDU predicts that agents near the parameter-region estimated in previous studies will end up taking risk until the end, independent of the outcome. This is not what we find in the data, as outcome valence (gain versus loss) causally affects participants' decision to continue taking risk.

V. Welfare Implications

In the following, we discuss the welfare implications resulting from dynamic inconsistency in sequential risk choices. To this end, we use Cumulative Prospect Theory as it best predicts our empirical findings outlined above. Following CPT, being dynamically inconsistent may result in two types of welfare loss that affect naïve and sophisticated agents differently. Naïve agents who invest in a sequential gamble do it because of an

ex ante illusion of welfare which they cannot achieve because of their dynamically inconsistent ex post behavior. The welfare of investing in the gamble with perfect foresight, i.e. correctly anticipating one's actual behavior, is lower than the welfare of rejecting the gamble altogether. Therefore, naïve agents who accept the gamble incur a welfare loss compared to sophisticated agents, who reject the gamble at the beginning and earn nothing. This type of welfare loss can thus be interpreted as *cost of naïvete*. In addition, both naïve and sophisticated agents incur a further type of welfare loss resulting from the absence of perfect commitment devices in the real world. This is because for dynamically inconsistent agents, commitment devices are the only way to achieve optimal behavior and thus maximize welfare. In the absence of a perfect commitment device subjects's choice is constrained to strategies that they can actually follow through. Following CPT, the constrained optimal choice for any CPT agent would be to reject the gamble and earn nothing. The difference between the maximum welfare possible with a perfect commitment device and the constrained maximum welfare, which is the welfare of non-participation, can thus be viewed as the *value of commitment*. In the following we use the CPT framework to measure these two types of welfare loss for various CPT agents.

We measure the welfare loss resulting from dynamic inconsistency in the experimental setting described in Section II, where the distribution of the gamble and the maximum length of the sequence is known. To measure welfare loss, for every CPT agent we calculate the certainty equivalent of the aggregate outcome distribution resulting from the ex ante plan and the ex post behavior, respectively. Outcome distributions are obtained from simulations as described in Section III, where the one-round investment amount is a numeraire. Hence, certainty equivalents are measured as multiples of the one-round investment amount.

[Insert Figure 11 about here]

Figure 11 reports the certainty equivalent of the ex ante outcome distribution (i.e.

value of commitment) and the ex post outcome distribution (i.e. cost of naïvete) for the representative agent as measured by Tversky and Kahneman (1992), as well as CPT agents with varying probability weighting and different levels of diminishing sensitivity compared to the representative agent. A positive certainty equivalent of the outcome distribution implied by the ex ante plan (see Panel B) indicates how much a sophisticated agent would be willing to pay for a perfect commitment device which would allow her to benefit from her ex ante plan as opposed to rejecting the gamble and earning nothing. Note that the value of commitment for the representative agent is close to zero by construction. This is because the maximum number of rounds is deliberately chosen such that the representative agent would be close to indifferent whether to accept the sequential gamble or not. Agents with stronger skewness preferences (i.e. lower δ) and lower sensitivity (i.e. higher α) than the representative agent would be willing to pay up to 166% of the one-round investment amount depending on the exact parameter combination, which corresponds to 6.4% of their total endowment.[12] In the absence of a commitment device both naïve and sophisticated agents incur a welfare loss corresponding to the value of commitment reported in Panel B. In addition, naïve agents further incur a welfare loss because they are naïve about their own ability to stick to their ex ante plan and accept the sequential gamble instead of rejecting it. These costs of naïvete are illustrated in Figure 11, Panel A. Notably, all CPT agents who would optimally select a loss-exit strategy incur positive costs of naïvete as indicated by the negative certainty equivalents for all parameter combinations with high skewness preferences and low sensitivity. For CPT agents in the region adjacent to the representative agent, the costs of naïvete are over 110% of the one-round investment amount or over 4.2% of their total endowment. The costs of naïvete are higher for CPT agents with higher preferences for skewness and lower sensitivity. Depending on the exact parameter combination, a CPT agent may incur a welfare loss corresponding to a certain loss of up to 208.5% of the

[12] A number of CPT agents with particularly low skewness preferences and high sensitivity would commit to a gain exit plan, as discussed in Section III. However, the certainty equivalent of a gain-exit plan for CPT agents for whom it is optimal is always lower than 3% of the one-round investment amount and thus barely distinguishable from zero.

one-round investment amount (or 8% of the total endowment).

Note that whether a subject incurs costs of naïvete or not depends not only on her subject-specific level of sophistication but also on the type of commitment device that is available. Soft commitment devices, which is the type of commitment devices typically available in the real world, have the disadvantage that a large majority of subjects are naïve about their lack of effectiveness, as shown in Section IV. In welfare terms, the introduction of soft commitment devices compared to no commitment devices would make a large majority of the population, who wrongfully believe in an effectiveness of soft commitment devices, incur significant costs of naïvete and only a small minority of the population, for whom soft commitment devices are indeed effective, gain a value of commitment. Hence, any regulation that introduces soft commitment devices for dynamic risky environments or addresses third-parties' discretionary power to design soft commitment devices for consumers and retail investors should take into account the trade-off between the benefits of such devices to a minority and the costs to the majority of consumers and retain investors.

VI. Dynamic Inconsistency in the Field

To test for dynamic inconsistency in the field, we use trading data from a large international online brokerage from June 2013 and August 2015. The data contains 187,521 traders from over 150 countries, 84% of whom are male. The dataset covers transactions of contracts for difference (CFD) — speculative (long and short) and self-initiated (non-advised). The majority of the transactions during this sample period are for CFDs in major currencies (e.g, EUR/USD, USD/JPY and GBP/USD). The majority of trades are levered at the time of purchase using margin provided by the brokerage (the median (average) leverage of the trades is 100:1 (163:1)). Relative to other brokerage data sets used by the literature (e.g., Barber and Odean (2001); Linnainmaa (2010)),

a distinguishing feature of our data set is that it contains limit orders (stop-loss and take-profit orders) and specifically the *revisions* traders make to such orders while their trading positions are open.

The setting of the online broker overlaps closely with our experimental design. First, our unique dataset allows us to measure the traders' ex-ante strategies and ex-post revisions as accurately as it is possible in a real-world setting. It contains data on stop-loss and take-profit limit orders that traders are required to submit at the time of opening every position. These limits are mandatory but they are not binding, as traders revise their limit orders after experiencing gains and losses on open positions. Thus, the empirical setup is most closely related to the soft commitment treatment in our experiment. We observe every revision of the limits which allows us to test for the link between gains versus losses and the decision to close a position. Second, the holding period of the transactions is short with a median of 3.6 hours (average of 3.5 days); hence, informational shocks are less likely to play a role relative to settings with longer holding lengths, analogous to the experimental setting. Third, the CFD of currency pairs are characterized by negative expected returns, hence similarly to our experiment, any dynamic inconsistency cannot potential arise from return aggregation as in Benartzi and Thaler (1995).

[Insert Figure 9 about here]

A. Results

Figure 9 shows the distribution of ex-ante plans. A plan is defined as a combination of a stop-loss and a take-profit limit order that traders submit when they initiate the position. The plans are categorized as loss-exit plan (gain-exit plan) if the stop-loss limit (take-profit limit) is closer to the reference level than the take-profit limit (stop-loss limit). We assume the reference level is the opening quoted price of the asset. This assumption is plausible because the brokerage displays position-level gains and losses relative to the

opening quote. The figure shows that the majority, approximately 40% of positions, follows a loss-exit strategy. This result is qualitatively similar to our experimental findings. Differences in the magnitude might be due to the fact that an equidistant default is the maximum limit size (same for take-profit and stop-loss orders), which may constitute a nudge towards equidistant plans.

[Insert Figure 10 about here]

Figure 10 illustrates how traders deviate from their ex-ante plans.[13] The figure displays the distribution of actions that traders undertake conditional on earning a cumulative (paper) gain since the opening of the position versus earning a cumulative (paper) loss. It shows that the most frequently undertaken action in times of losses is to postpone the planned exit by pushing away the stop-loss limit. Conversely, in times of gains, subjects are most likely to close the position before their take-profit limit has been reached. Moreover, we show that this result is almost identical in the subsample of traders whose ex-ante plan was categorized as "loss-exit." For this subsample of traders the ex-ante plan and subsequent ex-post deviation are consistent with our experimental findings and with the predictions of CPT.[14]

---

[13] Gains and losses are evaluated in real-time using all trades on the online broker's platform to estimate the bid and ask quotes of each underlying on a ten-minute frequency. The figure illustrates the distribution of different actions conditional on undertaking any action. This condition is necessary as we do not have any data on when traders are paying attention to their trading account.

[14] As a further robustness check, in Panel B we focus only on single positions, that have no contemporaneously open position that may distract the trader or introduce ambiguity into the reference point assumption. Single positions constitute 3% of the dataset. Even though selection effects are possible given the small sample size of the subsample, Figure 9 shows that single positions are characterized by similar ex-ante plan distribution as the rest. It should be noted, however, that traders' are more likely to stick to their ex-ante plan when trading single positions. This is indicated by a higher probability to close the position automatically upon triggering a limit relative to the probability of any other action and compared to all positions. This might be due to traders' distraction from their trading account that might be the reason why the trader has only left single position open. In support of this explanation, we observe that the median single position is held longer (5.4 hours compared to the median over all positions which is 3.6 hours).

It should be noted that traders in the real world do not have any binding restriction with respect to the holding time, unlike the subjects in our experiment, who are restricted to investing in the lottery for a maximum of 26 rounds. Under the assumption of an infinite planning horizon, however, agents are predicted to continue gambling independent of the outcome (see Ebert and Strack (2015)). In particular, they are expected to continue investing after earning gains, not to exit manually before their take-profit limit is triggered. This is because the longer the agent can invest in the lottery, the higher is the potential skewness that she can generate with a dynamic loss-exit strategy. If the planning period is long enough, the prospective skewness justifies to continue gambling even for CPT agents who find themselves in the gain domain and are therefore risk averse (in the concave part of the value function). A finite planning horizon, in contrast, restricts the potential skewness of a dynamic loss-exit strategy. Hence, as the number of remaining investment decisions decreases, so does the CPT agent's willingness to continue investing if she is in the gain domain. Our results show that even in a real-world setting with no explicit restriction of the planning horizon, traders behave consistent with the predictions of the model with a finite horizon (see Barberis (2012)). The finite horizon may arise from traders' latent restrictions, for example, limits in the provision of leverage, trader-specific trading style, traders' attention capacity etc. This further highlights the empirical validity of CPT as a positive model of dynamic decision-making under risk.

VII. Conclusion

We show that people are dynamically inconsistent when taking risks repeatedly while knowing that they have the option to stop whenever they like. In particular, they take 'bad' risks too often because they plan to stop as soon as they make small losses and continue as long as they win to reap large gains. However, they fail to execute their ex-ante plan and behave in an opposite way ex-post if they cannot fully commit to an ex-ante plan. Even though subjects are partially sophisticated to the extent that they

reject taking risks in the absence of any commitment device, they are easily lured into accepting risk if non-binding commitment devices are provided, which are just as effective as absent commitment devices for most people.

In the real world, people sometimes have commitment devices at disposal but they are non-binding at best. Our findings of a dynamic inconsistency taken together with people's prevailing naïvete about the effectiveness of non-binding commitment devices highlights the need for regulation to address the discretion of the industry (e.g., casinos, credence goods providers, financial brokers etc.) to design and, in particular, to soften the commitment devices available to their clients. Moreover, our findings highlight the need for regulation itself to be evidence-based in order to avoid unintended adverse effects on consumers.

Particularly, the results highlight the need for regulation of consumer financial product to be based on evidence to prevent unintended adverse effects. For example, consider the recently introduced "depreciation reporting rule", which is a part of the revised European market in financial instruments regulation (MiFID II) that came into force in January 2018.[15] The rule requires all wealth manager, brokers and financial advisers in Europe to immediately notify their clients in case their portfolio exhibits a loss of 10% of higher from the portfolio value at the beginning of the quarter. The reporting rule Although the intention behind it might have been different, the rule ultimately constitutes an exogenous non-binding loss limit at the level of 10%. Following our results, the notifications will be utterly redundant for the majority of investors as they are likely to ignore it. More importantly, however, the rule might change investors' ex ante perspective on investing as they would anticipate receiving notifications ex ante. Our results suggest that investors might be more likely to take risks after the introduction of the rule, most importantly even investing in assets with zero or negative risk premium, because they might erroneously presume that they will reduce or close their positions after receiving

[15] See Directive 2014/65/EU of the European Parliament on the market in financial instruments (MiFID II). For European legislation see `https://eur-lex.europa.eu/`.

a notification, thus limiting their losses at 10%. This way, instead of helping investors make better financial decisions, the regulation may pave the way into higher share of speculative trades and more (costly) dynamic inconsistency on the part of naïve investors. In addition, the finance industry provides investors with a variety of alert subscriptions, notifications and reminders on trading and information platforms, which may be used as non-binding limits and are currently largely unsupervised and unregulated. In light of our findings, regulation should address the finance industry's discretionary power to design and provide such tools to their clients.

VIII. Figures and Tables

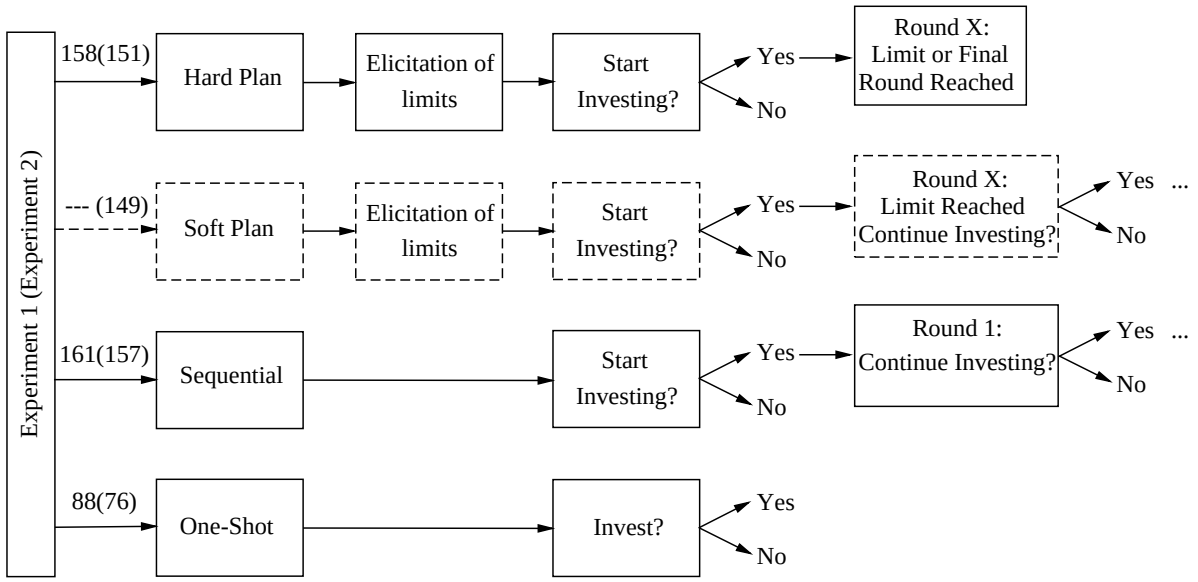
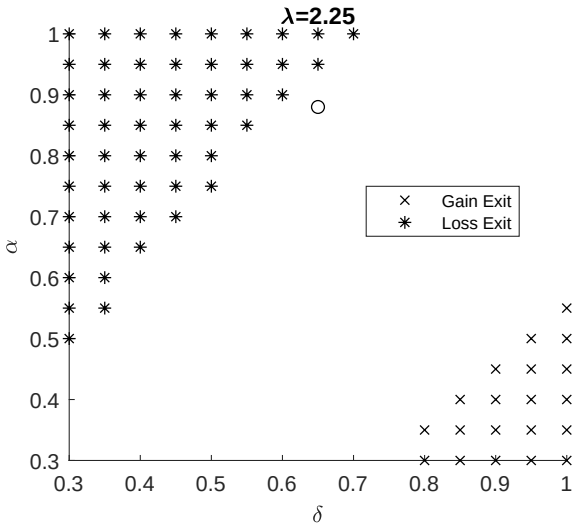
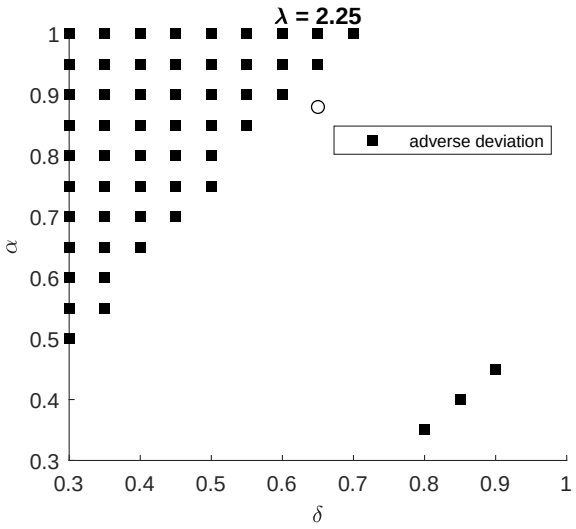


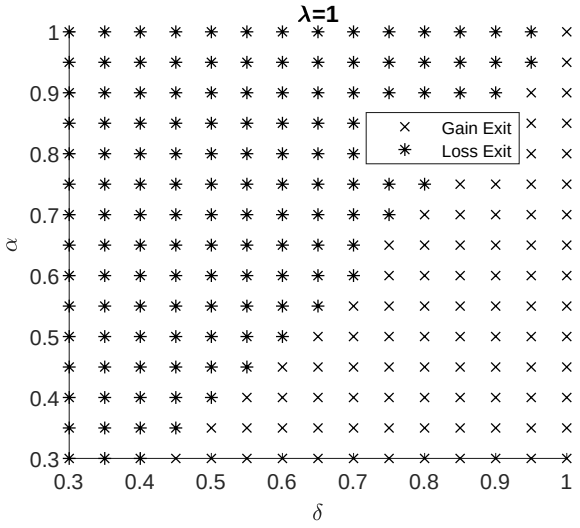
Figure 1. Experimental design and sample size. This figure shows the experimental design of Experiments 1 and 2. The sample sizes in each between-subject treatment in Experiment 1 (Experiment 2) are displayed (in parentheses).



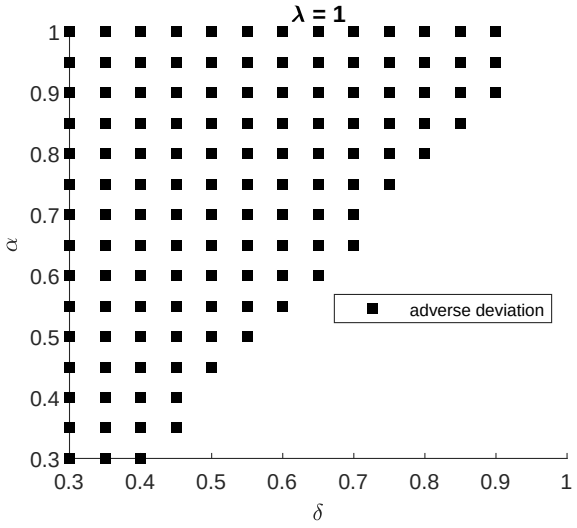
Panel A. Ex ante plan ($\lambda = 2.25$)



Panel B. Ex post deviation ($\lambda = 2.25$)

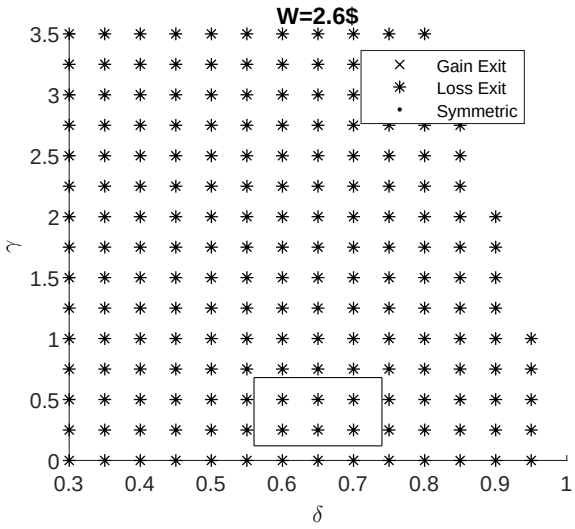


Panel C. Ex ante plan ($\lambda = 1.00$)

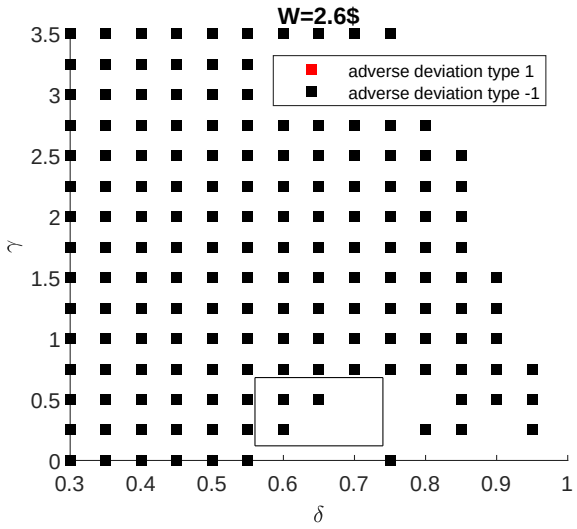


Panel D. Ex post deviation ($\lambda = 1.00$)

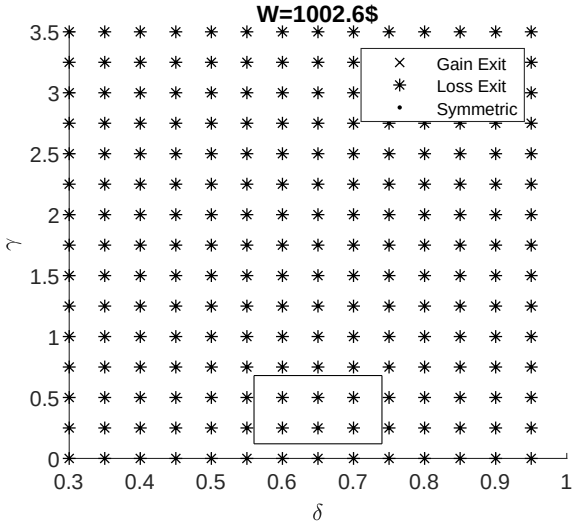
Figure 2. Theoretical predictions of Cumulative Prospect Theory. This figure illustrates the ex ante plan (Panel A and C) and the ex post deviation (Panel B and D) of agents with CPT preferences with different levels of probability weighting (δ) and diminishing sensitivity (α), i.e. the extent to which agents are risk-averse in the gain domain and risk-seeking in the loss domain. Panels A and B report the results for loss-averse agents (i.e., positive λ), while Panels C and D report the results for agents who are not loss-averse (i.e., $\lambda = 1$). Ex ante plans are categorized as "gain exit" ("loss exit"), which implies that the gain limit is closer to (further away from) the reference point. Ex post behavior is categorized as adverse deviation if the ex-post outcome distribution has a lower value than rejective the gamble at the beginning, in $t = 0$. The representative agent as estimated by Tversky and Kahneman (1992) is marked by a circle in Panels A and B.



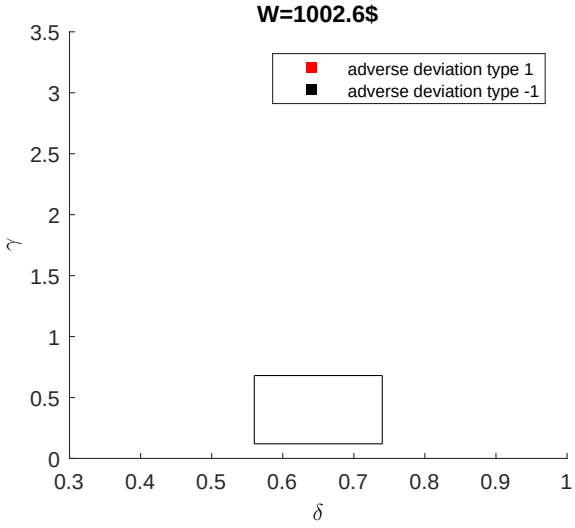
Panel A. Ex ante plan ($W = 2.6\$$)



Panel B. Ex post deviation ($W = 2.6\$$)



Panel C. Ex ante plan ($W = 1002.6\$$)



Panel D. Ex post deviation ($W = 1002.6\$$)

Figure 3. Theoretical predictions of Rank Dependent Utility. This figure illustrates the ex ante plan (Panel A and C) and the ex post deviation (Panel B and D) of agents with RDU preferences with different levels of probability weighting (δ) and risk aversion (γ). Ex ante plans are categorized as "gain exit" ("loss exit"), which implies that the gain limit is closer to (further away from) the initial wealth than the loss limit, and "symmetric", which implies equidistant limits. Ex post behavior is categorized as adverse deviation if the ex-post outcome distribution has a lower utility than the utility of rejecting the gamble at the beginning, in $t = 0$. In addition, "type 1" deviation implies that the agent almost certainly continues investing until the final round. Other types of deviation are denotes as "type -1". W indicates the assumption about initial wealth. The parameter-region estimated from previous experimental studies is marked by the box.

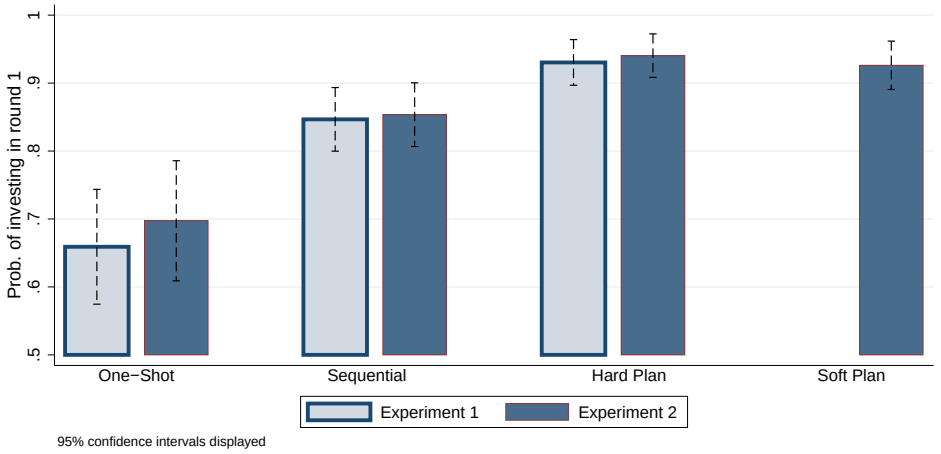
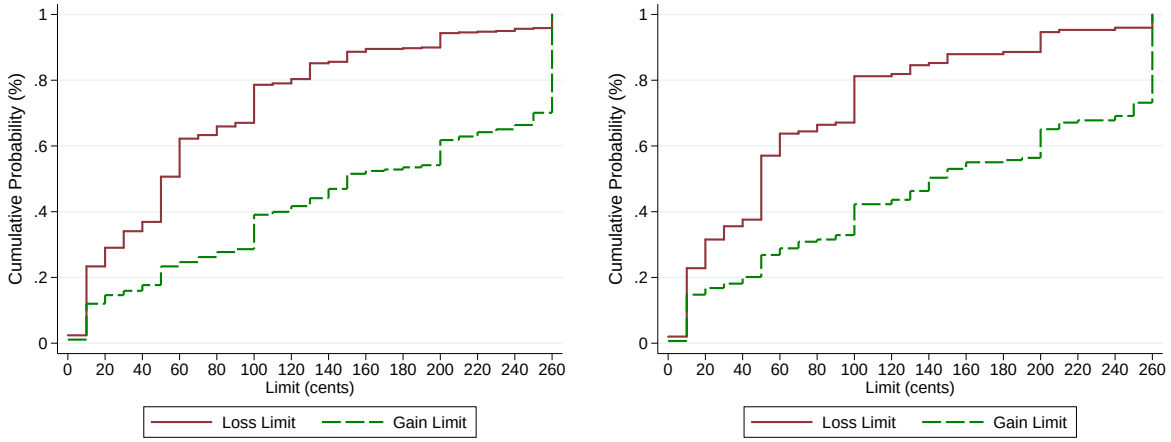


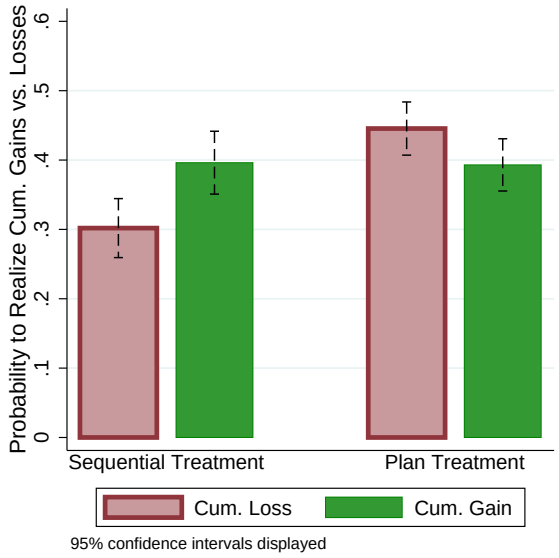
Figure 4. Entry decision. This figure shows the percentage of subjects in each experiment and treatment who choose to accept the gamble in the first round.



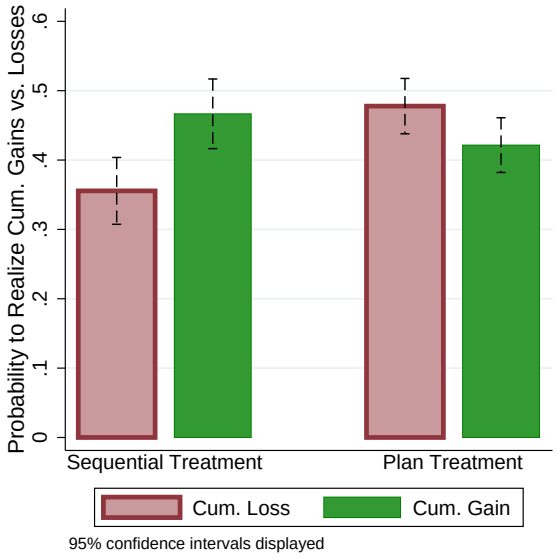
Panel A. All Subjects

Panel B. Soft Plan Treatment (Exp. 2)

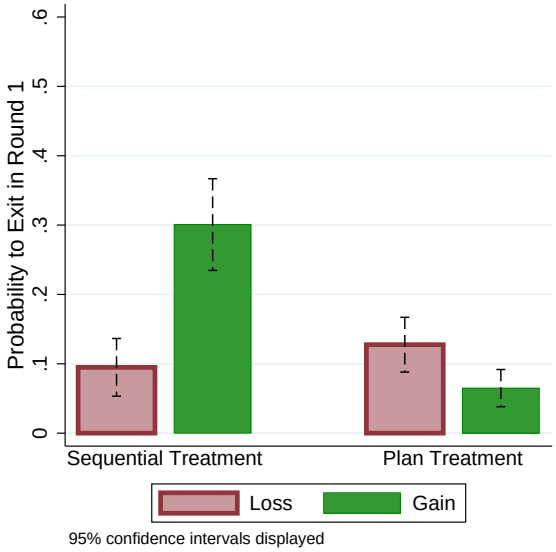
Figure 5. Ex-ante exit plan. This figure illustrates the exit strategy of the subjects in the plan treatments. It shows the cumulative distribution of loss and gain limits. Panel A covers both experiments and plan treatments (hard and soft). Panel B covers the subsample of subjects in the soft plan treatment in Experiment 2.



Panel A. All Subjects

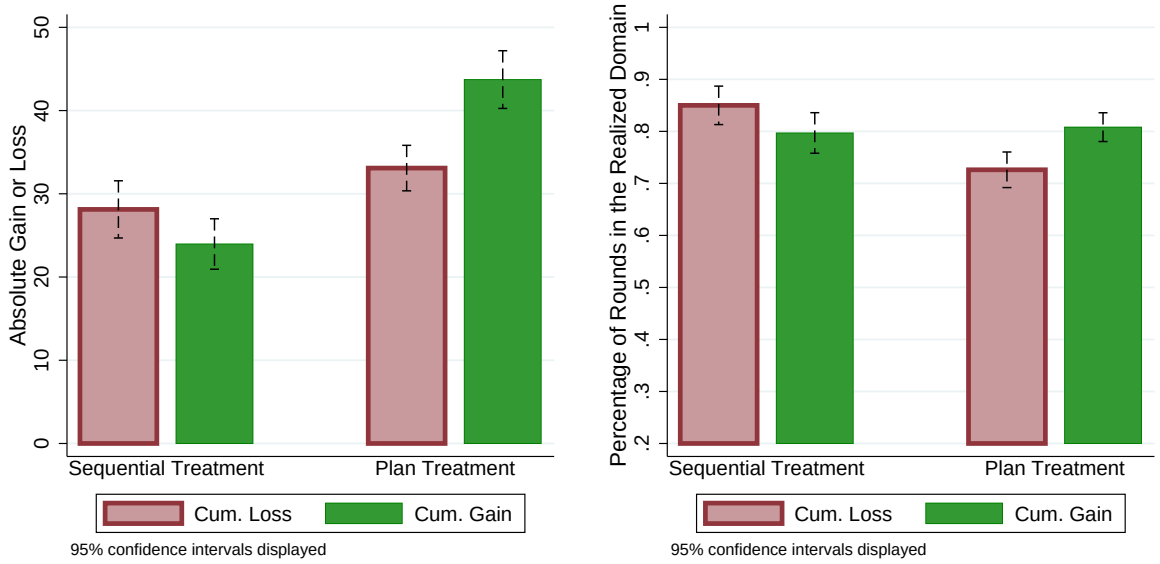


Panel B. Entered Lottery



Panel C. Exit in Round 1

Figure 6. Probability of realizing a gain versus loss — ex-ante versus ex-post. This figure illustrates the ex-post behavior in the sequential treatment and the ex-ante plan in the plan treatment after gains versus losses. The plan treatment includes the hypothetical cumulative outcomes of the subject in the soft plan treatment, hence in case the subject deviates from her plan we replace the actual outcome with the outcomes at the point of time when her limits were first triggered. Panel A shows the percentage of subjects who earn a cumulative gain versus a cumulative loss from all subjects in Experiments 1 and 2. Panel B excludes subjects who do not enter the lottery at the beginning. Panel C shows the percentage of subjects who enter the lottery and stop investing after the earnings in Round 1 (gains versus losses) are revealed. Corresponding regression results are shown in Table V.



Panel A. Absolute Cum. Gains/Losses

Panel B. Realization Reluctance

Figure 7. Path to realizing a gain versus loss — ex-ante versus ex-post. This figure illustrates differences in the outcomes and behavior between subjects who realize a cumulative loss versus gain in the plan (ex-ante) versus sequential treatment (ex-post) in Experiments 1 and 2. The plan treatment includes the hypothetical cumulative outcomes of the subject in the soft plan treatment, hence in case the subject deviates from her plan we replace the actual outcome with the outcomes at the point of time when her limits were first triggered. Panel A shows the absolute cumulative gains/losses in sequential versus plan treatments. Panel B illustrates differences in subjects' reluctance to realize their final cumulative outcome as measured by the percentage of rounds the subjects' cumulative gain has been in the domain that the subject ended up realizing. Corresponding regression results are shown in Table VI.

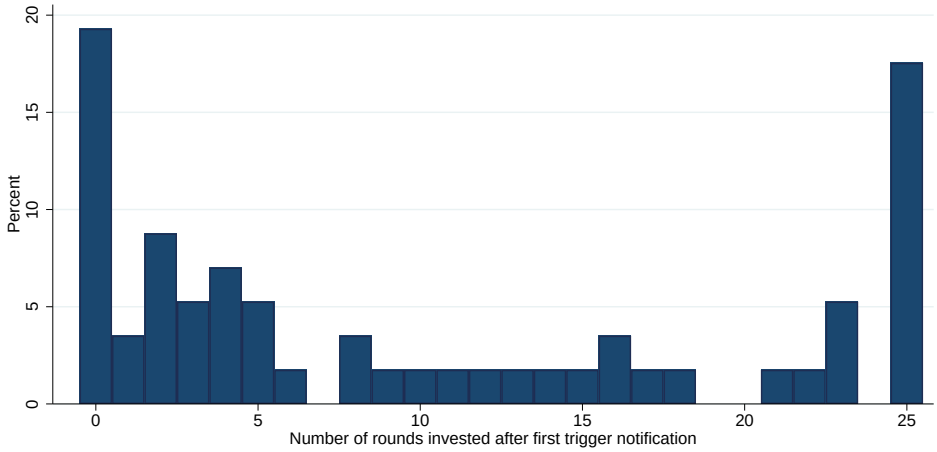
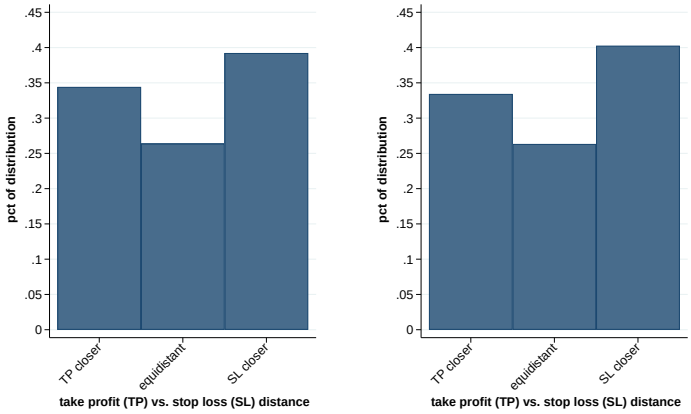
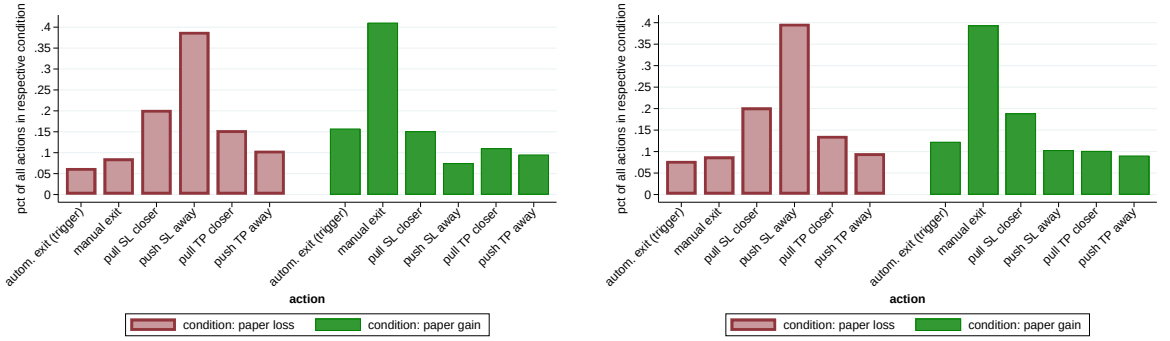


Figure 8. Ex-post deviation in soft plan treatment — how many and for how long? This figure shows the sample distribution of the duration of investment (in rounds) after a limit has been triggered for the first time. The data for this figure consists of 57 observations of subjects in the soft plan treatment whose limits are triggered during the experiment. A duration of zero indicates that the subject sticks to her ex ante plan. The maximum duration is 25 rounds.



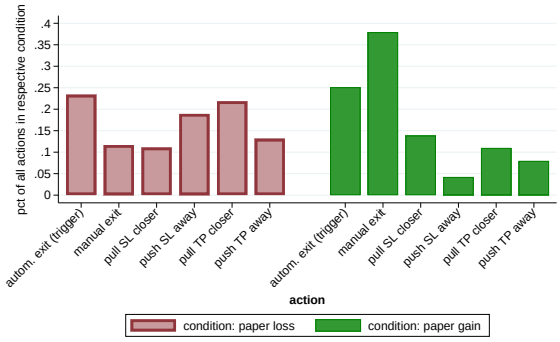
Panel A. All Positions Panel B. Single Positions

Figure 9. Online traders' ex-ante plan This figure illustrates the distribution of ex-ante strategies that consist of a pair of limits which are set at the opening of the position (mandatory) – a stop-loss (SL) and a take-profit (TP) limit. The positions are categorized as follows. First, a loss-exit plan is given by a stop-loss limit that is closer to the opening price than the contemporaneously submitted take-profit limit. Second, a gain-exit plan is given by a stop-loss limit that is further away from the opening price than the corresponding take-profit limit. Third, symmetric plan is characterized by equidistant limits. The percentage of positions categorized as ex-ante symmetric may be overestimated due to equidistant default limits suggested by the trading platform. Panel A includes all positions. Panel B includes only positions for which the trader does not have a contemporaneously open other position.



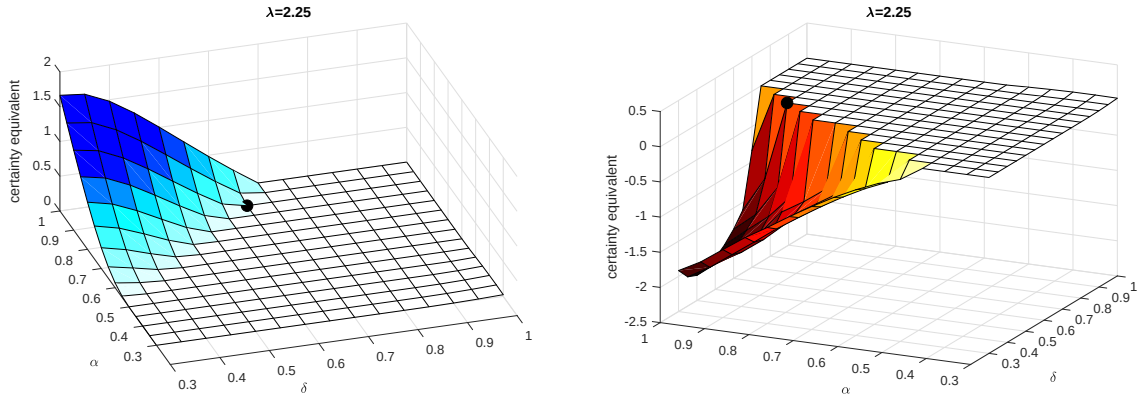
Panel A. All Positions

Panel B. Positions with loss-exit plan (ex ante)



Panel C. Single Positions

Figure 10. Online traders' reaction to paper gains versus paper losses This figure illustrates the distribution of actions undertaken by FOREX traders conditional on the non-realized (i.e. paper) outcome of their positions. Paper gains and paper losses are calculated based on respective bid and ask prices constructed from the trading activity on the platform. Panel A shows the traders' reaction to paper gains and paper losses in all positions. Panel B includes only positions that were started with a loss-exit plan (i.e. initial stop-loss limit closer to opening price than initial take-profit limit). Panel C includes only positions for which the trader does not have a contemporaneously open other position.



Panel A. Ex ante plan

Panel B. Ex post behavior

Figure 11. Welfare implications of ex post deviation: value of commitment and costs of naïvete. This figure illustrates the certainty equivalent of agents with CPT preferences with different levels of probability weighting (δ) and diminishing sensitivity (α), i.e. the extent to which agents are risk-averse in the gain domain and risk-seeking in the loss domain, and loss-aversion of $\lambda = 2.25$, corresponding to the representative agent as estimated by Tversky and Kahneman (1992). The certainty equivalent is measured as a multiple of the one-round investment amount (which corresponds to 1/26 of the total endowment). Panel A reports the certainty equivalents of the outcome distribution which would be generated by the agent's ex-ante plan. It can be interpreted as the *value of commitment* from the point of view of a sophisticated agent, who would participate in the lottery if she can commit to her ex ante plan and stay out of it (and earn zero) otherwise. Panel B reports the certainty equivalent of the ex post outcome distribution. It can be interpreted as *costs of naïvete* that agents endure if they participate in the sequential lottery without a commitment device as opposed to staying out of it (and earning zero).

Table I
Overview of Demographics

This table reports sample statistics of demographic characteristics elicited with an entry-level questionnaire before the main experimental task. Dummy variables are indicated with "(D)" and the range of categorical variables is indicated in parentheses. Columns (1) and (2) present the sample statistics of Experiment 1, while columns (3) and (4) show the respective results for Experiment 2. The $\underline{z}$-Statistic of a nonparametric Mann-Whitney test is reported in column (5).

	Experiment 1		Experiment 2		Mann-Whitney
	(1) μ	(2) σ	(3) μ	(4) σ	(5) $\underline{z}$-stat $(\mu_1 - \mu_2 = 0)$
Age	34.55	10.34	34.88	11.00	0.04
Male (D)	0.60	0.49	0.56	0.50	1.18
Statistical Skills (1-6)	3.31	1.29	3.52	1.30	-2.59
Study Business (D)	0.14	0.35	0.17	0.38	-1.14
Study Comp Sciences (D)	0.18	0.39	0.18	0.38	0.16
Study Econ (D)	0.04	0.19	0.04	0.19	-0.20
Study Math (D)	0.02	0.15	0.02	0.14	0.16
Study Statistics (D)	0.00	0.05	0.00	0.06	-0.35
Study Psychology (D)	0.04	0.21	0.06	0.23	-0.83
Highest Education (1-6)	3.42	1.02	3.53	1.03	-1.56

Table II
Sequence and Commitment Effects on Entry Decision

This table reports the coefficients of Probit regressions of the decision whether or not to start investing in round 1. The main independent variables are dummy variables for sequential treatment (D^{seq}), hard and soft plan treatment ($D^{hardplan}$ and $D^{softplan}$). Panel A shows results including the one-shot treatment as a reference group. Panel B shows results excluding the one-shot treatment, hence the reference group is the sequential treatment. Columns (1) and (2) show results based on the combined dataset of Experiments 1 and 2 where the soft plan treatment, which is unique for Experiment 2, is excluded. We include demographic variables elicited in an entry-level questionnaire. These variables are age, gender, study field (dummies), highest level of education, self-reported statistical skills. Standard errors are cluster-robust. t-Statistics are in parentheses. *, ** and *** indicate statistically significant at the 10%, 5%, and 1% level, respectively.

Panel A

	Experiments 1 & 2		Experiment 1		Experiment 2	
	(1) mfx	(2) mfx	(3) mfx	(4) mfx	(5) mfx	(6) mfx
D^{seq}	0.125***	0.125***	0.136***	0.139***	0.102***	0.097***
	(4.418)	(4.433)	(3.409)	(3.586)	(2.775)	(2.710)
$D^{hardplan}$	0.230***	0.229***	0.241***	0.239***	0.198***	0.195***
	(7.351)	(7.356)	(5.511)	(5.695)	(4.765)	(4.784)
$D^{softplan}$					0.177***	0.174***
					(4.399)	(4.386)
Demographics	No	Yes	No	Yes	No	Yes
Pseudo R^2	0.077	0.090	0.080	0.130	0.071	0.093
N	791	791	407	407	533	533

Panel B

	Experiments 1 & 2		Experiment 1		Experiment 2	
	(1) mfx	(2) mfx	(3) mfx	(4) mfx	(5) mfx	(6) mfx
$D^{hardplan}$	0.088***	0.089***	0.087**	0.087**	0.083**	0.083**
	(3.435)	(3.520)	(2.397)	(2.553)	(2.462)	(2.522)
$D^{softplan}$					0.065**	0.065**
					(2.006)	(2.030)
Demographics	No	Yes	No	Yes	No	Yes
Pseudo R^2	0.029	0.044	0.027	0.086	0.027	0.040
N	627	627	319	319	457	457

Table III
Ex-Ante Plan

This table illustrates the exit strategy in the plan treatments of Experiments 1 and 2. Panel A reports the share of subjects who have a loss-exit and gain-exit strategy. A loss-exit (gain-exit) strategy is defined as lower (greater) loss limit than gain limit. Column (1) to (4) reports the results for all subjects in the plan treatments. Columns (5) to (8) reports the results only for subjects who invest in the lottery at the beginning. Panel B reports aggregate statistics to illustrate the magnitude of the difference between gain and loss limits. "Ratio" refers to the ratio between the gain and loss limit ($\frac{gain}{loss}$) and "Difference" refers to their difference ($gain - loss$). t-statistics of Wald tests for $H0 : Ratio = 1$ and $H0 : Diff = 0$, respectively, are in parentheses. Panel C reports the mean and the skewness of the aggregate outcome distribution that results from subjects' gain and loss limits. The outcome distribution for each subject (each set of gain and loss limits) results from 100,000 independently simulated outcome paths.

	All Subjects				Entered Lottery			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Experiment	1&2	1	2 (hard)	2 (soft)	1&2	1	2 (hard)	2 (soft)
N	458	158	151	149	427	147	142	138
Panel A. Strategy Categorization								
Lossexit	80.3%	85.4%	78.1%	77.2%	80.8%	86.4%	78.9%	76.8%
Symmetric	12.7%	10.1%	14.6%	13.4%	12.2%	8.8%	14.1%	13.8%
Gainexit	7.0%	4.5%	7.3%	9.4%	7.0%	4.8%	7.0%	9.4%
Panel B. Aggregate Statistics								
Ratio	3.81***	3.67***	4.51***	3.24***	3.63***	3.56***	4.25***	3.05***
	(10.71)	(6.17)	(6.45)	(6.19)	(10.11)	(5.79)	(6.11)	(5.87)
Difference	80.26***	85.00***	81.06***	74.43***	79.98***	86.26***	80.14***	73.12***
	(19.67)	(13.84)	(10.51)	(10.13)	(19.01)	(13.63)	(10.22)	(9.52)
Panel C. Outcome Distribution								
Mean	-0.03	-0.15	0.10	0.14	0.04	-0.09	-0.07	0.01
Skewness	0.31	0.30	0.33	0.33	0.31	0.28	0.32	0.29

Table IV
Commitment Effect on Outcome Distribution

This table reports sample distribution parameters of the outcome distributions in different treatments. The outcome is the cumulative gain or loss from the beginning of the sequential lottery until the subject chooses to stop investing. Panel A compares the plan treatments from Experiments 1 and 2 to the sequential treatments. For the soft plan treatment in Experiment 2 we take the outcome at the point of time when one of the limits was triggered for the first time. In case no limit was triggered, we take the final outcome. Panel B shows the results for subjects enter the lottery. Panel C focuses only on the soft plan treatment in Experiment 2 and compares the hypothetical outcome that would have been reached with the limits (i.e. ex ante) to the actual outcome that was reached at the end (i.e. ex post). We report the p-values of Wald tests comparing the means of the distributions to zero. In addition, we use Jarque-Bera test for the skewness of the outcome distribution. *, ** and *** indicate statistically significant at the 10%, 5%, and 1% level, respectively.

Panel A: Between-Subject Test (Experiment 1 and 2)

	N	Mean	Skewness
Plan Treatment	458	2.445	0.268**
		(0.217)	(0.019)
Sequential Treatment	318	1.006	0.060
		(0.515)	(0.656)

Panel B: Between-Subject Test (Experiment 1 and 2)

	N	Mean	Skewness
Plan Treatment	427	1.185	0.247**
		(0.515)	(0.037)
Sequential Treatment	270	2.623	0.037
		(0.217)	(0.798)

Panel C: Within-Subject Test (Experiment 2)

	N	Mean	Skewness
Soft Plan Treatment (Ex Ante)	138	1.812	0.471**
		(0.618)	(0.023)
Soft Plan Treatment (Ex Post)	138	2.319	0.280
		(0.569)	(0.166)

Table V
Probability of Realizing Gain versus Loss — Ex-Ante versus Ex-Post

This table reports the coefficients of OLS and the marginal fixed effects of Probit regressions. The dependent variable in Columns (1) and (2) is a dummy variable that equals one if the subject's final outcome is a cumulative gain as opposed to a cumulative loss (zeros are excluded). In Columns (3) and (4), the dependent variable is a dummy variable that equals one if the subject stops investing in Round 1 after the first outcome of the lottery is revealed (excluding subjects who did not start investing at all). The main independent variables are dummy variables for sequential treatment (D^{seq}), drawing a gain versus loss in Round 1 (D^{1gain}). The plan treatment includes the hypothetical cumulative outcomes of the subject in the soft plan treatment, hence in case the subject deviates from her plan we replace the actual outcome with the outcomes at the point of time when her limits were first triggered. We include demographic variables elicited in an entry-level questionnaire. These variables are age, gender, study field (dummies), highest level of education, self-reported statistical skills. Standard errors are cluster-robust. t-statistics are in parentheses. *, ** and *** indicate statistically significant at the 10%, 5%, and 1% level, respectively. The results of this table are illustrated in Figure 6.

	Realizing a Cum. Gain versus Cum. Loss		Stop Investing in Round 1	
	(1)	(2)	(3)	(4)
D^{seq}	0.098**	0.089**	-0.099***	-0.112***
	(2.376)	(2.154)	(-2.615)	(-2.960)
D^{1gain}			-0.081**	-0.092***
			(-2.312)	(-2.653)
$D^{seq} \times D^{1gain}$			0.287***	0.305***
			(4.884)	(5.203)
Demographics	No	Yes	No	Yes
Pseudo R^2	0.007	0.017		
R2			0.040	0.060
N	606	606	697	697

Table VI
Path to Realizing a Cumulative Gain versus Loss

This table reports the coefficients of OLS regressions. The dependent variables are the absolute cumulative gains/losses (Columns (1) and (2)), and the subjects' reluctance to realize a gain or a loss (Columns (3) and (4)). The reluctance to realize gains or losses is measured by the percentage of rounds the subject has had a paper cumulative gain (loss) if the subject realized a cumulative gain (loss) at the end. The main independent variables are dummy variables for sequential treatment (D^{seq}), a dummy variable for realizing a cumulative gain versus loss (D^{gain}). The plan treatment includes the hypothetical cumulative outcomes of the subject in the soft plan treatment, hence in case the subject deviates from her plan we replace the actual outcome with the outcomes at the point of time when her limits were first triggered. We include demographic variables elicited in an entry-level questionnaire. These variables are age, gender, study field (dummies), highest level of education, self-reported statistical skills. Standard errors are cluster-robust. t-statistics are in parentheses. *, ** and *** indicate statistically significant at the 10%, 5%, and 1% level, respectively. The results of this table are illustrated in Figure 7.

	Absolute Cumulative Gain/Loss		Realization Reluctance	
	(1)	(2)	(3)	(4)
D^{gain}	10.634***	11.295***	0.082***	0.085***
	(3.979)	(4.327)	(3.078)	(3.165)
D^{seq}	-4.963*	-4.016	0.124***	0.126***
	(-1.875)	(-1.506)	(4.079)	(4.036)
$D^{seq} \times D^{gain}$	-14.791***	-15.937***	-0.135***	-0.135***
	(-3.852)	(-4.124)	(-3.224)	(-3.173)
Demographics	No	Yes	No	Yes
R2	0.087	0.119	0.030	0.046
N	606	606	606	606

REFERENCES

Andrade, Eduardo B., and Ganesh Iyer, 2009, Planned versus actual betting in sequential gambles, Journal of Marketing Research 46, 372–383.

Arditti, F., 1967, Risk and the required return on equity, The Journal of Finance 22, 19–36.

Barber, Brad M, and Terrance Odean, 2000, Trading is hazardous to your wealth: The common stock investment performance of individual investors, The Journal of Finance 55, 773–806.

Barber, Brad M, and Terrance Odean, 2001, Boys will be boys: Gender, overconfidence, and common stock investment, The Quarterly Journal of Economics 116, 261–292.

Barberis, Nicholas, 2012, A model of casino gambling, Management Science 58, 35–51.

Barberis, Nicholas, and Ming Huang, 2008, Stocks as lotteries: The implications of probability weighting for security prices, American Economic Review 98, 2066–2100.

Barkan, Rachel, and Jerome R. Busemeyer, 1999, Changing plans: Dynamic inconsistency and the effect of experience on the reference point, Psychonomic Bulletin & Review 6, 547–554.

Benartzi, Shlomo, and Richard H Thaler, 1995, Myopic loss aversion and the equity premium puzzle, The Quarterly Journal of Economics 110, 73–92.

Camerer, Colin F., and Teck-Hua Ho, 1994, Violations of the betweenness axiom and nonlinearity in probability, Journal of Risk and Uncertainty 8, 1.

DellaVigna, S., and U. Malmendier, 2006, Paying not to go to the gym, American Economic Review 96, 694–719.

Dertwinkel-Kalt, Markus, and Mats Köster, forthcoming, Salience and skewness preferences, Journal of the European Economic Association .

Dixit, Avinash K, John JF Sherrerd, et al., 1990, Optimization in economic theory (Oxford University Press on Demand).

Ebert, Sebastian, and Philipp Strack, 2015, Until the bitter end: On prospect theory in a dynamic context, American Economic Review 105, 1618–1633.

Eraker, Bjørn, and Mark Ready, 2015, Do investors overpay for stocks with lottery-like payoffs? An examination of the returns of OTC stocks, Journal of Financial Economics 115, 486–504.

Fischbacher, Urs, Gerson Hoffmann, and Simeon Schudy, 2017, The causal effect of stop-loss and take-profit orders on the disposition effect, Review of Financial Studies 30, 2111–2129.

Frederick, S., G. Loewenstein, and T. O'Donoghue, 2002, Time discounting and time preference: A critical review, Journal of Economic Literature 40, 351–401.

Gneezy, Uri, and Jan Potters, 1997, An experiment on risk taking and evaluation periods, The Quarterly Journal of Economics 112, 631–645.

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, 38–54.

Heimer, R., and A. Imas, 2019, Doing less with more, Mimeo .

Hogarth, R., and E. Soyer, 2015, Communicating forecasts: The simplicity of simulated experience, Journal of Business Research 68, 1800–1809.

Imas, Alex, 2016, The realization effect: Risk-taking after realized versus paper losses, American Economic Review 106, 2086–2109.

Kahneman, Daniel, and Amos Tversky, 1979, Prospect theory: An analysis of decision under risk, Econometrica 47, 263–91.

Kaufmann, C., M. Weber, and E. Haisley, 2013, The role of experience sampling and graphical displays on one's investment risk appetite, Management Science 59, 323–340.

King, Michael R, and Dagfinn Rime, 2010, The $4 trillion question: what explains FX growth since the 2007 survey?, BIS Quarterly Review .

Linnainmaa, Juhani T, 2010, Do limit orders alter inferences about investor performance and behavior?, The Journal of Finance 65, 1473–1506.

Maynard, Douglas C., and Milton D. Hakel, 1997, Effects of objective and subjective task complexity on performance, Human Performance 10, 303–330.

Odean, Terrance, 1998, Are investors reluctant to realize their losses?, The Journal of Finance 53, 1775–1798.

O'Donoghue, Ted, and Matthew Rabin, 1999, Doing it now or later, American Economic Review 89, 103–124.

Peköz, Erol A., 2002, Samuelson's fallacy of large numbers and optional stopping, Journal of Risk and Insurance 69, 1–7.

Polkovnichenko, Valery, 2005, Household portfolio diversification: A case for Rank-Dependent Preferences, Review of Financial Studies 18, 1467–1502.

Quiggin, John, 1982, A theory of anticipated utility, Journal of Economic Behavior & Organization 3, 225–243.

Samuelson, Paul A., 1963, Risk and uncertainty: A fallacy of large numbers, Scientia .

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

Yaari, Menahem E., 1987, The dual theory of choice under risk, Econometrica 55, 95–115.

Appendix A. Theoretical Predictions

In this section, we formally derive the predictions of Cumulative Prospect Theory (CPT), Rank Dependent Utility (RDU), and Expected Utility Theory (EUT) on dynamic inconsistency in sequential risk-taking.[16]

Cumulative Prospect Theory

Our analysis of CPT largely follows Barberis (2012). A decision-maker considers the following gamble $L = (p_{-m}, x_{-m}; ..; p_{-1}, x_{-1}; p_0, x_0; p_1, x_1; ...p_n, x_n)$, where p_i corresponds to the likelihood of attaining outcome x_i. Outcomes are ordered such that $-m... - 1$ correspond to those below the reference point x_0, here assumed to be the status quo, and $1...n$ correspond to those above the reference point. We follow Tversky and Kahneman (1992) in assuming that utility derived from this gamble can be represented by:

$$V(L) = \sum_{-m}^{n} \pi_i^{CPT} v(x_i), \tag{A1}$$

where

$$\pi_i^{CPT} = \begin{cases} w(p_i + \cdots + p_n) - w(p_{i+1} + \cdots + p_n) & \text{for } 0 \leq i \leq n, \\ w(p_{-m} + \cdots + p_i) - w(p_{-m} + \cdots + p_{i-1}) & \text{for } -m \leq i < 0, \end{cases} \tag{A2}$$

We also follow Tversky and Kahneman (1992) in assuming the following form for the probability weighting function:

$$w(p) = \frac{p^\delta}{(p^\delta + (1-p)^\delta)^{1/\delta}} \tag{A3}$$

[16] Note that CPT is a special case of the generalized formulation of RDU as proposed by Quiggin (1982) and Yaari (1987). Here, we derive predictions of RDU without the assumption of reference dependence or loss aversion using the formulation employed in Polkovnichenko (2005) .

and value function

$$v(x) = \begin{cases} x^{\alpha} & \text{for } x \geq 0, \\ -\lambda(-x)^{\alpha} & \text{for } x < 0 \end{cases} \tag{A4}$$

where $\alpha, \delta \in (0, 1)$ and $\lambda \geq 1$.

In our setting, $L = (1/2, -10; 1/2, 10)$. The decision-maker faces a sequence of choices. In each round $t \in \{0, ..., 26\}$, she can accept or reject the gamble. If she rejects the gamble, no more gambles are offered. If she accepts it, the outcome is revealed and the decision-maker is offered the same choice again (up through round 26, when the sequence ends). When evaluating this choice problem, the decision-maker chooses a plan s from the set of available plans $S_{t,j}$ in round t and outcome node j. For a given (t, j), the subscript $j \in \{1, t+1\}$ corresponds to the distance of the outcome node (t, j) from the top node of a column in a binomial tree of all potential outcomes that could have occurred by that round t. For example, $S_{1,2}$ corresponds to the set of available plans available after the decision-maker accepted the first gamble and lost. Each plan $s \in S_{t,j}$ is a mapping from every potential outcome of the sequence of gambles from round t onward to actions $a \in \{continue, exit\}$.

Each s generates a random variable $\tilde{G}_s$, which corresponds to the accumulated gains or losses conditional on s being carried out. For example, take $s \in S_{0,1}$ where the decision maker plans to accept the first gamble, continue if she wins in $t = 1$ and then exits regardless of the next outcome in $t = 2$, and exiting in $t = 1$ if she loses the first gamble. This plan corresponds to $\tilde{G}_s \sim (1/2, -10; 1/4, 0; 1/4, 20)$. She chooses plan s which maximizes utility, $max_{s \in S_{t,j}} V(\tilde{G}_s)$. Absent a commitment device, in each round t the decision-maker evaluates the choice problem and re-optimizes given her set of available plans.

Non-linear probability weighting makes it difficult to solve the problem analytically; there is no known analytical solution for a dynamic setting with an arbitrary T. We follow Barberis (2012) in solving the decision problem numerically.

We run simulations to determine the ex-ante optimal plan (in $t = 0$) and the ex-post behavior (in $t > 0$) of each agent. An agent is defined by a unique parameter combination of probability weighting (δ), diminishing sensitivity of the value function (α), and loss aversion (λ). Following the design of the commitment device described in Section II, we define an ex-ante strategy as a combination of a loss limit and a gain limit. For each agent we simulate 10,000 independent paths, each consisting of 26 iid draws from a fair symmetric gamble. A strategy transforms the simulated paths into an outcome distribution.

The optimal plan s^* for each agent is the one that is connected to the outcome distribution with the highest expected value among all possible strategies, as given by the objective function in Equation A1. The agent accepts the sequential gamble if the expected value of the optimal strategy is higher than the value of exiting, which is normalized to zero. If the agent accepts the gamble in the first round, she revisits her decision in every subsequent round. For this purpose, the agent compares the expected value of continuing to accept the gamble, assuming that she will adhere to the ex-ante optimal strategy, with the value of exiting. The value of continuing to gamble is determined by running 10,000 new simulations to determine the updated outcome distribution. The value of exiting is given by the value of the accumulated gains or losses since the beginning. We assume that the reference point is the initial endowment of \$2.6, hence the agent does not update the reference point until the final period when the outcome is paid out.[17] This assumption is consistent with prior experimental evidence (see Imas (2016)).

Results (CPT)

Figure 2, Panels A, C, and E present the findings on the ex-ante optimal plan in $t = 0$. Though the figures illustrate results across a broad range of parameters, we focus on the representative agent with $\alpha = .88$, $\delta = .65$ and $\lambda = 2.25$, as estimated in Tversky and Kahneman (1992). Note that in the presence of dynamic inconsistency,

[17] We present results when this assumption is relaxed below.

the ex-ante decision depends not only on the preference parameters but also on agents's sophistication and the availability of commitment devices. In the following, we analyze the ex-ante decisions of naïve agents, who erroneously believe that they will stick to their ex-ante optimal strategy, as well as sophisticated agents who have a commitment device at their disposal to make sure that they will stick to their ex-ante optimal strategy. Later in this section we discuss how the ex-ante decisions of these two types of agents differ from the ex-ante decision of sophisticated agents without commitment devices. Several findings are obtained.

[Insert Figure 2 about here]

First, the agent would accept the first gamble in a sequence with endogenous exit even though she would reject the same gamble in isolation.

Second, the combination of non-linear probability weighting, diminishing sensitivity and loss aversion determines the ex-ante optimal plan s^*. The optimal plan can be classified as a "loss-exit" strategy for the representative agent. This plan is also optimal for agents with moderate non-linear probability weighting and diminishing sensitivity. In contrast, agents who weigh probabilities close to linearly and have high levels of diminishing sensitivity ($\alpha << 1$), the optimal strategy is a "gain-exit" strategy. Loss aversion plays a straightforward role in determining the proportion of agents who accept the first gamble as part of an optimal plan, as opposed to not entering in the first place. [18]

Figure 2, Panels B, D, and F present the findings on the ex-post behavior. First, agents who accepted the first gamble as part of a "loss-exit" plan deviate from this strategy. Notably, this includes the representative agent. Instead of exiting after initial losses, they

[18] Note that this result and the estimated parameter combinations that accept the sequential gamble largely overlap with the findings of Barberis (2012) even though we restrict the choice set of ex-ante strategies to include only strategies that can be expressed as a combination of two limits, whereas Barberis (2012) optimizes over all possible strategies that can be expressed using a binomial tree. This indicates that our experimental design choice to simplify the plan elicitation is not very restrictive.

end up chasing the losses further by accepting subsequent gambles. These agents exit too early after experiencing small gains (relative to their strategy). Note that while ex-ante strategies are largely determined by the extent of probability weighting, the deviation in ex-post behavior is driven by diminishing sensitivity, i.e. the extent to which agents are risk-averse in the gain domain and risk-seeking in the loss domain. Critically, agents who accept risk as part of a "loss-exit" strategy end up with an outcome distribution that has a *lower* value than rejecting the first gamble in $t = 0$. This represents substantial dynamic inconsistency between planned and actual behavior in this setting.[19]

The predicted dynamic inconsistency generates predictions on initial choices as a function of sophistication and the availability of opportunities to commit to a plan. Agents with "loss-exit" optimal plans but who are aware that they will deviate also understand that the choice to enter yields less utility in expectation than rejecting the first gamble. These agents will only accept the first gamble if offered the opportunity to commit to their optimal plan. Agents who are naive about their dynamic inconsistency will accept the first gamble regardless of commitment opportunities. If the proportion of sophisticated agents is high enough, this leads to the prediction that a greater number of participants will begin gambling when offered an opportunity to commit to a loss or gain limit than when no such opportunities are available.

It is important to highlight that our assumption that the agent does not update the reference point until the final round is critical for predictions on ex-post behavior. An alternative assumption where the reference point updates after every round would not predict an asymmetrical response after accumulated gains and losses. This is because the agent would be in a similar situation as in $t = 0$ in every round, but with a fewer number of prospective rounds. Due to the lower number of prospective rounds, the expected value

[19] They also exit after breaking even if the remaining number of rounds is below an agent-specific minimum required number of rounds to enter the gamble. This type of deviation does not affect the skewness of the outcome distribution but reduces its standard deviation and increases its kurtosis compared to the distribution of the ex-ante optimal plan. Note that the resulting deviation between the value of the ex-ante and ex-post outcome distributions cannot explain why agents with an ex-ante "loss-exit" strategy would regret accepting the first gamble ex-post.

of a "loss-exit" strategy is lower than it was in the beginning of the sequence. Once the number of rounds falls below the agent-specific number that would prompt her to accept the first gamble, she exits. This leads to the prediction that the agent is just as likely to exit after a loss as after a gain. For instance, for the representative CPT agent, the agent-specific minimum required number of rounds is 26, hence she would exit after the first round independent of the outcome because the number of remaining rounds is too low. In general, the closer an agent is to the white-colored area in Figure 2, Panel A, the higher is her agent-specific number of minimum required periods, hence the sooner the agent would exit the lottery independent of its outcome. As outlined in Section IV, the prediction that subjects exit the lottery independent of their gains and losses is not borne out in the data.

Rank-Dependent Utility

Rank-Dependent Utility (RDU) was introduced by Quiggin (1982) and Yaari (1987). We follow Polkovnichenko (2005) and assume the following functional form:

$$\sum_{-m}^{n} \pi_i^{RDU} u(W + x_i), \tag{A5}$$

where W denotes the initial wealth before the first round, $u(.)$ is a power utility function of the form

$$u(W+x) = \begin{cases} \dfrac{(W+x)^{1-\gamma}}{1-\gamma} & \text{for } \gamma \geq 0 \ \& \ \gamma \neq 1, \\ ln(W+x) & \text{for } \gamma = 1, \end{cases} \tag{A6}$$

and

$$\pi_i^{RDU} = w(p_{-m} + \cdots + p_i) - w(p_{-m} + \cdots + p_{i-1}) \tag{A7}$$

For consistency, we assume the probability weighting function w is the same as in (3).

The simulations of the RDU ex-ante optimal plans and ex-post behavior are conducted in a similar way to those for CPT. In contrast to CPT, RDU requires an additional assumption regarding the agents' wealth. Linking this analysis to our experiment, we assume that the agent's wealth comprises their initial endowment of \$2.6. To illustrate how the model's predictions depend on the wealth assumption, we also run the simulations assuming that subjects have additional wealth of \$1,000.

Results (RDU)

Figure 3, Panels A and C present the findings on the ex-ante behavior for an RDU agent. Two main results are obtained. First, as in the case of CPT, some agents would accept the sequential fair gamble with endogenous exit even though they would not accept a single play of the gamble in isolation. Note that this result depends critically on the wealth assumption. Assuming a wealth level of \$2.6, only agents with low levels of risk aversion and strongly non-linear probability weighting would accept the gamble for a single round. In contrast, assuming a wealth level of \$1,000 prior to the experiment leads to the prediction that all agents in Figure 3 with $\delta < 0.9$ would also accept the gamble in isolation. Second, all agents who accept the first gamble do so as part of a "loss-exit" strategy.

[Insert Figure 3 about here]

Figure 3, Panels B and D present the findings on the ex-post behavior for an RDU agent. The main result is that while RDU does predict deviations from the "loss-exit" strategy for some agents, the deviations are always symmetric in response to gains and losses. This is in contrast to the results from CPT which predict asymmetric deviations, such that the agent is more likely to stop after a gain than after a loss.

There are two main types of deviations under RDU, which we refer to as 'type 1' and 'type -1'. In the 'type 1' deviation, some decision-makers who accept the first

gamble deviate from their "loss-exit" strategy to continue investing until the final round independent of the gamble outcomes. From an ex-ante perspective, the utility of this type of deviation is lower than from the "loss-exit" strategy for some agents (marked red in Figure 3). However, for a significant proportion of agents, the 'type 1' deviation still generates an outcome distribution which yields greater utility than rejecting the first gamble. This includes agents with preference parameters in the region of estimates from prior studies (e.g., Tversky and Kahneman (1992) and Camerer and Ho (1994)).[20] 'Type -1' deviations comprise all other types of deviations, such as the agent exits when the initial wealth is reached and the number of remaining rounds is sufficiently low.

Expected Utility Theory

Expected Utility Theory (EUT) is a special case of RDU for $\delta = 1$. The utility function for $\delta = 1$ is characterized by constant relative and decreasing absolute risk aversion with skewness preferences (see Arditti (1967)). Predictions for both ex-ante plans and ex-post behavior are illustrated in Figure 3. It is clear that EUT with skewness preferences does not predict that participants will accept the first gamble in a sequence while rejecting the gamble in isolation. It also follows trivially that EUT does not predict any dynamic inconsistency.[21]

[20] This result is dependent on the assumption about initial wealth. Assuming initial wealth \$1,000, all agents in Figure 3 who accept the initial gamble will continue until the end, independent of the outcome.

[21] For the case of gambles with positive expected value, Peköz (2002) shows that skewness preferences in combination with endogenous exit can explain the Samuelson paradox (Samuelson (1963)) as subjects will follow a "loss-exit" strategy to generate positive skewness. This is not the case for the fair gamble which is used in our experiments.

Appendix B. List of Variables

Variable	Description
Age	Demographic variable elicited in entry-level questionnaire elicited in entry-level questionnaire.
Male	Dummy variable that equals one for male.
Statistical skills	Categorical variable elicited as follows: "How would you rate your statistical knowledge? Please choose a category between 1 ("very bad") and 6 ("very good")."
Study field	Categorical variable elicited as follows: "Your field of study?"
Highest education	Categorical variable elicited as follows: "Your highest level of education: 1 (below high school), 2 (high school), 3 (college), 4 (bachelor), 5 (master), 6 (PhD or above)".
Entry decision	Dummy variables which equals one if the subject accepts the gamble before the first round.
D^{seq}	Dummy variable that equals one for sequential treatment
$D^{hardplan}$	Dummy variable that equals one for hard plan treatment
$D^{softplan}$	Dummy variable that equals one for soft plan treatment
D^{1gain}	Dummy variable that equals one for drawing a gain in round 1 conditional on accepting the gamble.
D^{gain}	Dummy variable that equals one for realizing a cumulative gain and zero for realizing a cum loss.
Complexity	Equally-weighted average response to four questions regarding the perceived complexity of the main task following Maynard and Hakel (1997): "I found this to be a complex task"; "This task was mentally demanding"; "This task required a lot of thought and problem-solving"; "I found this to be a challenging task". The responses are elicited on a Likert-type scale from 1 "totally disagree" to 7 "totally agree". The perceived complexity is elicited after the main task.

Appendix C. Experimental Instructions

Appendix CI. One-shot Treatment

`Screen: Instructions`

You have 10 cents. You can choose to invest 10 cents in the following lottery or to keep it:

With a chance of 1/2 (50%) the lottery will "succeed" and you will earn an additional 10 cents, for a total of 20 cents. With a chance of 1/2 (50%) the lottery will "fail" and you will lose the 10 cents you invested.

Click "Start" to make several random draws from the distribution of the lottery: [The subject is required to make 10 draws from an individual stratified sample before proceeding to the next screen.]

`Screen: Investment decision`

You can now choose whether or not to invest 10 cents in the lottery.

Do you want to invest? [Yes]/[No]

Appendix CII. Sequential Treatment

`Screen: Instructions (1)`

You can choose to invest 10 cents in the following lottery or to keep it:

With a chance of 1/2 (50%) the lottery will "succeed" and you will earn an additional 10 cents, for a total of 20 cents. With a chance of 1/2 (50%) the lottery will "fail" and you will lose the 10 cents you invested.

Click "Start" to make several random draws from the distribution of the lottery: [The subject is required to make 10 draws from an individual stratified sample before proceeding to the next screen.]

`Screen: Instructions (2)`

The experiment consists of 26 successive rounds. You have 260 cents in total to invest with. You can invest 10 cents per round in the lottery for up to 26 rounds. At the beginning you will choose whether or not to invest in the first round. After learning the outcome of your investment (whether you won or lost), you will choose whether to invest again or nor. You can stop investing at any time. Once you decide to stop investing, this part of the experiment will end.

Your earnings for this part of the experiment are as follows: At the end, we will count the number of rounds you have won (n) and the number of rounds you have lost (m). Your total gain or loss is given by the difference between these numbers multiplied by your investment per round which is 10 cents.

- If n $\geq$ m, you have earned a total gain of (n - m)$\times$ 10 cents. In this case, you will receive your initial endowment plus the amount of your total gain.
- If n $<$ m, you have endured a total loss of (m - n)$\times$ 10 cents. In this case, you will

receive the rest of your initial endowment after deducting the amount of your total loss.

Please click "Next" to proceed with the first round.

`Screen: Round 1 of 26: Investment decision`

You can now choose whether or not to invest 10 cents in the lottery in round 1.

Do you want to invest? [Yes]/[No]

dots

`Screen: Round X of 26: Result`

[This screen is displayed conditional on investing in this round.]

In round X you have earned a gain/endured a loss of [...] cents. [The outcome is colored in red or green for loss or gain, respectively.]

In the first X rounds you have earned a total gain/endured a total loss of [...] cents. [The outcome is colored in red or green for loss or gain, respectively.]

Please click "Next" to proceed to the next round.

Appendix CIII. Hard and Soft Plan Treatments

`Screen: Instructions (1)`

You can choose to invest 10 cents in the following lottery or to keep it:

With a chance of 1/2 (50%) the lottery will "succeed" and you will earn an additional 10 cents, for a total of 20 cents. With a chance of 1/2 (50%) the lottery will "fail" and you will lose the 10 cents you invested.

Click "Start" to make several random draws from the distribution of the lottery: [The subject is required to make 10 draws from an individual stratified sample before proceeding to the next screen.]

`Screen: Before we move on...`

Please think of two arbitrary numbers (integers) between 0 and 260. [Note: Your earnings do not depend on your responses to this question.]

My first number:...

My second number:...

`Screen: Instructions (2)`

You have 260 cents in total to invest with. You will choose whether or not to invest 10 cents in the lottery over a series of up to 26 rounds. But first we ask you to indicate what is the maximum amount of losses or gains you would be willing to take before stopping. These are your loss limit and your gain limit. If you choose to start investing 10 cents, you will keep investing 10 cents in each subsequent round until your total gain or loss reaches your gain limit or loss limit respectively.

- You can think of the **loss limit** as the most of your endowment that you are willing

to lose.

- You can think of your **gain limit** as the amount of gains you would be happy to walk away with, without having to risk any more.

Example: At the beginning of the experiment we asked you for two arbitrary numbers and you gave us the numbers [y_1] and [y_2]. Let us assume, your loss limit is [y_1] cents and your gain limit is [y_2] cents. After every round, we will count the number of rounds you won and the number of rounds you lost so far to determine your total gain or loss. Your loss limit is reached if your total loss reaches [y_1] cents. In other words, it is reached as soon as you have lost [$y_1/10$] rounds more often than won. Your gain limit is reached if your total gain reaches [y_2] cents. In other words, it is reached as soon as you have won [$y_2/10$] rounds more often than lost.

Please note, that there is no guarantee that your gain limit or your loss limit will be reached during the course of the experiment as the lottery outcomes are completely random and independent.

Your earnings for this part of the experiment are as follows: At the end, we will count the number of rounds you have won (i.e., n) and the number of rounds you have lost (i.e., m) before you stopped investing.

- If you have earned a total gain (if n > m), you will receive your initial endowment of 260 cents plus the amount of your total gain of (n - m)$times$10 cents.
- If you have endured a total loss (if m > n), you will receive the rest of your initial endowment after deducting the amount of your total loss of (m - n)$\times$10 cents.

Please indicate your loss limit and your gain limit. [Note: You will choose whether or not to start investing afterwards.]

Loss limit (in cents):...

Gain limit (in cents):...

`Screen: Instructions (3)`

[This screen is displayed conditional on being in the **soft plan treatment**.]

Your gain and loss limits are not binding and will not be enforced if you start investing. This means that we will inform you immediately if either your gain limit of [...] cents or your loss limit of [...] cents is reached before the final round. You will then choose whether to continue or stop investing.

This part of the experiment ends if:

- you decide not to invest in the first round (see next page), or
- you start investing and you decide to stop investing after being informed that one of your limits is triggered, or
- you reach the final 26th round.

Click "Next" to proceed.

`Screen: Instructions (3)`

[This screen is displayed conditional on being in the **hard plan treatment**.]

Your gain and loss limits are binding and will be enforced if you start investing. This means that you will automatically stop investing if either your gain limit of [...] cents or your loss limit of [...] cents is reached before the final round.

This part of the experiment ends if:

- you decide not to invest in the first round (see next page), or

- you start investing and either your gain limit or your loss limit has been reached or exceeded, or
- you reach the final 26th round.

`Screen: Round 1 of 26: Investment decision`

[This screen is displayed conditional on being in the **soft plan treatment**.]

You can now choose whether or not to start investing 10 cents in the lottery over a series of up to 26 rounds.

Do you want to start investing? [Yes]/[No]

`Screen: Investment decision`

[This screen is displayed conditional on being in the **hard plan treatment**.]

You can now choose whether or not to start investing 10 cents in the lottery over a series of up to 26 rounds until you automatically stop.

Do you want to start investing? [Yes]/[No]

`Screen: Round X of 26: Result`

[This screen is displayed conditional on being in the **soft plan treatment**.]

Your gain limit/loss limit was triggered in round X.

For the first X rounds you have earned a total gain/endured a total loss of [...] cents. [The outcome is colored in red or green for loss or gain, respectively.]

`Screen: Round X+1 of 26: Investment decision`

[This screen is displayed conditional on being in the **soft plan treatment**.]

You can now choose whether or not to continue investing 10 cents in the lottery. In case you continue, we will inform you as soon as either one of your limits is triggered again or the final round is reached.

Do you want to continue investing? [Yes]/[No]