

Discussion Paper Series – CRC TR 224

Discussion Paper No. 274
Project C 01

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

Rawley Heimer¹
Zwetelina Iliewa²
Alex Imas³
Martin Weber⁴

March 2021

¹ Carroll School of Management, Boston College (e-mail: heimer@bc.edu)

² University of Bonn (e-mail: zwetelina.iliewa@uni-bonn.de)

³ Booth School of Business, University of Chicago (e-mail: aleximas@chicagobooth.edu)

⁴ Department of Finance, University of Mannheim, and the Centre for European Policy Research
(e-mail: weber@bank.bwl.uni-mannheim.de)

Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)
through CRC TR 224 is gratefully acknowledged.

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

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

This draft: March 1, 2021

ABSTRACT

We document a robust dynamic inconsistency in risky choice. Using a unique brokerage dataset and two preregistered experiments, we compare people’s initial risk-taking plans to their subsequent decisions. In both settings, people accept risk as part of a “loss-exit” strategy—planning to continue taking risk after gains and stopping after losses. Actual behavior follows the *reverse* pattern, deviating from initial strategies by cutting gains early and chasing losses. More individuals accept risk when offered a commitment to their initial strategy. Our results help reconcile seemingly contradictory findings on risk-taking in static versus dynamic contexts. We discuss implications for theory and welfare.

* Rawley Heimer is at the Carroll School of Management, Boston College (e-mail: heimer@bc.edu). Zwetelina Iliewa is at the University of Bonn (e-mail: zwetelina.iliewa@uni-bonn.de). Alex Imas is at the Booth School of Business, University of Chicago (e-mail: aleximas@chicagobooth.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 Hunt Allcott, Sandro Ambuehl, Nick Barberis, Erik Eyster, Cary Frydman, Daniel Gottlieb, Johannes Haushofer, Peter Kelly, Camelia Kuhnen, John List, Pietro Ortoleva, Collin Raymond, Alex Rees-Jones, Hersh Shefrin, Kelly Shue, Dmitry Taubinsky, and Florian Zimmermann for helpful suggestions and illuminating discussion. We also thank seminar participants at Boston College, Middlebury College, the Radboud University Nijmegen, Simon Fraser University, the University of Colorado, Boulder, University of Bonn, and the University of Washington, as well as participants at the Annual Meeting of the Society for Experimental Finance, the Finance in the Cloud Conference, the Workshop on Behavioral and Experimental Economics, the NBER Behavioral Finance Conference, SFS Cavalcade North America, and the EFA Annual Meeting. Pascal Kieren and Yanghua Shi provided excellent research assistance. Zwetelina Iliewa would like to acknowledge that this project was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy - EXC 2126/1- 390838866 and through CRC / TR 224 (Project C01). All remaining errors are our own.

I. Introduction

People are often confronted with risky decisions in dynamic environments, such as whether to purchase a stock or take out a loan. A crucial feature of such environments is the option to revise one’s initial decisions and reevaluate choices over time. For example, an investor observing her stock’s performance can choose to hold the position, purchase more shares, or sell them. A borrower can repay her loan, roll it over, or borrow more. However, people’s behavior in such dynamic environments appears to run counter to their decisions in one-shot, static settings. People tend to avoid taking risk when it is presented in isolation, overwhelmingly rejecting even positive-expected-value bets and displaying a seemingly anomalous preference for safer options (Kahneman and Tversky, 1979). These findings are difficult to reconcile with those from dynamic environments, where individuals appear risk-seeking for bets that are offered as part of a sequence, even if the cumulative risk has a zero or *negative* expected value.¹ For example, retail investors in financial markets tend to trade too much, exhibiting excessive trading volume even in settings with no risk premium or negative expected returns (see e.g., Barber and Odean, 2000; Heimer and Simsek, 2019; Liu et al., 2020).²

Unlike the case of static settings, peoples’ behavior in dynamic settings is determined both by their *planned choices* and their *actual choices* in response to experiencing gains and losses. Planned choices correspond to a fully contingent strategy characterizing how much risk to take in response to the resolution of prior decisions. For example, a person may purchase a risky asset with the intention of keeping it if the price goes up and selling it if the price goes down. The choice to take on risk in the first place—what we term the ‘entry decision’—depends on whether the expected outcomes of a given plan are more attractive than a safer alternative. The person’s actual choices in response to losses and gains may follow her plan—or they may not. Machina (1989) studies this question formally in his influential survey of dynamic choice

¹ Note that this discrepancy between one-shot and sequential risk-taking is conceptually 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 dynamic versus one-shot environments. This is because i) in dynamic environments feedback is provided after every choice and ii) an MLA agent would reject fair or negative expected-value risk regardless of feedback frequency (Langer and Weber, 2008).

² For similar evidence from the lab, see Andrade and Iyer (2009) and Imas (2016).

under uncertainty. He notes that while the standard expected utility framework predicts that planned and actual choices will be the same, dynamic inconsistency can potentially emerge from many models of non-expected utility. He argues that dynamic inconsistency is undesirable from a normative perspective, and proposes that in order to eliminate it, theory should amend the assumption that prior decisions do not affect future considerations (i.e. consequentialism).

This paper presents an in-depth empirical investigation of why and how people take risk in dynamic settings, examining what strategies motivate them to start taking on risk in the first place, whether their actual choices deviate from their plans, and if so, what these deviations look like. Using evidence from both the lab and field, we identify a robust dynamic inconsistency in decision-making under uncertainty. People are initially motivated to take risk as part of a “loss-exit” strategy, which involves continuing to take on risk after gains and to stop earlier after losses. Notably, this strategy generates a more positively-skewed outcome distribution than any of the individual gambles in isolation. In contrast, people’s actual behavior follows the *reverse* pattern: they cut their gains early and continue to chase losses. This “gain-exit” strategy represents a systematic deviation from people’s own stated plans. Moreover, they appear at least partially aware of their dynamic inconsistency and are more likely to begin taking on risk if it is possible to commit to a strategy.

Interpreting the behavioral data through the lens of theory helps rationalize the discrepancy in risk-taking between static and dynamic environments. We find that people are indeed more likely to accept risk when it is presented as part of a dynamic sequence of choices than if the same gamble is offered in isolation. We run a horse race between different models of decision-making under risk and show that a dynamic framework featuring probability weighting, reference dependence, and diminishing sensitivity such as Cumulative Prospect Theory is most consistent with the observed behavioral patterns (Barberis, 2012; Tversky and Kahneman, 1992). Finally, we apply the techniques outlined in Bernheim and Taubinsky (2018) for behavioral welfare analysis to shed light on the welfare consequences of dynamic inconsistency in our setting. Our results highlight the fact that, although dynamic inconsistency in the domain of risk may not be normatively appealing (Machina, 1989), it is an important feature of the data. Although much of the theoretical literature has focused on understanding the conditions under which dynamic consistency is preserved (e.g., Volij, 1994; Karni and Schmeidler, 1991), our paper

offers a starting point for understanding the drivers of dynamic inconsistency in practice.

We begin our investigation by examining planned versus actual behavior in a naturalistic setting, using a dataset from a large online brokerage with approximately 190,000 traders from over 150 countries. The unique feature of this dataset is that the brokerage mandates that traders submit ex-ante strategies for every position that they open. When purchasing an asset, traders are required to submit an exit strategy after gains (take-profit order) and after losses (stop-loss order). Take-profit and stop-loss orders correspond to limits on how much a trader is willing to gain or lose, respectively, before exiting the position. Importantly, the dataset also tracks all subsequent revisions to these limits until the position is closed, as well as whether positions are manually closed before triggering a gain or loss limit. The combination of initial limits, subsequent revisions, and manual exits allows us to characterize the traders' ex-ante risk-taking strategies and compare these strategies to actual behavior in response to gains and losses. For example, by placing a loss limit of 10% and a gain limit of 20%, the trader opens the position with a risk-taking strategy that pairs a willingness to lose 10% for the chance of gaining 20%. The trader can revise this strategy by changing one of the limits after seeing gains and losses, e.g., moving the loss limit to 20%, or by manually closing the position before the limits are hit, e.g., selling the asset after a 5% gain.

We document a significant discrepancy between people's planned and actual behavior. The majority of traders' ex-ante strategies can be classified as "loss-exit" plans, in which the average loss limit is smaller than the corresponding gain limit. This implies that traders open new positions with the intention of exiting after smaller losses relative to gains. Note that this strategy seems to contradict the well-known disposition effect whereby traders hold losers longer than winners (Shefrin and Statman, 1985). However, traders' subsequent choices follow a pattern *opposite* of their intended plans. After experiencing losses, investors revise their loss limits to allow the price to decrease further. For example, when the position is currently at a 5% loss, a trader modifies her loss limit by pushing it away from 10% to 20%. On the other hand, when experiencing gains, investors are most likely to manually exit the position before the gain limit is triggered. For example, when the trader has set a 20% gain limit and the position hits a 10% gain, she will manually close it at 10% before the gain limit is reached. This pattern demonstrates a discrepancy between investors' "loss-exit" plans when opening a position and

their subsequent behavior, which is instead consistent with a “gain-exit” strategy.

The financial setting is unique because it allows us to compare people’s ex-ante risk-taking strategies to their subsequent decisions in an environment with significant stakes and frequent feedback. Our findings suggest a dynamic inconsistency in risk-taking behavior. To facilitate identification and help pin down the mechanism, we designed an experimental paradigm that generates data rich enough to isolate dynamic inconsistency and interpret it through the lens of theory. As our Appendix A formally outlines, this requires the experiment to have the following features: (i) the ability to elicit incentivized ex-ante strategies and compare them to ex-post behavior, (ii) elicit initial choices to begin taking risk—‘entry’ decisions—as a function of the number of rounds and availability of commitment opportunities, and (iii) a long enough sequence of gambles such that strategies can significantly affect skew over final outcomes compared to the one-shot gamble.

In two pre-registered experiments ($N = 940$), participants are offered the choice to accept or reject a sequence of *fair* symmetric gambles. They are provided feedback after every decision and have the choice to stop anytime. We use a mixture of between-subject and within-subject designs 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 rounds and whether we elicit their strategies prior to the initial choice. Similar to the field setting, strategies are elicited in the form of loss (gain) limits, which correspond to the most a participant is willing to lose (gain) before refraining from taking on further risk. In addition to being intuitive and easy to explain, under mild assumptions, these limits are sufficient to characterize participants’ risk-taking plans. In one treatment, the elicited strategies are binding; we refer to this condition as a “hard plan” because it provides participants with a guarantee that their preferred strategies will be followed. In the “soft plan” treatment, participants are reminded of their initially preferred limits but can deviate from them, similar to the financial setting. Finally, in a separate “sequential” treatment, participants make decisions without stating their ex-ante strategies.

We find that participants are significantly more likely to accept risk when it is part of a larger sequence of gambles than in isolation. This confirms the discrepancy in risk-taking between static and dynamic settings within the same paradigm. More than 80% of participants’ ex-ante strategies can be classified as “loss-exit” plans. In contrast, only 7% of strategies can be classified as “gain-exit.” Strikingly, the average participant initially accepts risk with a gain limit that is 3.81 times higher than her loss limit. In contrast, participants’ actual choices follow the *reverse* “gain-exit” pattern: they are significantly more likely to stop after winning than after losing, replicating the behavioral pattern we observe in the field.

The experimental treatments also allow us to examine whether people are aware of their dynamic inconsistency or not.³ We find that a significant proportion of participants appear sophisticated about their dynamic inconsistency: people are more likely to begin taking on risk when they are provided with a commitment opportunity. In contrast, we find no evidence that participants are sophisticated about the inefficacy of soft commitment opportunities in disciplining their behavior. Though most participants deviate from their strategies in a manner similar to when no commitment is available—which demonstrates dynamic inconsistency within-subject—they are equally likely to take on risk as when commitment is binding. This suggests that participants’ sophistication about the efficacy of different commitment opportunities may be limited.

Taking stock of the entire set of findings, our Appendix A formally shows that they are consistent with the dynamic predictions of models that incorporate probability weighting, diminishing sensitivity, and gain-loss asymmetry such as Cumulative Prospect Theory (CPT)—first proposed by Tversky and Kahneman (1992) and extended to a dynamic setting with finite rounds by Barberis (2012). There, we show that the observed behavioral data cannot be rationalized by frameworks that incorporate only diminishing sensitivity (Expected Utility Theory, EUT thereafter) or probability weighting (Rank Dependent Utility, RDU thereafter) without reference dependence (Quiggin, 1982; Yaari, 1987).⁴ The dynamic version of CPT predicts that the same person will reject a single fair gamble while accepting the same gamble as part of

³ Following the time preference literature, we classify those who are aware of their dynamic inconsistency as ‘sophisticated’ and those who are not as ‘naïve’.

⁴ We also show that our pattern of results is not consistent with models of quasi-hyperbolic discounting (O’Donoghue and Rabin, 1999) or naïveté about reference point updating (Strub and Li, 2020).

a dynamic sequence. The framework also generates both the pattern in planned choices and deviation in actual choices, as well as the demand for commitment that we observe in our data (Barberis, 2012).

Assessing the welfare consequences of dynamic inconsistency in risky choice is not straightforward because both planned and actual decisions appear to be driven by non-normative factors. We follow the methodology outlined in Bernheim and Taubinsky (2018) to analyze welfare in this setting. As a first step, we attempt to characterize the welfare-relevant domain by running a separate study that re-framed decisions between treatments but kept the underlying choice structure constant.⁵ Our results provide support for using ex-ante choices as the welfare-relevant benchmark. We then proceed to use simulations to quantify the welfare consequences of dynamic inconsistency. This analysis suggests that a representative naïve agent who accepts the sequential gamble with the belief that she will stick to her “loss-exit” plan incurs a utility cost equivalent to losing more than *one hundred and ten percent* of the one-shot investment with certainty. Abstracting away from our attempt to assess welfare through the lens of a specific theory, we view this exercise as a more general contribution for evaluating welfare in a normatively-ambiguous domain—where all decisions are potentially driven by psychological frictions.

The theoretical interpretation of our results suggests that the option to stop taking on risk in response to gains and losses is a critical feature of dynamic 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. This has significant implications for interpreting prior findings and policy design, as well as for generating new predictions on the role of commitment. First, our results provide support for a mechanism that links seemingly disparate phenomena—such as differences in risk-taking in static versus dynamic environments and the disposition effect—within a unified framework. While the disposition effect, i.e., the propensity to realize gains earlier than losses, has been one of the most widely studied phenomena in finance (see Kaustia, 2010, for review), its costs are typically quantified in strictly financial terms. Our findings offer direct evidence for

⁵ Chetty, Looney, and Kroft (2009), Taubinsky and Rees-Jones (2018), and Allcott and Taubinsky (2015) use a similar strategy of re-framing decisions to calculate the welfare costs of behavioral frictions in the domains of sales taxes, income taxes, and the market for energy efficient light bulbs, respectively.

the hypothesis introduced by Barberis (2012) that the disposition effect is actually inconsistent with traders’ ex-ante preferences.⁶ As a result, the welfare consequences of the disposition effect potentially extend beyond calculating financial losses.

Moreover, our findings suggest that loss and gain limits—which are prominent and oft-used features in financial markets—may serve the dual purpose of attracting investors through their perceived role as commitment opportunities. However, the vast majority of these limits can be revised ex-post. Such soft commitment is also featured in regulation aimed at limiting the scope for unintended losses. For example, 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. Our experimental results show that the presence of soft commitment opportunities leads a substantial fraction of individuals who would have avoided risk absent commitment opportunities to accept it instead. Such an ‘illusion of commitment’ is costly, as these same individuals end up systematically deviating from their non-binding strategy. This suggests that policy and regulations employing non-binding commitment may potentially backfire by encouraging investors to take on more risk than they otherwise would without effectively preventing them from chasing losses. On the other hand, 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 would 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.

Related Literature We are not the first to examine planned choices in dynamic risk-taking environments. In the lab, Andrade and Iyer (2009), Ploner (2017), and Imas (2016) find that people plan to bet more after a gain than a loss, while Barkan and Busemeyer (2003) find that people want to bet more after a loss than a gain. Dertwinkel-Kalt, Frey, and Köster (2020) study strategies in a dynamic optimal stopping problem through the lens of Saliency Theory (Bordalo, Gennaioli, and Shleifer, 2012). We contribute to this literature by providing evidence for “loss-exit” strategies from the field and comparing planned behavior to ex-post choices within the

⁶ As further evidence of this mechanism, Bernard, Weber, and Loos (2020) show that the disposition effect amongst stock traders increases with skewness of the assets. In the lab, both Nielsen (2019) and Merkle, Müller-Dethard, and Weber (2020) show that loss-chasing is eliminated when risk is negatively skewed.

same incentivized setting. The latter allows us to identify dynamic inconsistency and interpret it through a theoretical lens.⁷

The paper is closely related to the theoretical work on dynamic models of risky choice (Machina, 1989). Barberis (2012) and Ebert and Strack (2015) present the predictions of CPT in dynamic environments, demonstrating that it generates dynamic inconsistency. The former studies a setting with a finite number of rounds in discrete time while the latter considers an infinite-horizon setting. While Barberis (2012) demonstrates that CPT agents exhibit an outcome-dependent dynamic inconsistency—stopping earlier (later) than planned after gains (losses)—Ebert and Strack (2015) predict that CPT agents will take on risk until bankruptcy regardless of whether they see gains or losses. In addition to discussing these predictions further, the Appendix A presents a more in-depth review of the theory literature on dynamic choice under uncertainty.

Our paper also contributes to the literature on dynamic inconsistency between planned and actual behavior more broadly. 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.⁸ In contrast to the significant body of work on time preferences, few papers have explored dynamic inconsistency in choices under risk, and those that do have largely been theoretical (e.g., Barberis, 2012; Ebert and Strack, 2015; Machina, 1989).

⁷ In this prior work, the decision to initially take risk was either forced (Ploner, 2017; Barkan and Busemeyer, 2003) or coerced (Andrade and Iyer, 2009). All studies use four rounds or less, which as discussed in Appendix A, can explain the somewhat contradictory results. Imas (2016) compares planned and actual behavior only in the domain of losses using a within-subject hypothetical planning stage. Dertwinkel-Kalt, Frey, and Köster (2020) also elicit hypothetical, within-subject plans, but do not study dynamic inconsistency since ex-post choices (as opposed to strategies) cannot be revised.

⁸ Unlike a dynamic framework with probability weighting and diminishing sensitivity, models of time inconsistency do not predict outcome-specific deviations between planned and actual behavior. While it may be possible to generate deviations where people gamble for longer than they intended, frameworks such as present-biased preferences (O’Donoghue and Rabin, 1999) would not predict the discrepancy between “loss-exit” strategies and “gain-exit” behavior observed in our data. As demonstrated in Appendix A, present bias predicts that people would not accept risk in our dynamic setting, and even conditional on entry, the majority would not have “loss-exit” plans.

The rest of the paper proceeds as follows. Section II describes the field setting and presents results on risk-taking behavior. Section III describes the experimental design. Section IV presents the results and Section V discusses the potential welfare consequences. Section VI outlines the implications of our findings and concludes. Finally, Appendix A formally compares the predictions of various models of dynamic choice under uncertainty.

II. Dynamic Inconsistency in the Field

We first examine the dynamics of risky decision-making in the field. We use trading data from a large international online brokerage with 187,521 traders from June 2013 until August 2015 (summary statistics are presented in Table I). The data contain traders from all six major continents and over 150 countries. Its broad geographic coverage is novel relative to other individual-investor datasets used in the literature, which tend to be confined to a single country. Similar to other studies of retail traders (e.g., Barber and Odean, 2001), 82% of traders are male.

[INSERT TABLE I ABOUT HERE]

The brokerage offers contracts for difference (CFD), which are derivatives contracts that pay the difference between the open and close price of an instrument and involve no actual receipt of the underlying asset. Traders can open long or short positions in the assets and all transactions are self-initiated (non-advised). Most 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. Leverage is a common feature of these markets because currency prices tend to be much less volatile than other securities, such as individual stocks; hence, leverage is needed to match the risk/return profile of other securities. Prior work, such as Heimer (2016), have shown that traders in these markets exhibit many of the same behavioral patterns as common stock traders that have been studied in the literature (starting with Barber and Odean, 2000). Moreover, the global daily market volume for currency CFDs (FOREX) is large. In past years, the volume has been roughly

equivalent to the entire NYSE family of stock exchanges (King and Rime, 2010).

A unique feature of our dataset relative to brokerage datasets used in prior work is that it allows us to identify traders' ex-ante risk-taking strategies when they open new positions. The brokerage requires all traders to set loss and gain limits (stop-loss and take-profit orders, respectively) for every position that they open.⁹ Each gain (loss) limit corresponds to the most a trader is willing to gain (lose) as part of her ex-ante strategy when buying an asset. For example, the investor may purchase a stock while setting the gain limit at 20% and a loss limit at 10%. Once a limit is hit (e.g., the price declines by 10%), the position is closed automatically at the price specified by the order.

The brokerage also records all of the *revisions* that traders make to the limits after a position is opened. Though traders are required to enter gain and loss limits when they open a position, these limits are not binding — after opening the position, traders can revise them after experiencing gains and losses, and can manually close positions before the limits are hit. These order revisions are at the traders' discretion, are not influenced by the brokerage, and can be revised as often as the trader would like until she closes the position manually or one of the limits is triggered. Once the limit is triggered it cannot be revised.

These features of the dataset allow us to compare ex-ante plans to actual choices in a real-world setting. Because we observe traders' revisions to their limit orders, we can study how experiencing gains versus losses affects the decision to either continue holding a position or close it. Moreover, several features of the setting make it likely that our findings on traders' ex-post decisions are driven by changes in prices—i.e., gains and losses—rather than omitted variables. First, 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. Second, the CFD of currency pairs have been shown to yield negative expected returns for active retail traders (Heimer and Simsek, 2019); in turn, the initial willingness to take on risk in this setting cannot arise from return aggregation as in Benartzi and Thaler (1995)—a motive that requires risk to have a positive expected value.

⁹ Some papers—for example, Linnainmaa (2010), which uses individual investor data from Finland—also study traders' limit orders. However, such papers analyze *buy* limit orders, which are orders to purchase (as opposed to sell) an asset once it hits a specified price.

A. Results

We start by defining a trader’s risk-taking strategy as the combination of the gain and loss limit when she opens a new position. Strategies are classified as “loss-exit” (“gain-exit”) if the loss limit (gain limit) is closer to the purchase price than the gain limit (loss limit). We classify traders’ strategies as “neutral” when the position is opened with loss and gain limits that are the same distance from the purchase price. The brokerage sets the “neutral” strategy as the default.

We find that the majority of traders use “loss-exit” as their modal strategy when deciding to open a new position. Figure 1 presents the proportion of traders whose modal strategy is “loss-exit”, “gain-exit,” or “neutral”. The largest proportion of traders (46%) use “loss-exit” as their modal strategy. Substantially fewer adopt “gain-exit” (32%) and “neutral” (22%) as their modal strategies (all differences significant at the 1% level).

We also show that traders’ preference for “loss-exit” strategies is robust to the characteristics of the trade and to the characteristics of the trader. Table II presents regression analyses where the unit of observation is an individual trade. To capture traders’ active decisions, the regressions include the sample of trades that have been revised away from equidistant limits. The dependent variable equals one if the trade uses a “loss-exit” strategy. The coefficient on the constant term captures the fraction of positions that use a “loss-exit” strategy. We find that significantly more positions are opened as part of a “loss-exit” strategy across all specifications. Panel A illustrates that trader characteristics, such as gender, country-of-origin, and experience, do not moderate this result. Panel B shows that the result is also robust to the characteristics of each individual trade, such as the position’s leverage, direction, capital, and instrument.

[INSERT TABLE II ABOUT HERE]

Next, we examine how traders’ behavior after the position is opened responds to accumulated gains and losses. We define gains and losses relative to the individual asset’s initial purchase price. This is a reasonable assumption because the brokerage displays position-level gains and losses relative to this opening purchase price.

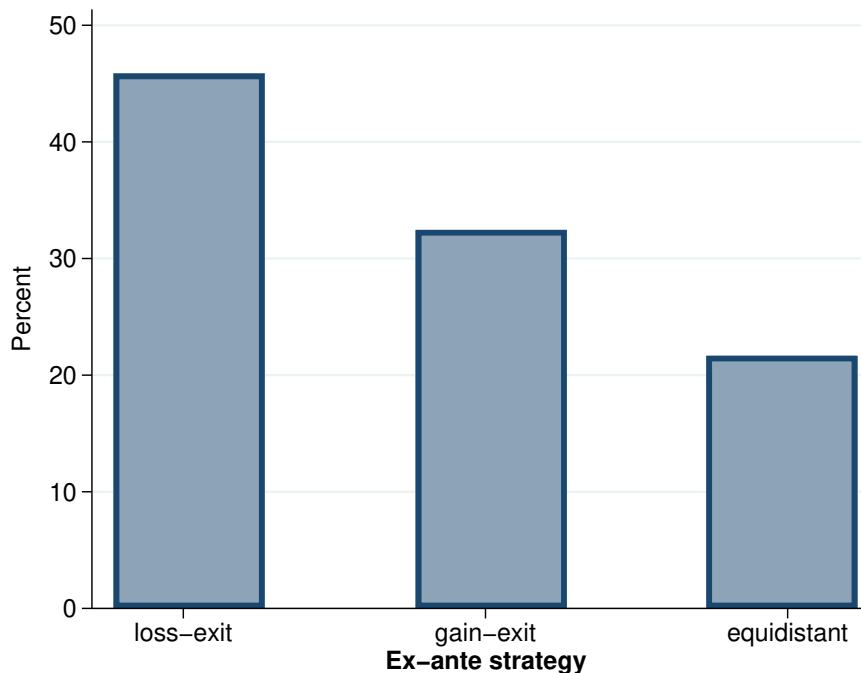


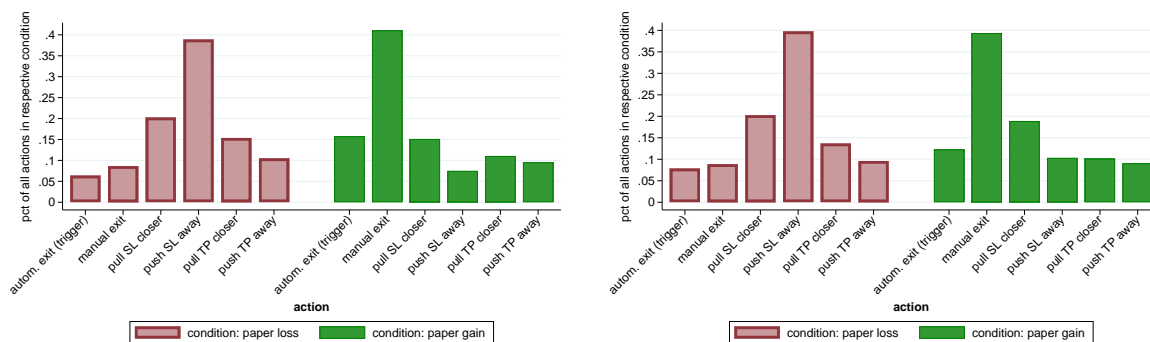
Figure 1. Ex-ante strategies. This figure illustrates the proportion of traders whose modal strategy is “loss-exit”, “gain-exit” or “neutral”. A “loss-exit” strategy corresponds to a loss limit that is closer to the opening price than the contemporaneously submitted gain limit. A “gain-exit” strategy corresponds to a loss limit that is further away from the opening price than the corresponding gain limit. A neutral strategy implies equidistant limits.

Figure 2 presents traders’ behavior in response to paper (i.e., unrealized) gains and losses.¹⁰ Note that traders can respond to gains and losses by allowing their initial gain and loss limits to trigger, closing the positions manually before these limits are hit, or revising the limits towards or away from the current market price. Panel A presents the distribution of actions for all trades on the brokerage, split by whether the position had accumulated a loss or a gain. Panel B presents the distribution of actions only for positions that were opened as part of a “loss-exit” strategy.

We find that the most frequent response to an accumulated loss is to revise the strategy by pushing the loss limit away. For example, a trader who opened the position with the intention of limiting her losses to 10% revises it to 15% in response to accumulating losses. We observe

¹⁰ Gains and losses are evaluated in real-time at a ten-minute frequency using all of the trades on the brokerage to estimate the bid and ask quotes of each underlying asset.

such revisions of loss limits nearly 20 percentage points more often than the next most common choice. This is in contrast to traders’ behavior when they have accumulated gains—traders are most likely to manually close the position before the gain limit is reached. Traders are nearly 25 percentage points more likely to manually sell a winning position than the next most frequent action. These manual exits are rare for positions in the loss domain. Figure 2, Panel B presents ex-post behavior in the sub-sample of trades whose ex-ante strategies are categorized as “loss-exit”. The distribution of decisions is the same as in the full sample, implying that actual choices follow the opposite pattern from planned choices for the same position.



Panel A. All Positions

Panel B. Positions with loss-exit strategies

Figure 2. Ex-post behavior. This figure illustrates the distribution of actions undertaken by traders in response to experiencing paper gains and losses. 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 strategies (i.e., initial loss limit closer to opening price than initial gain limit).

To further corroborate this pattern, in Figure 3 we explore the dynamics of order revisions and how they relate to the propensity to experience losses. Focusing on trades with at least one limit revision, the figure plots the fraction of positions that are realized at a loss (compared to a gain) against the total number of limit revisions on that position. The size of each point indicates that total number of trades that have the specific number of limit revisions before being sold. The figure shows that a substantial majority of limit revisions are associated with losses. Of positions with exactly one revision, nearly 75% are losses. Apart from an initial dip for the first few revisions—which importantly never comes close to reaching 50%—the association between realized losses and number of revisions is positive. This implies that the greater the number of limit revisions, the more likely the position is experiencing a loss.

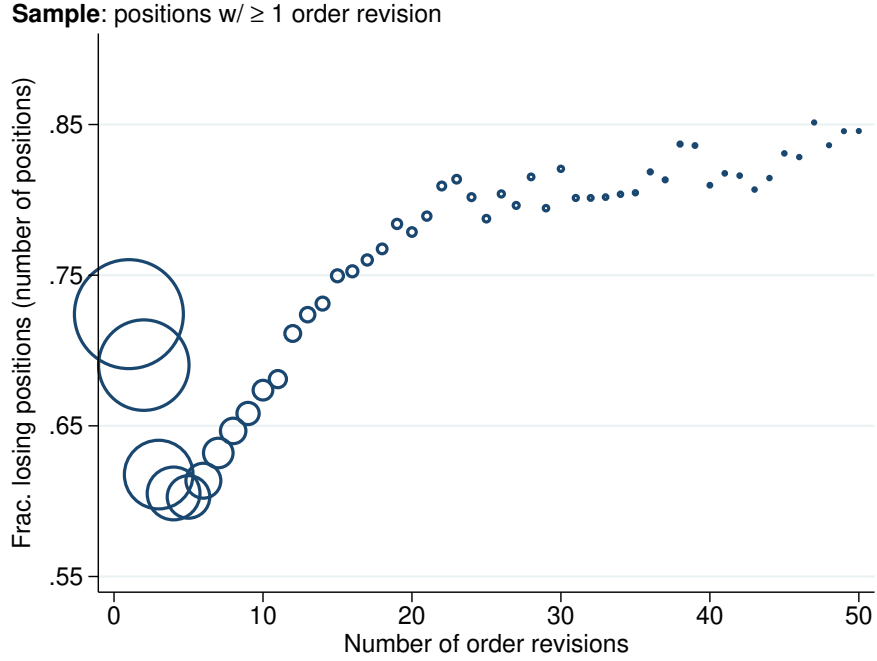


Figure 3. Path of limit revisions and the propensity to have losing positions. This figure plots the fraction of position that execute at a loss against the total number of revisions to loss and gain limits on a given trade. The figure contains positions with at least one revision. The size of each dot indicates the total number of trades that have a given number of limit revisions.

In sum, we find that traders allow larger losses to accumulate and realize gains too early compared to their initial “loss-exit” plans. These decisions result in a distribution of outcomes that skews in the *opposite* direction of traders’ intended strategies when they open new positions.

III. Experimental Design

Results from the field setting show a substantial discrepancy between people’s risk-taking intentions and actual behavior in dynamic settings. We developed an experimental design that allows us to identify whether this discrepancy is generated by an underlying dynamic inconsistency and to formally test and distinguish between models of decision-making under risk. We conducted two preregistered experimental studies with a total of 940 subjects recruited from

Mechanical Turk.¹¹ We designed these experiments to: (i) elicit incentivized initial strategies and compare them to subsequent behavior both within the same participant and across subjects, (ii) examine ‘entry’ decisions—the initial choice to take on risk or not—as a function of number of rounds and access to commitment opportunities, and (iii) study sophistication about dynamic inconsistency given the availability of ‘hard’ and ‘soft’ commitment opportunities.

Participants face binary choices of whether or not to invest a portion of their endowment in fair symmetric gambles that have an expected value of zero. Each gamble features a simple 50/50 chance that the investment either doubles or is lost. If the participant chooses not to invest, she keeps that portion of the endowment. We verify that subjects understand the decision by having them draw ten observations from a stratified sample before making a choice.¹²

There are 3 (4) between-subject treatments in Experiment 1 (Experiment 2), as illustrated in Figure 4. Both experiments feature treatments with one round (One-Shot) and treatments with multiple rounds. Experiment 1 has two multi-round treatments. In the Sequential treatment, participants begin taking on risk knowing that they will receive feedback after every round and can adjust their choice accordingly. In the Hard Plan treatment, participants first choose a strategy and then decide to take on risk knowing that this strategy will be carried out. Experiment 1 allows us to test whether the decision to take on risk is affected by the number of prospective opportunities, identify any dynamic inconsistency between planned and actual choices, and examine participants’ sophistication about it by examining differences in entry rates between treatments. Experiment 2 replicates the three treatments in Experiment 1. It also includes a Soft Plan treatment that elicits participants’ ex-ante strategies, but unlike the Hard Plan treatment, participants can potentially deviate from these strategies ex-post. This treatment allows us to test for dynamic inconsistency within-subject, as well as to examine whether setting a non-binding plan affects the ‘entry’ decision and ex-post behavior.

¹¹ Mechanical Turk is increasingly being used as a recruitment platform for experiments in economics and the broader social sciences (e.g., Martínez-Marquina, Niederle, and Vespa, 2019; Gneezy et al., 2020; DellaVigna and Pope, 2018). Studies have shown that results from the laboratory broadly replicate on Mechanical Turk (see Hauser, Paolacci, and Chandler, 2018; Paolacci and Chandler, 2014, for reviews). Our pre-registration can be found at <https://aspredicted.org/blind.php?x=x954rp> and <https://aspredicted.org/blind.php?x=tn4dt4>.

¹² Kaufmann, Weber, and Haisley (2013) and Hogarth and Soyer (2015) provide evidence on the benefits of sampling for the understanding of probabilities.




Figure 4. 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).

In the One-Shot treatment, participants receive an endowment of 10 cents and decide whether or not to invest in a single gamble. In the multi-round treatments (i.e., Sequential, Hard Plan and Soft Plan), participants face a sequence of the same investment decisions over a maximum of 26 rounds. This number of rounds was chosen for theoretical considerations because it allows us to differentiate between different models of dynamic decision-making under risk.¹³ Each participant receives an initial endowment of \$2.60 at the beginning of the experiment and decides whether to invest 10 cents in each round or to keep it, such that every investment decision was equivalent to the One-Shot treatment.¹⁴ Once a participant chooses to stop investing, she cannot re-enter and the main part of the experiment is over.¹⁵

In the Sequential treatment, participants made decisions round by round and were given

¹³ As discussed further in the Appendix A, the number of rounds was chosen to facilitate differentiation between different models of dynamic decision-making.

¹⁴ In order to rule out that differential endowments would drive differences in entry rates between treatments, a separate experiment included a modified One-Shot treatment where participants were given a \$2.60 initial endowment as in the multi-round treatments, but made the same one-shot decision to invest 10 cents as in the baseline One-Shot treatment. As outlined in Section IV, the results are unchanged, suggesting that wealth effects are unlikely to explain differences in entry rates.

¹⁵ The experimental stakes, which worked out to approximately \$12 an hour, were on the higher end of those typically offered on the platform.

feedback on the previous outcome (gain or loss) in between rounds. Specifically, each participant first decides whether or not to invest in the first round; we refer to this as the ‘entry’ decision. If she chooses to invest, the outcome of the first round is revealed and she decides whether or not to invest in the next 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.

Both Plan treatments have the same structure as the Sequential treatment. The main difference is that participants enter their desired risk-taking strategies before making their ‘entry’ decision. We elicit participants’ strategies by asking them to indicate a gain limit (i.e., the minimum gain at which they would prefer to stop gambling rather than continue) and a loss limit (i.e., the maximum loss at which they would prefer to stop gambling rather than continue). We use these gain and loss limits to infer participants’ risk-taking strategies. For example, a gain and loss limit of +10 and -10, respectively, correspond to the outcome “neutral” strategy of exiting after the first round; a gain limit of +20 and a loss limit of -10 corresponds to a “loss-exit” strategy of taking on more risk after winning in the first round, and exiting after losing in the first round. After entering both limits, participants decide whether or not to begin taking risk. In the Hard Plan treatment, we inform participants that they would automatically stop investing if either of their limits are triggered—hence, the limits are fully incentivized and binding. In the Soft Plan treatment, we inform participants that they will be notified when either of their limits are triggered. Participants can then decide whether or not to continue investing. Hence, the limits correspond to a non-binding plan. If participants decide to continue investing after triggering their limits, they are notified in subsequent rounds when either of the limits are triggered again. The limits themselves cannot be revised during the course of the experiment.¹⁶ Note that even though they are non-binding, the limits in the Soft Plan treatment are incentivized as they are elicited *before* the participant is randomly assigned to either the Hard Plan or the Soft Plan treatment—at the time of elicitation, the participant faces a 50% probability that the limits will be binding.¹⁷

¹⁶ Fischbacher, Hoffmann, and Schudy (2017) propose optional and amendable gain and loss limits as an intervention to reduce the disposition effect. In contrast to this study, the limits in our experiments are not optional to prevent any selection effects and cannot be revised to ensure that they accurately capture participants’ ex-ante strategies.

¹⁷ Note that this was the case in Experiment 2 which included the Soft Plan treatment. Limits were fully incentives in Experiment 1 which only had the Hard Plan treatment. As we show in the next section, decreasing the probability of the limits being binding did not affect participants’ reported strategies.

We programmed both experiments using oTree (Chen, Schonger, and Wickens, 2016) and ran them each in four batches to isolate potential day-of-week effects and Friday/holiday 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 (e.g., age, gender, highest level of education), as well as self-reported level of statistical skills. After the main task, participants complete an exit questionnaire that elicits other control variables.¹⁸ Fewer than 5% of participants exited the experiment after being randomized into treatments. Hence, there is little room for selective attrition. Table III presents participant demographics. Except for participants in Experiment 2 reporting higher perceptions of their own statistical skills, there are no significant differences in demographic characteristics between Experiments 1 and 2.

[INSERT TABLE III ABOUT HERE]

IV. Experimental Results and Discussion

A. *Accepting a Fair Gamble (Entry Decision)*

We begin by studying participants' initial willingness to accept risk, i.e., the *entry decision*. First, we examine whether participants are more likely to take risk if it is part of a dynamic sequence of choices. Second, we analyze whether participants value the ability to commit to an ex-ante strategy. Finally, we examine potential differences between binding and non-binding commitment opportunities.

We find that participants are more likely to accept the fair gamble as part of a sequence than a one-shot gamble offered in isolation. Figure 5 displays the proportion of participants in each of the treatments who accept risk in the first round. We find that the entry rate is substantially lower in the One-Shot treatment compared to any of the multi-round treatments.¹⁹

¹⁸ Appendix B provides a list of all variables.

¹⁹ Note that because subjects receive the endowment at the beginning of the experiment, investment rates may be partly driven by a house money effect (Thaler and Johnson, 1990). Since this effect is not expected to vary by treatment, our analyses focus on differences between treatments rather than absolute levels.

Table IV, Panel A, displays estimates of the binary entry decision across both Experiment 1 and 2. The main independent variables are dummy variables for each of the multi-round treatments (i.e., Sequential, Hard Plan and Soft Plan treatments). We present marginal effects (mfx), which measures the difference between the probability of making the initial investment in the respective multi-round treatment compared to the One-Shot treatment, which is the reference treatment in Panel A. We display all regressions with and without demographic control variables. We estimate cluster-robust standard errors for the regression coefficients in this table, as well as in all subsequent regressions that use the experimental data. The differences between the multi-round and the One-Shot treatments are all significant at the 1% level in both experiments.




Figure 5. Entry decision. This figure shows the percentage of participants in each treatment who accept risk in the first round.

[INSERT TABLE IV ABOUT HERE]

Next, we compare the entry rates between the Sequential and Hard Plan treatments and find that participants are more likely to accept the multi-round fair gamble if they can commit to an ex-ante strategy (see Figure 5). Table IV, Panel B, presents the mfx from Probit regressions of the binary entry decision excluding the One-Shot treatment. In these regressions, the reference treatment is the Sequential treatment, in which participants can revise their choices each round. We find that participants are significantly more likely to initially accept risk if they can commit to an ex-ante strategy with a binding commitment device. Notably, entry rates are also higher

in the Soft Plan treatment, where the commitment opportunity is non-binding. We discuss the implications of this latter result in the next subsection.

Alternative Explanations

In our experiments, taking the fair gamble many times does not result in a higher expected value nor a lower probability of experiencing losses. The expected value remains at zero regardless of one’s strategy or the number of rounds played. Thus, the higher tendency to accept risk for multiple rounds cannot be explained by loss aversion and narrow bracketing, i.e., myopic loss aversion, as in the case of positive-expected-value gambles with infrequent feedback (see Gneezy and Potters, 1997).

Next, we provide evidence that the higher entry rates in the multi-round treatments are unlikely to be driven by biased expectations—specifically, erroneous beliefs that choosing a different risk-taking strategy can affect one’s expected earnings. By design, we hold expectations about the single-round gamble constant by giving participants full information about the gamble and communicating it in a simple, straightforward way—both by description and through experienced sampling. Because the stochastic process is a martingale, it follows that the expected value of any strategy is equal to that of the single-round gamble.²⁰ However, participants may erroneously believe that they can increase the gamble’s expected value by following a “loss-exit” strategy. To test this conjecture, we ran an additional treatment where we clarified that the expected value of *any* strategy is zero—equal to that of the one-round gamble. We ran this treatment along with a replication of the baseline Hard Plan treatment ($N=314$). There was no significant difference in entry rates between the baseline Hard Plan treatment (95.39%) and the new treatment with information about the expected value of strategies (95.68%). As such, over-optimism about the expected value of one’s strategy is unlikely to be driving the higher level of risk-taking in the multi-round treatments compared to the One-Shot treatment.²¹

²⁰ See Doob’s optional sampling theorem.

²¹ Furthermore, to rule out that differential endowments would drive differences in entry rates between treatments, we ran a separate study ($N = 122$) with the One-Shot treatment as described here, a modified One-Shot treatment, and the Hard Plan treatment (between-subject design). In the modified One-Shot treatment, participants were given a \$2.60 initial endowment as in the multi-round treatment, but made the same One-Shot decision to invest 10 cents as in the original One-Shot treatment. Results are unchanged: entry rates in both One-Shot treatments (67.5% in baseline One-Shot treatment and 55.8% in modified One-Shot treatment) were lower than in the Hard Plan treatment (94.9%). Differences with respect to the Hard Plan treatment were statistically significant for both the baseline One-Shot treatment (p -value: .002)

Participants’ perceptions of task complexity are also unlikely to explain the higher levels of risk-taking in the Plan treatments compared to the Sequential treatment. In particular, because the Plan treatments restrict the strategy-choice set to strategies that can be described by a pair of limits, it can be hypothesized that the perceived task-complexity is lower than the task-complexity in the Sequential treatment, which may explain why participants in the former have a higher propensity 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). We find that perceptions of complexity do not significantly differ between the conditions. A related alternative explanation is that the higher entry rate in the plan treatments is caused by decision aversion because the number of decisions participants make in those treatments (i.e., 3) is on average smaller than in the Sequential treatment. However, this implies that the entry rates would have been highest in the treatment with the lowest number of decisions—the One-Shot treatment, which is not the case.

Finally, a sunk cost fallacy with respect to the cognitive effort required to make an explicit ex-ante plan is unlikely to explain differences in entry rates. In a separate study, which Section V describes in detail, we reverse the order of eliciting the plan and the entry decision for the Hard Plan treatment. This should remove potential sunk cost effects prior to the entry choice. We find no significant difference in the entry rates between the baseline Hard Plan treatment (95.39%) and the treatment with reverse order (95.63%).

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

In this section, we examine the dynamic strategies that participants choose ex-ante—whether they accept risk as part of a “loss-exit” strategy, similar to the field setting, or choose a different risk-taking strategy. We then look at participants’ ex-post behavior in the absence of a commitment opportunity to identify dynamic inconsistency. The specific pattern of deviations also allows us to distinguish between different theoretical explanations, e.g., whether the deviations are outcome-specific or not. In the subsections that follow, we analyze the ex-ante strategies and the ex-post behavior both between-subjects (comparing the Sequential and the

and the modified One-Shot treatment (p -value $< .001$).

Hard Plan treatments) and within-subjects (comparing ex-ante versus ex-post behavior in the Soft Plan treatment).

B.1. Ex-Ante Strategies

Figure 6 illustrates the cumulative distribution of gain and loss limits in our sample. Panel A reports results from all participants in both the Hard and Soft Plan treatments. Panel B restricts the sample to participants in the Soft Plan treatment. We find striking differences in loss and gain limits: the distribution of gain limits first-order stochastically dominates the distribution of loss limits. Table V further shows that the average participant who accepts the initial gamble sets a gain limit that is 3.81 times higher than her loss limit. This corresponds to an average difference that is more than 30% of the total endowment. As in the field setting, we define strategies as “loss-exit” if the loss limit is closer to the endowment than the gain limit, as “gain-exit” if the loss limit is further from the endowment than the gain limit, or neutral “symmetric” if both limits are equidistant. The table shows that the overwhelming majority of participants (80.8%) begin taking risk as part of a “loss-exit” strategy, whereas only 7% do so as part of a “gain-exit” strategy.

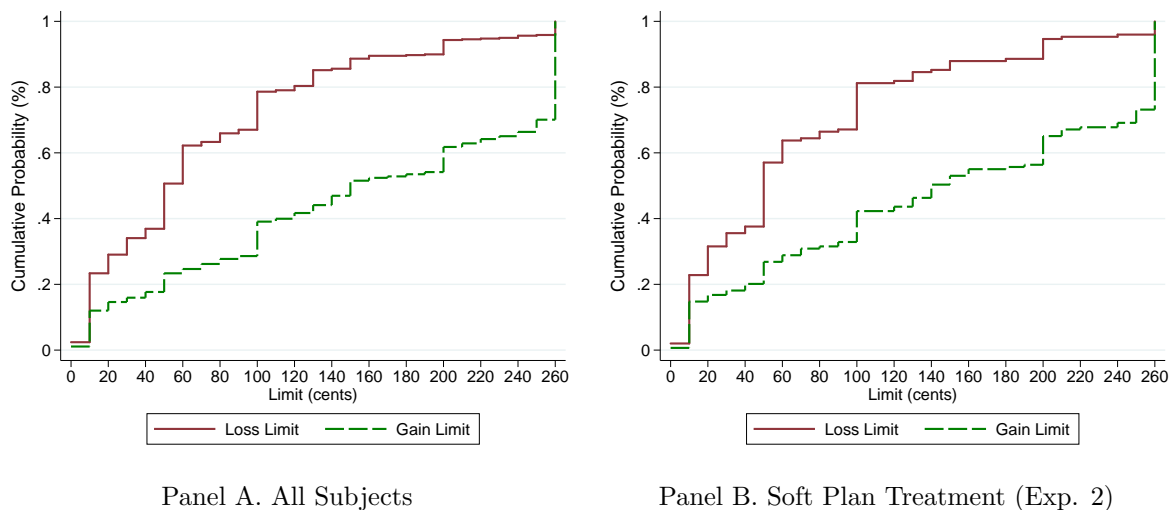


Figure 6. Ex-ante strategy. This figure illustrates participants’ strategies in the Hard Plan and Soft Plan treatments. It shows the cumulative distribution of loss and gain limits. Panel A reports results from both experiments and commitment treatments. Panel B reports results from the Soft Plan treatment only.

[INSERT TABLE V 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 V. To calculate this number, we run 100,000 independent outcome paths for each participant in our sample and determine the individual expected final outcome distribution, which we then aggregate to form the expected final outcome distribution of the representative participant. Note that given the symmetric gamble, positive skewness can only result from an *outcome-dependent* dynamic strategy. The skewness of the realized outcome distributions of final outcomes in the Hard and Soft Plan treatments is close to the skewness of the expected outcome distributions—between 0.247 and 0.471. In Table VI, we report the skewness of the realized outcome distribution among all participants in the Hard and Soft Plan treatments (Panel A), only amongst those who enter (Panel B), and the hypothetical outcome distribution if participants in the Soft Plan treatment would have stuck to their ex-ante strategies (Panel C). In all cases, skewness is positive and statistically significant at the 5% level. The table also shows that skewness in the Sequential treatment is close to zero and not statistically significant at the 10% level, in contrast to the skewness of outcomes generated by ex-ante strategies in the Hard and Soft Plan treatments. Importantly, we find that the *actual* outcome distributions of participants in the Soft Plan treatment does not exhibit any significant skewness, unlike the hypothetical outcome distributions implied by their ex-ante strategies. These results suggest that participants in the Sequential and Soft Plan treatments deviate from their desired “loss-exit” strategies. The following subsection explores this deviation in greater detail.

[INSERT TABLE VI ABOUT HERE]

B.2. *Ex-Post Behavior*


We now analyze participants’ ex-post behavior and compare it to their ex-ante strategies. In particular, we study whether participants’ ex-post decisions depend on the valence of the prior outcome (i.e., whether risk-taking choices differ after experiencing gains versus losses), and whether this pattern is consistent with their ex-ante “loss-exit” plans.

The cleanest test for dynamic inconsistency comes from comparing ex-ante strategies to behavior after feedback in the first round—before endogenous exit decisions have a chance to accumulate. The first-round decision to exit can be conditioned on the simple chance outcome of experiencing a gain or a loss. Since the likelihood of experiencing either outcome is purely determined by chance, any differences in behavior can be attributed to a differential reaction to losses versus gains. In contrast, because participants can choose to exit in any round, behavior in later rounds may be increasingly a function of unobservable differences in exit preferences as opposed to a reaction to cumulative gains and losses. Figure 7, Panel A presents the exit probabilities after first round losses and gains implied by participants’ plans (right) versus their actual behavior (left). We find a striking difference: while participants’ planned behavior implies a significantly higher probability of exiting after a loss than a gain, their actual exit rates follow the opposite pattern—participants are three times more likely to exit after a gain than a loss.

Table VII, columns (1) and (2), displays these results using OLS regressions with main effects of the Sequential treatment compared to the commitment treatments, a Gain dummy, and the treatment-outcome interaction. The analysis shows that the higher probability of exiting after a gain than a loss represents a substantial deviation from participants’ ex-ante strategies in response to first round outcomes. The interaction term measures the difference in probabilities and it is statistically significant at the 1% level. In sharp contrast to participants’ “loss-exit” plans, this behavioral pattern is consistent with a “gain-exit” strategy in ex-post choices—a dynamic inconsistency between planned and actual behavior.


The differences in first-round exit rates translate into differences in overall exit rates with respect to cumulative earnings. Figure 7, Panel B and Panel C show that participants in the Sequential treatment are more likely to end up with a cumulative gain than a cumulative loss even though the gamble is symmetric. Table VII, Columns (3) and (4), 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. We find that the probability of realizing a cumulative gain is significantly higher in the Sequential treatment than in the outcomes generated by participants’ ex-ante strategies.

[INSERT TABLE VII ABOUT HERE]



Panel A. Round 1 Exit

Panel B. All Participants

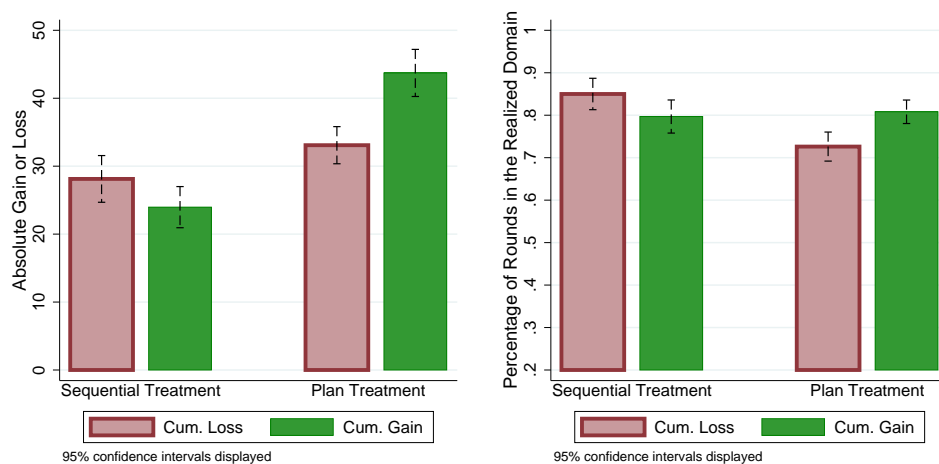


Panel C. Participants Who Enter

Figure 7. Probability of realizing a gain versus loss—planned versus actual behavior. This figure illustrates the actual behavior in the Sequential treatment compared to the ex-ante plans in the Plan treatments as a function of prior gains and losses. The Plan treatments include the hypothetical cumulative outcomes of the participant in the Soft Plan treatment. Panel A shows the percentage of participants who enter and then stop investing after feedback in the first round. Panel B shows the percentage of participants who earn a cumulative gain versus a cumulative loss across Experiments 1 and 2. Panel C excludes participants who choose not to enter.

We now examine whether participants' ex-post cumulative earnings in the experiment differ from those implied by their ex-ante strategies. Figure 8 compares the absolute cumulative gains and the absolute cumulative losses in the Sequential treatment and the Plan treatments (Panel A). It also shows what percentage of rounds participants are willing to experience gains

versus losses before rejecting further risk. For both measures, the pattern between gains and losses in the Sequential treatment is the reverse of the one in the Plan treatments. In the Plan treatments, participants leave more room for large gains and restrict the scope of their losses (Panel A). To achieve that, they allow for larger gains to accumulate compared to losses, which are cut sooner (Panel B). In contrast, the average participant in the Sequential treatment ends up realizing larger losses than gains (Panel A) and staying in the loss domain for longer than in the gain domain (Panel B). Table VIII shows that the difference-in-difference in both measures is statistically significant at the 1% level, providing evidence that the dynamic inconsistency between planned and actual behavior leads to differences in cumulative earnings.



Panel A. Absolute Cum. Gains/Losses

Panel B. Realization Reluctance

Figure 8. Cumulative Gains and Losses—ex-ante versus ex-post. This figure illustrates differences in the outcomes and behavior between participants who realize a cumulative loss versus gain in the Plan versus Sequential treatments (ex post) across Experiments 1 and 2. The Plan treatments includes the hypothetical cumulative outcomes of participants in the Soft Plan treatment; hence in case the participant 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 participants’ reluctance to realize their final cumulative outcome as measured by the percentage of rounds the participants’ cumulative outcome has been in the domain that they ended up realizing.

[INSERT TABLE VIII ABOUT HERE]

We now examine dynamic inconsistency within-subject by looking at participants’ behavior in the Soft Plan treatment, in which participants are allowed to revise their decisions after being

notified that a limit has been reached. For the 57 participants whose limits were triggered, 80.7% decide to continue investing and thus deviate from their ex-ante strategies. The majority (70.2%) of the triggered limits are loss limits.²² Strikingly, the most common type of deviation is to begin overriding the loss limit down in the beginning and to do so until the very end. This pattern replicates the findings from the field which shows that people are significantly more likely to revise their loss limit down compared to their gain limits.

Returning to the decision to take on risk, the difference in entry rates between the Plan and Sequential treatments suggests that participants are at least partially sophisticated about their dynamic inconsistency. However, though people may be aware of their dynamic inconsistency, they may also over-estimate the efficacy of soft commitment for disciplining future deviations. To measure the scope of sophistication across both domains, we examine entry rates under binding and a non-binding commitment opportunities. Figure 5 and Table IV, Panel B, show that a non-binding commitment opportunity also significantly increases the likelihood that participants accept initial risk; the entry rate for non-binding commitment is roughly the same as with binding commitment (the difference is not significant at the 10% level). Yet as outlined above, soft commitment is not effective in mitigating dynamic inconsistency. This is particularly important because non-binding commitment is frequently offered in many real-world settings as part of the choice set (e.g., revisable loss and gain limits) or incorporated into policy aimed at improving outcomes for non-institutional investors (e.g., the MiFID II regulation in Europe). As we discuss in Section VI, such soft-commitment opportunities can have unintended adverse consequences for consumer welfare.

Finally, one potential concern for our analysis of dynamic inconsistency is that differences in entry decision may bias inference about ex-post choices in the first round. To address this we estimate lower bounds for the discrepancy between planned and actual decisions. To do so, we (i) make the most conservative assumption about selection bias in the Sequential treatment, (ii) run simulations of the Plan treatments assuming the same type of selection, and (iii) rerun the difference-in-difference analysis using the simulated Plan treatments against the Sequential treatment. Specifically, we first assume that entry decisions in the Sequential treatment lead us to overestimate differences between the treatments, which would be the case if participants

²² This should be expected given that most participants set loss limits closer than the gain limits.

who chose to enter in the Plan treatments but not in the the Sequential treatment preferred a “loss-exit” to a “gain-exit strategy. Therefore, we run a simulation of the Plan treatments by excluding the 10% of participants with the most pronounced “loss-exit” strategies (i.e., lowest loss limits and highest gain limits). We repeat the difference-in-difference analysis of exit behavior after the first round as presented in Table VII, Columns (1) and (2). The coefficients on the interaction term decrease slightly to 0.246 (instead of 0.287) without demographic controls and 0.264 (instead of 0.305) with demographic controls, but are still significant at the 1% level. Because participants with the most extreme “loss-exit” strategies were excluded from the analyses, these coefficients represent lower bounds of the differences between planned and actual behavior. Finally, the design of our Soft Plan treatment allows us to rule out selection issues because participants form initial strategies *and* make ex-post choices, which allows us to identify dynamic inconsistency within-subject.²³

C. *Interpreting the Findings*

To summarize, we find that people begin taking risk as part of a “loss-exit” strategy, planning to stop earlier after losses than after gains. However, peoples’ actual choices follow the opposite pattern—they take on more risk after losses than after gains. We find some evidence that people are somewhat sophisticated about this dynamic inconsistency, exhibiting a greater initial willingness accept risk when offered an opportunity to commit to their strategy. Finally, we document a greater willing to accept risk when it is part of a sequence of choices than when the same one-round gamble is offered in isolation.

In Appendix A, we formally derive the dynamic predictions of Expected Utility (with and without skewness preferences), Rank-Dependent utility, Quasi-Hyperbolic Discounting, and Cumulative Prospect Theory in our setting. The combination of findings outlined above is most consistent with CPT. The model predicts that agents may reject a single fair gamble while accepting the same gamble when it is a part of a dynamic sequence. This greater willingness to accept risk is due to the agent’s ability to change the final outcome distribution through her

²³ Moreover, because entry rates were essentially the same and close to 100% in the Hard and Soft Plan treatments, we can also compare the binding initial strategies in the former to ex-post behavior in the latter as further evidence for dynamic inconsistency.

risk-taking strategy. A “loss-exit” strategy increases the positive skew over potential earnings compared to a gamble in isolation. Probability weighting—which leads the agent to overweight unlikely outcomes—makes it more attractive to accept risk as part of this strategy than in isolation. However, CPT also predicts that agents will systematically deviate from their “loss-exit” plan and exhibit “gain-exit” behavior. Diminishing sensitivity of the value function generates risk-seeking behavior in the loss domain and risk-aversion in the gain domain. This will lead the agent to stop taking on risk earlier after gains and to continue for longer after losses, relative to her initial strategy.²⁴ Finally, an agent who is 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.

To illustrate the intuition for these predictions, consider an agent facing a sequence of two 50:50 gambles, each with an upside of G and downside of $-G$. A “loss-exit” strategy generates 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 an outcome-independent strategy of accepting both gambles regardless of a gain or loss, $(2G, 1/4; 0, 1/2; -2G, 1/4)$. *Probability weighting* prompts the agent to overweight smaller probabilities, which leads her to find the gamble more attractive as part of a “loss-exit” strategy than when offered in isolation. The same agent who accepts the first bet as part of this strategy systematically deviates from it by stopping too early after winning and continuing on too late after losing (relative to her plan). There are two reasons for this deviation. First, *diminishing sensitivity* of the value function makes it less attractive to take on risk after a gain than after a loss. Second, after one round, a strategy that involves accepting the next gamble does not generate as much positive skew as the one which included the previous gamble—the agent is now looking at a prospect with even odds. With the probability weighting motive being less pronounced, diminishing sensitivity leads to “gain-exit” behavior.

²⁴ This specific form of outcome-dependent dynamic inconsistency distinguishes CPT from other models which also predict dynamic inconsistency, but where either ex-ante strategies or ex-post behavior do not depend on whether the agent experiences gains or losses (e.g., Rank Dependent Utility).

V. Welfare Implications

In this section we examine the potential welfare consequences of dynamic inconsistency in choice under risk. Our empirical results show that people prefer loss-exit strategies before they begin taking risk. However, after experiencing gains and losses, they systematically deviate from their plans to take on more risk after losses and stop earlier after gains—exhibiting gain-exit behavior. While peoples’ strategies generate a positively-skewed distribution over final earnings, their actual choices—which are inconsistent with their ex-ante preferences on how to respond to prior realizations—generate a more negatively-skewed distribution. If one set of choices can be classified as a “mistake”, then the dynamic inconsistency will have welfare consequences that can be assessed using a best-fit model.

However, Bernheim and Taubinsky (2018) underscore that it is unclear whether either set of choices can be classified as a mistake without additional data. One cannot appeal to normative standards when interpreting the data through CPT because both planned and actual choices are subject to psychological frictions. In order to assess agents’ welfare in this normatively-ambiguous domain, we adopt the techniques described in Bernheim and Taubinsky (2018). There, the authors outline how the behavioral welfare framework of Bernheim and Rangel (2009) can be used to classify mistakes in empirical applications.

The initial stage of the analysis requires identifying which choices *merit deference*, i.e., should be included in the welfare relevant domain. A large literature in neuroscience and psychology has shown that the neurocognitive constructs of negative urgency and compulsivity channel excess attention to prior outcomes in sequential risk-taking (for review, see Zhang and Clark, 2020). This research argues that a large fraction of outcome-dependent choices—differential responses to prior gains and losses—reflects a failure to appropriately balance activity within the neural system in how conflicting motivational states are coded (Campbell-Meiklejohn et al., 2008). Particularly in the case of prior losses, excessive attention spurs loss chasing that can lead to ruinous financial outcomes (Hodgins, Stea, and Grant, 2011). This evidence suggests that the welfare relevant domain should favor ex-ante planned choices over ex-post decisions in response to experienced losses and gains, which exhibit chasing of the former and the early

realization of the latter.²⁵

To provide empirical supporting for using participants’ ex-ante plans as the welfare-relevant benchmark, we ran a separate study ($N = 151$) that used the re-framing technique outlined in Bernheim and Taubinsky (2018) to identify a ‘frame-invariant’ welfare measure. The study examined ex-post choices using a decision frame that emphasized final outcome distributions over prior realizations of risk.²⁶ It begins in the same way as the Sequential treatment. However, after revealing whether the participant had lost or gained in the first round, we asked her to state the most that she would be willing to lose and gain in total. Importantly, the maximum loss and gain numbers include the participant’s *current* earnings, and as a result, this set of choices maps directly onto those in our standard Sequential treatment for the same round. For example, after a one-round loss, a participant stating that the most she’d be willing to lose is one round’s endowment would imply immediate exit; stating a number that corresponds to two rounds of losses implies accepting at least one more round of the lottery. We refer to this method of eliciting choices as the *outcome frame*. Importantly, the outcome frame does not alter opportunities relative to the standard *sequential choice frame*—participants’ choices in the outcome frame correspond to decisions of either immediately exiting or continuing to take on risk in the same manner as in the sequential choice frame. Rather, the outcome frame is designed to shift participants’ attention away from prior gains and losses, instead focusing attention on the final outcomes of their choices. This approach is analogous to that of Chetty, Looney, and Kroft (2009), Allcott and Taubinsky (2015), and Taubinsky and Rees-Jones (2018), who conduct welfare analysis using the Bernheim-Rangel framework by re-framing choices without affecting the underlying opportunities.

First, in contrast to ex-post choices in the sequential frame, we find that subjects in the outcome frame followed a “loss-exit” strategy—stating that they would be willing to take significantly more risk after gains than after losses. The average numbers corresponding to prospective gains were nearly three times higher than for prospective losses ($p < .001$). Importantly, the outcome frame appeared to successfully shift attention away from prior outcomes. There were no

²⁵ Bernheim and Rangel (2004) similarly rely on evidence from neuroscience and cognitive psychology to argue that the welfare relevant domain should not include choices made in the presence of substance-related environmental cues, e.g., the choice to indulge in excessive drinking after walking past one’s favorite bar.

²⁶ The study was pre-registered here: <https://aspredicted.org/blind.php?x=vw3uu8>.

significant differences in participants' choices after experiencing a loss or a gain. In both cases, participants' choices implied earlier exit after losses than gains (gain numbers set at 2.36 versus 3.34 times higher than loss numbers, respectively, $p = .60$). Furthermore, participants' *actual* choices in the outcome frame generate a positively-skewed distribution over final outcomes (see Table VI, Panel D), which is similar to the outcome distribution implied by participants' ex-ante plans (see Table VI, Panels B and C). Despite presenting participants with the same opportunities, their choices in the outcome frame were significantly different than the ex-post choices in the sequential frame, as indicated by substantially higher propensities of exiting earlier after losses and taking on more risk after gains in the former than in the latter ($p < .01$).

Based on this evidence, we now proceed to assess the welfare costs of dynamic inconsistency. We use the CPT framework to assess welfare because it is most consistent with our empirical findings. As outlined in Barberis (2012), being dynamically inconsistent leads to two types of potential welfare losses for a CPT agent, which affect naïve and sophisticated individuals differently. Some naïve agents begin investing because of a mistaken belief that they will stick to their ex-ante strategy. For sophisticated agents aware of their dynamic inconsistency, the utility of investing in the first gamble is lower than rejecting it. As a result, naïve agents who accept the gamble and deviate from their intended strategy incur a larger welfare loss compared to sophisticated agents, who reject the gamble from the beginning. We refer to this form of welfare loss as the *cost of naiveté*. Both naïve and sophisticated agents can potentially incur another type of welfare loss, which stems from the opportunity cost of not having access to binding commitment opportunities. For dynamically inconsistent agents, binding commitment is the only method of implementing their ex-ante utility-maximizing strategy. We refer to the utility difference between implementing the preferred ex-ante strategy through binding commitment and rejecting risk due to a lack of commitment opportunities as the *value of commitment*.

We run simulations to assess the welfare costs resulting from dynamic inconsistency for our experimental setting in which the distribution of the gamble and the maximum length of the sequence is known. We calculate utility using the set of CPT parameters $\{\alpha, \delta, \lambda\}$, where α corresponds to the diminishing sensitivity parameter, δ corresponds to the probability weighting parameter, and λ corresponds to loss aversion. To measure welfare, we calculate the certainty equivalent of the aggregate outcome distribution resulting from the ex-ante strategy

and the ex-post behavior, respectively. We obtain the outcome distributions from simulations in which the one-round investment amount is a numeraire (see Appendix A for further details about the simulations). Hence, certainty equivalents are measured as multiples of the one-round investment amount. We refer to a CPT agent with preference parameters $\{\alpha, \delta, \lambda\} = \{0.88, 0.65, 2.25\}$ —median estimates from Tversky and Kahneman (1992)—as a ‘representative agent.’

Figure 9 presents the certainty equivalents of the ex-ante strategy (i.e., value of commitment) and the ex-post behavior (i.e., cost of naïveté). Positive values in Panel A indicate how much a sophisticated agent would be willing to pay for a binding commitment device that guarantees execution of her ex-ante strategy.²⁷ Agents with stronger skewness preferences (i.e., lower δ) and higher sensitivities (i.e., higher α) than the ‘representative agent’ would be willing to pay up to 166% of the one-round endowment depending on the parameter combination.

In the absence of a commitment device, both naïvé and sophisticated agents incur a welfare loss that corresponds to the value of commitment reported in Figure 9, Panel A. Naïvé agents incur another potential welfare loss because they are unaware of their inability to implement their ex-ante strategy and accept the initial gamble rather than rejecting it. These costs of naïveté are illustrated in Figure 9, Panel B. Notably, all CPT agents who would optimally select a “loss-exit” strategy incur costs of naïveté as indicated by the negative certainty equivalents for all parameter combinations with high skewness preferences and low sensitivity. Around the parameter region of the representative agent, the costs of naïveté are over 110% of the one-round endowment amount. Importantly, the costs of naïveté are even higher for CPT agents with greater probability weighting and lower sensitivity. Depending on the parameter combination, a CPT agent may incur a cost that corresponds to a certain loss of up to 208.5% of the one-round endowment. Note that the relationship between probability weighting and naïveté is important from a policy perspective—a positive relationship would imply that those who bear the highest costs of dynamic inconsistency are also the ones most prone to it.

²⁷ Note that the value of commitment for the ‘representative agent’ is close to zero by construction. This is because the maximum number of rounds was deliberately chosen such that this agent would be close to indifferent between accepting or rejecting the first gamble.

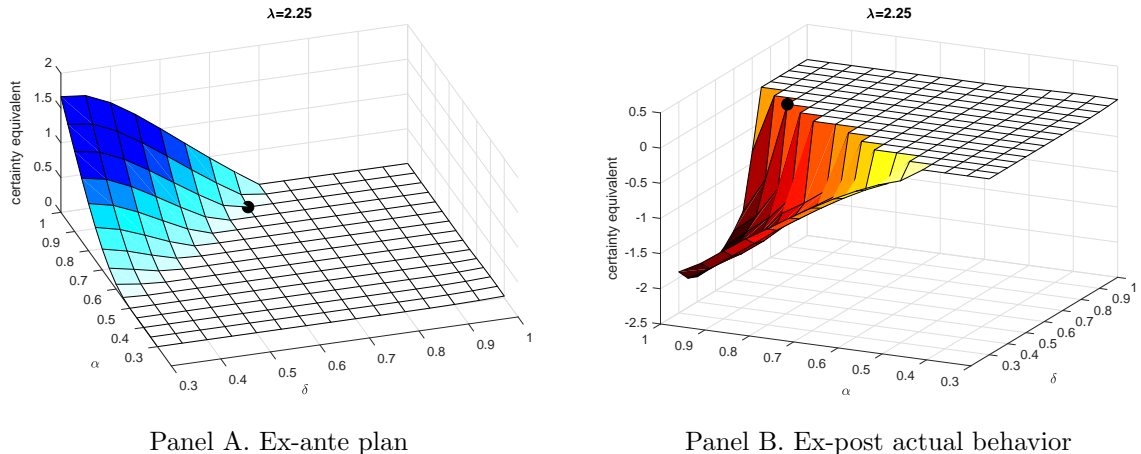


Figure 9. Welfare implications of dynamic inconsistency: value of commitment and costs of naïveté. This figure illustrates the certainty equivalent of agents with CPT preferences with different levels of probability weighting (δ) and diminishing sensitivity (α), i.e., the extent to which agents are risk-averse in the gain domain and risk-seeking in the loss domain. Loss-aversion is taken as fixed at $\lambda = 2.25$, corresponding to the ‘representative’ level 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 begin taking risk if she could commit to her ex-ante strategy, and would reject risk otherwise. Panel B reports the certainty equivalent of the ex-post outcome distribution. It can be interpreted as *costs of naïveté* that agents endure if they begin to take on risk without a commitment device as opposed to rejecting it.

VI. Conclusion

We show that people are dynamically inconsistent when taking risk repeatedly while knowing that they have the option to stop at any time. In particular, they begin to take on risk with the strategy of stopping after small losses and continuing after gains. However, their actual behavior exhibits the opposite pattern—people cut gains early and chase their losses. We find at least partial evidence for sophistication about this dynamic inconsistency. Interpreting this data through the lens of theory suggests that people accept risk offered as part of a sequence that they would reject in isolation because the dynamic environment allows them to form plans that makes the distribution of potential outcomes more attractive. Lastly, we demonstrate that naïveté about deviations from these plans imply significant welfare costs.

The evidence on dynamic inconsistency and people’s naïveté about the effectiveness of non-binding commitment opportunities highlights the discretion of industry (e.g., casinos, credence goods providers, financial brokers, etc.) to design and soften the commitment opportunities available to their clients. For example, ‘soft’ commitment opportunities may have unintended effects when bundled with regulation that targets consumer financial products. Consider the recently introduced “depreciation reporting rule”—part of the revised European market in financial instruments regulation (MiFID II) implemented in January 2018.²⁸ The rule requires all European wealth managers, brokers, and financial advisers to immediately notify their clients when their portfolio loses at least 10% of its value relative to the beginning of the quarter. This constitutes an exogenous non-binding loss limit at the level of 10%. Although the rule was intended to protect retail investors, our results suggest that the notifications will be redundant for the majority of investors as they are likely to override them. More importantly, however, the rule might change investors’ ex-ante choices regarding the type and amount of risk to take: they may seek out types of risk that they would otherwise avoid—such as volatile assets with a zero or negative risk premium—because of naïveté regarding the effectiveness of the rule in disciplining behavior. Instead of helping investors make better financial decisions, the regulation may exacerbate the types of losses that it was designed to protect retail traders from.

Additionally, our finding that people are more likely to accept risk when they can commit to an initial strategy may have implications for the observed low stock market participation rate amongst US households (Mankiw and Zeldes, 1991). Despite a high historic returns, less than half of households own stocks directly or indirectly (Ameriks and Zeldes, 2004). Part of the reason that people abstain from participating in asset markets may be sophistication about their own dynamic inconsistency; namely, people believing that after deciding to invest, they will accept more losses than they are willing to bear from an ex-ante perspective. To that end, offering households opportunities to commit to an investment strategy may attract more of them to participate in financial markets. That being said, we do not claim that this mechanism is the only, or even the dominant, factor that explains stock market participation.²⁹

²⁸ 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/>.

²⁹ Other factors such as historical reasons or attitudes toward capital markets, heterogeneity in risk aversion, and fixed entry costs have shown to be important determinants of stock market participation (Gomes and Michaelides, 2005).

Finally, our results relate to the work on self-control, impulsivity, and financial decision-making. Papers have linked proxies for impulsivity such as propensity to smoke (Uhr, Meyer, and Hackethal, 2019), drink alcohol (Ben-David and Bos, 2021), or procrastinate (Brown and Previtro, 2016) to increased trade frequency and inferior financial performance. The form of dynamic inconsistency studied in the current paper may or may not be linked to the types of self-control proxies considered in these papers. An important avenue of future research would link dynamic inconsistency in choice under uncertainty to other measures of impulsivity, which would potentially increase the scope for targeted policy interventions.

VII. Tables

Table I
Summary Statistics of Trading Data

This table reports summary statistics from the brokerage data. Trading experience is measured at the time of the trader’s last action in the brokerage data. A “loss-exit” (“gain-exit”) strategy is when the stop-loss (take-profit) order is a smaller distance from the opening spot price than is the take-profit (stop-loss) order.

<i>Panel A: Trader characteristics</i>				
	N_{total}	Mean	Median	Std dev.
Female	159,668	0.18		
Trading experience (years)	187,521	0.88	0.33	1.21
Location ...				
Africa	187,099	0.11		
Asia	187,099	0.21		
Europe	187,099	0.54		
North America	187,099	0.054		
Oceania	187,099	0.022		
South America	187,099	0.076		
Total number of trades	187,521	83.0	13	392.6
<i>Panel B: Position-level statistics</i>				
	N_{total}	Mean	Median	Std dev.
Long position	15,571,278	0.51		
Holding period (hours)	15,571,278	83.4	3.60	448.6
Leverage (XX:1)	15,571,278	163.4	100	142.6
Initial stop-loss/take-profit strategy ...				
gain-exit strategy	15,571,278	0.344		
equidistant limits	15,571,278	0.264		
loss-exit strategy	15,571,278	0.392		

Table II
Traders' Ex-Ante Strategies

This table reports the coefficients of OLS regressions using brokerage data. The dependent variable equals one if the trade has an ex-ante “loss-exit” strategy in which the stop-loss order is a smaller distance from the opening spot price than is the take-profit order. Panel A includes independent variables that reflect trader characteristics. Panel B includes independent variables related to the characteristics of each trade. Standard errors, in parentheses, are clustered by trader *, ** and *** indicate statistically significant at the 10%, 5%, and 1% level, respectively.

<i>Panel A</i>				
	Ex-ante strategy (loss-limit = 1)			
	(1)	(2)	(3)	(4)
Constant	0.533*** (0.0050)	0.534*** (0.0056)	0.521*** (0.0042)	0.535*** (0.0074)
female		-0.00480 (0.012)		
trader location (Africa omitted)				
Asia			0.00158 (0.0076)	
Europe			0.00440 (0.0054)	
N Amer			0.0225*** (0.0058)	
Oceania			0.0206** (0.0098)	
S Amer			0.0920* (0.048)	
trading experience ² (years)				0.00324** (0.0014)
trading experience (years)				-0.0102 (0.0069)
R2	-8.1e-11	0.000012	0.0022	0.00023
N	11,465,145	11,322,186	11,437,756	11,465,145

Table II
Traders' Ex-Ante Strategies

<i>Panel B</i>				
	Ex-ante strategy (loss-limit = 1)			
	(5)	(6)	(7)	(8)
Constant	0.540*** (0.0052)	0.531*** (0.014)	0.569*** (0.0026)	0.582*** (0.024)
long position	-0.0140*** (0.0039)			
currency pair groups (EUR/USD omitted)				
USDpairs		-0.00224 (0.014)		
EURpairs		0.00874 (0.014)		
JPYpairs		0.0242* (0.014)		
position leverage (400:1 omitted)				
2:1			-0.222 (0.21)	
5:1			0.161*** (0.058)	
10:1			-0.0608*** (0.0068)	
25:1			-0.0886*** (0.0049)	
50:1			-0.0566*** (0.0034)	
100:1			-0.0408*** (0.013)	
200:1			-0.0163** (0.0082)	
log(position capital)				-0.0158** (0.0062)
R2	0.00020	0.00019	0.0042	0.0021
N	11,465,145	11,465,145	11,465,145	11,465,145

Table III
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 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) 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 IV
Sequence and Commitment Effects on Entry Decision

This table reports the marginal effects (mfx) 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 commitment 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.

<i>Panel A</i>						
	Experiments 1 & 2		Experiment 1		Experiment 2	
	(1) mfx	(2) mfx	(3) mfx	(4) mfx	(5) mfx	(6) mfx
D^{seq}	0.125*** (4.418)	0.125*** (4.433)	0.136*** (3.409)	0.139*** (3.586)	0.102*** (2.775)	0.097*** (2.710)
$D^{hardplan}$	0.230*** (7.351)	0.229*** (7.356)	0.241*** (5.511)	0.239*** (5.695)	0.198*** (4.765)	0.195*** (4.784)
$D^{softplan}$					0.177*** (4.399)	0.174*** (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
<i>Panel B</i>						
	Experiments 1 & 2		Experiment 1		Experiment 2	
	(1) mfx	(2) mfx	(3) mfx	(4) mfx	(5) mfx	(6) mfx
$D^{hardplan}$	0.088*** (3.435)	0.089*** (3.520)	0.087** (2.397)	0.087** (2.553)	0.083** (2.462)	0.083** (2.522)
$D^{softplan}$					0.065** (2.006)	0.065** (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 V
Ex-Ante Strategies

This table illustrates the ex-ante strategies in the Hard and Soft plan treatments across Experiments 1 and 2. Panel A reports the share of participants who have a loss-exit, gain-exit, or symmetric neutral strategies. 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 participants. Columns (5) to (8) reports the results only for those who initially choose to take on risk. 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 $H_0 : Ratio = 1$ and $H_0 : Diff = 0$, respectively, are in parentheses. Panel C reports the mean and the skewness of the aggregate outcome distribution that results from participants’ gain and loss limits in expectation. The outcome distribution for each participant (each set of gain and loss limits) results from 100,000 independently simulated outcome paths.

Experiment	All Subjects				Entered Lottery			
	(1) 1&2	(2) 1	(3) 2 (hard)	(4) 2 (soft)	(5) 1&2	(6) 1	(7) 2 (hard)	(8) 2 (soft)
N	458	158	151	149	427	147	142	138
<i>Panel A. Strategy Categorization</i>								
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%
<i>Panel B. Aggregate Statistics</i>								
Ratio	3.81*** (10.71)	3.67*** (6.17)	4.51*** (6.45)	3.24*** (6.19)	3.63*** (10.11)	3.56*** (5.79)	4.25*** (6.11)	3.05*** (5.87)
Difference	80.26*** (19.67)	85.00*** (13.84)	81.06*** (10.51)	74.43*** (10.13)	79.98*** (19.01)	86.26*** (13.63)	80.14*** (10.22)	73.12*** (9.52)
<i>Panel C. Expected Outcome Distributions</i>								
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 VI
The Effect of Commitment on Outcome Distributions

This table reports sample distribution parameters of the realized outcome distributions across the different treatments. The outcome is the cumulative gain or loss from the beginning of the sequential lottery until the participant chooses to stop investing. Panel A compares the Hard and Soft Plan treatments from Experiments 1 and 2 to the Sequential treatment. 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 participants who choose to initially take on risk. Panel C focuses only on the Soft Plan treatment 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). Panel D tests the parameters of the outcome distribution in an additional experiment, using an outcome frame to elicit subjects' gain and loss limits after the outcome of the first round has been revealed (as described in Section V). 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.

<i>Panel A: Between-Subject Test (Experiment 1 and 2)</i>			
	N	Mean	Skewness
Commitment Treatment	458	2.445 (0.217)	0.268** (0.019)
Sequential Treatment	318	1.006 (0.515)	0.060 (0.656)
<i>Panel B: Between-Subject Test (Experiment 1 and 2)</i>			
	N	Mean	Skewness
Commitment Treatment	427	2.623 (0.217)	0.247** (0.037)
Sequential Treatment	270	1.185 (0.515)	0.037 (0.798)
<i>Panel C: Within-Subject Test (Experiment 2)</i>			
	N	Mean	Skewness
Soft Commitment (Ex-Ante)	138	1.812 (0.618)	0.471** (0.023)
Soft Commitment (Ex-Post)	138	2.319 (0.569)	0.280 (0.166)
<i>Panel D: Outcome Frame</i>			
	N	Mean	Skewness
Outcome Frame	151	4.967 (0.062)	0.612*** (0.003)

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

This table reports the coefficients of OLS and the marginal effects of Probit regressions. In Columns (1) and (2), the dependent variable is a dummy variable that equals one if the participant stops investing in Round 1 after the first outcome of the gamble is revealed (excluding participants 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 dependent variable in Columns (3) and (4) is a dummy variable that equals one if the participant's final outcome is a cumulative gain as opposed to a cumulative loss (zeros are excluded). The Plan treatments includes the hypothetical cumulative outcomes of the participant in the Soft Plan treatment; hence in cases where the participant 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.

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

Table VIII
Cumulative Gains and Losses

This table reports the coefficients of OLS regressions. The dependent variables are the absolute cumulative gains/losses (Columns (1) and (2)), and the participants' 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 participant has had a paper cumulative gain (loss) conditional on her realizing 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 treatments include the hypothetical cumulative outcomes of the participant in the Soft Plan treatment; hence in case the participant 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 8.

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

REFERENCES

- Allcott, Hunt, and Dmitry Taubinsky, 2015, Evaluating behaviorally motivated policy: Experimental evidence from the lightbulb market, *American Economic Review* 105, 2501–38.
- Ameriks, John, and Stephen P Zeldes, 2004, How do household portfolio shares vary with age, Technical report, working paper, Columbia University.
- 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 Wei Xiong, 2012, Realization utility, *Journal of Financial Economics* 104, 251–271.
- 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.
- Barkan, Rachel, and Jerome R Busemeyer, 2003, Modeling dynamic inconsistency with a changing reference point, *Journal of Behavioral Decision Making* 16, 235–255.
- Ben-David, Itzhak, and Marieke Bos, 2021, Impulsive consumption and financial wellbeing: Evidence from an increase in the availability of alcohol, *Review of Financial Studies* Forthcoming.
- Benartzi, Shlomo, and Richard H Thaler, 1995, Myopic loss aversion and the equity premium puzzle, *The Quarterly Journal of Economics* 110, 73–92.
- Bernard, Sabine, Martin Weber, and Benjamin Loos, 2020, How variance and skewness shape retail investors’ selling behavior, *mimeo* .

- Bernheim, B Douglas, and Antonio Rangel, 2004, Addiction and cue-triggered decision processes, American Economic Review 94, 1558–1590.
- Bernheim, B Douglas, and Antonio Rangel, 2009, Beyond revealed preference: choice-theoretic foundations for behavioral welfare economics, The Quarterly Journal of Economics 124, 51–104.
- Bernheim, B Douglas, and Dmitry Taubinsky, 2018, Behavioral public economics, in Handbook of Behavioral Economics: Applications and Foundations 1, volume 1, 381–516 (Elsevier).
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2012, Saliency theory of choice under risk, The Quarterly Journal of Economics 127, 1243–1285.
- Brown, Jeffrey R, and Alessandro Previtro, 2016, Saving for retirement, annuities and procrastination, Technical report, Working Paper.
- Camerer, Colin F., and Teck-Hua Ho, 1994, Violations of the betweenness axiom and nonlinearity in probability, Journal of Risk and Uncertainty 8, 167–196.
- Campbell-Meiklejohn, Daniel K, Mark W Woolrich, Richard E Passingham, and Robert D Rogers, 2008, Knowing when to stop: the brain mechanisms of chasing losses, Biological Psychiatry 63, 293–300.
- Chen, Daniel, Martin Schonger, and Chris Wickens, 2016, otreeâan open-source platform for laboratory, online, and field experiments, Journal of Behavioral and Experimental Finance 9, 88–97.
- Chetty, Raj, Adam Looney, and Kory Kroft, 2009, Saliency and taxation: Theory and evidence, American Economic Review 99, 1145–77.
- Dai, Hengchen, Katherine L Milkman, and Jason Riis, 2014, The fresh start effect: Temporal landmarks motivate aspirational behavior, Management Science 60, 2563–2582.
- DellaVigna, S., and U. Malmendier, 2006, Paying not to go to the gym, American Economic Review 96, 694–719.
- DellaVigna, Stefano, and Devin Pope, 2018, What motivates effort? evidence and expert forecasts, The Review of Economic Studies 85, 1029–1069.

- Dertwinkel-Kalt, Markus, Jonas Frey, and Mats Köster, 2020, Optimal stopping in a dynamic salience model, mimeo .
- Ebert, Sebastian, and Philipp Strack, 2015, Until the bitter end: On prospect theory in a dynamic context, American Economic Review 105, 1618–1633.
- Evers, Ellen, and Alex Imas, 2019, Mental accounting, similarity, and preferences over the timing of outcomes, Similarity, and Preferences Over the Timing of Outcomes (September 12, 2019) .
- 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.
- Gneezy, Uri, Silvia Saccardo, Marta Serra-Garcia, and Roel van Veldhuizen, 2020, Bribing the self, Games and Economic Behavior 120, 311–324.
- Gomes, Francisco, and Alexander Michaelides, 2005, Optimal life-cycle asset allocation: Understanding the empirical evidence, The Journal of Finance 60, 869–904.
- Hauser, David, Gabriele Paolacci, and Jesse J Chandler, 2018, Common concerns with MTurk as a participant pool: Evidence and solutions .
- Heimer, Rawley, and Alp Simsek, 2019, Should retail investors’ leverage be limited?, Journal of Financial Economics 132, 1–21.
- Heimer, Rawley Z., 2016, Peer Pressure: Social Interaction and the Disposition Effect, The Review of Financial Studies 29, 3177–3209.
- Hodgins, David C, Jonathan N Stea, and Jon E Grant, 2011, Gambling disorders, The Lancet 378, 1874–1884.

- 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.
- Karni, Edi, and David Schmeidler, 1991, Atemporal dynamic consistency and expected utility theory, Journal of Economic Theory 54, 401–408.
- 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.
- Kaustia, Markku, 2010, Disposition effect, Behavioral finance 169–189.
- King, Michael R, and Dagfinn Rime, 2010, The \$4 trillion question: what explains FX growth since the 2007 survey?, BIS Quarterly Review .
- Koszegi, Botond, and Filip Matejka, 2018, An attention-based theory of mental accounting, CERGE-EI mimeo .
- Laibson, David, 1997, Golden eggs and hyperbolic discounting, Quarterly Journal of Economics 112, 443–477.
- Langer, Thomas, and Martin Weber, 2008, Does commitment or feedback influence myopic loss aversion?: An experimental analysis, Journal of Economic Behavior & Organization 67, 810–819.
- Linnainmaa, Juhani T, 2010, Do limit orders alter inferences about investor performance and behavior?, The Journal of Finance 65, 1473–1506.
- Liu, Hongqi, Cameron Peng, Wei A Xiong, and Wei Xiong, 2020, Resolving the excessive trading puzzle: An integrated approach based on surveys and transactions, Technical report, National Bureau of Economic Research.

- Machina, Mark J, 1989, Dynamic consistency and non-expected utility models of choice under uncertainty, Journal of Economic Literature 27, 1622–1668.
- Mankiw, N Gregory, and Stephen P Zeldes, 1991, The consumption of stockholders and non-stockholders, Journal of financial Economics 29, 97–112.
- Martínez-Marquina, Alejandro, Muriel Niederle, and Emanuel Vespa, 2019, Failures in contingent reasoning: The role of uncertainty, American Economic Review 109, 3437–74.
- Maynard, Douglas C., and Milton D. Hakel, 1997, Effects of objective and subjective task complexity on performance, Human Performance 10, 303–330.
- Merkle, Christoph, Jan Müller-Dethard, and Martin Weber, 2020, Closing a mental account: The realization effect for gains and losses, Experimental Economics 161, 68–78.
- Nielsen, Kirby, 2019, Dynamic risk preferences under realized and paper outcomes, Journal of Economic Behavior & Organization 161, 68–78.
- O’Donoghue, Ted, and Matthew Rabin, 1999, Doing it now or later, American Economic Review 89, 103–124.
- Paolacci, Gabriele, and Jesse Chandler, 2014, Inside the turk: Understanding mechanical turk as a participant pool, Current Directions in Psychological Science 23, 184–188.
- Peköz, Erol A., 2002, Samuelson’s fallacy of large numbers and optional stopping, Journal of Risk and Insurance 69, 1–7.
- Ploner, Matteo, 2017, Hold on to it? an experimental analysis of the disposition effect., Judgment & Decision Making 12.
- 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 .

- Shefrin, Hersch, and Meir Statman, 1985, The disposition to sell winners too early and ride losers too long: Theory and evidence, The Journal of Finance 40, 777–790.
- Strub, Moris S, and Duan Li, 2020, Failing to foresee the updating of the reference point leads to time-inconsistent investment, Operations Research 68, 199–213.
- Taubinsky, Dmitry, and Alex Rees-Jones, 2018, Attention variation and welfare: theory and evidence from a tax salience experiment, The Review of Economic Studies 85, 2462–2496.
- Thaler, Richard, and Eric Johnson, 1990, Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice, Management Science 36, 643–660.
- Tversky, Amos, and Daniel Kahneman, 1992, Advances in prospect theory: Cumulative representation of uncertainty, Journal of Risk and Uncertainty 5, 297–323.
- Uhr, Charline, Steffen Meyer, and Andreas Hackethal, 2019, Smoking hot portfolios? self-control and investor decisions .
- Volij, Oscar, 1994, Dynamic consistency, consequentialism and reduction of compound lotteries, Economics Letters 46, 121–129.
- Yaari, Menahem E., 1987, The dual theory of choice under risk, Econometrica 55, 95–115.
- Zhang, Ke, and Luke Clark, 2020, Loss-chasing in gambling behaviour: neurocognitive and behavioural economic perspectives, Current Opinion in Behavioral Sciences 31, 1–7.

Appendix A. Theoretical Predictions

In this section, we formally derive the predictions of models with skewness preferences (Cumulative Prospect Theory (CPT), Rank Dependent Utility (RDU), and Expected Utility Theory (EUT)) and quasi-hyperbolic discounting for dynamic risky choice.³⁰

Cumulative Prospect Theory

Our analysis of CPT and the terminology used largely follows Barberis (2012), who originally derived the predictions of CPT in a dynamic setting with finite rounds.³¹ A decision-maker considers the gamble $L = (p_{-m}, x_{-m}; \dots; p_{-1}, x_{-1}; p_0, x_0; p_1, x_1; \dots; p_n, x_n)$, where p_i corresponds to the likelihood of attaining outcome x_i . Outcomes are ordered such that $-m \dots -1$ correspond to those below the reference point x_0 , here assumed to be the status quo, and $1 \dots n$ 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:

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

where

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

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

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

³⁰ 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). In turn, one could think of our theoretical exercise as highlighting the necessary features for matching the empirical findings.

³¹ As outlined below, we depart from Barberis (2012) in the way that we formulate ex-ante plans—an extension that allows us to link the theoretical results closer to our empirical setting.

and value function

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

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

In our setting, let $L = (1/2, -10; 1/2, 10)$. The decision-maker faces a sequence of choices. In each round t , 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. 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 $a \in \{\textit{continue}, \textit{exit}\}$.

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 opportunity, 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.

In a departure from Barberis (2012), we restrict the agents' initial plans to the subset of plans that can be described by a pair of limits—a gain and a loss limit—as is the case in our experimental design. Furthermore, we run simulations for a large number of rounds which allows us to determine the number of rounds that a representative CPT agent would require in order

to accept the first gamble in our empirical setting. We find that this number is 26 rounds.³²³³ Thus we examine the behavioral predictions given a dynamic sequence of 26 potential gambles, such that $t \in \{0, \dots, 26\}$.

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). 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 take on risk 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.³⁴ This assumption is consistent with prior experimental evidence (see Imas, 2016).

Results (CPT)

³² Barberis (2012) shows that a representative CPT agent will require a minimum of 26 rounds to accept the first gamble in a sequence, with a plan to exit as soon as she endures any losses and to continue as long as she earns gains. It is possible, however, that this plan is not exactly the optimal plan for a representative CPT agent and she would accept the sequential gamble as part of a different plan with a lower number of rounds. The fact that our analysis results in the same minimum round requirement suggests that this particular loss-exit plan is very close to optimal.

³³ Previous experimental studies by Andrade and Iyer (2009), Barkan and Busemeyer (1999), and Ploner (2017) analyze dynamics of risk-taking over fewer than four rounds. This small number of rounds makes it difficult to differentiate between different models of decision-making, and may explain the inconsistent results regarding participants' planned choices.

³⁴ We present results when this assumption is relaxed below.

Figure A1, 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 $\lambda = 2.25$ —the median estimates in Tversky and Kahneman (1992). Note that in the presence of dynamic inconsistency, the ex-ante decision depends not only on the preference parameters but also on agent’s sophistication and the availability of commitment opportunities. 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 opportunity at their disposal. 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 opportunities. Several findings are obtained.

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, for agents who weigh probabilities close to linearly and have high levels of diminishing sensitivity ($\alpha \ll 1$), 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.³⁵

Figure A1, 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 end up chasing the losses further by accepting subsequent gambles. These agents also 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

³⁵ 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 suggests that our experimental design choice to simplify the strategy elicitation is not very restrictive.

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.³⁶

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 naïve 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 in some manner 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 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.³⁷ In

³⁶ 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.

³⁷ Similarly, one could assume that an agent is naïve about future reference point updating, making ex-ante choices as if her reference point did not update until the final round. Strub and Li (2020) study this problem in a prospect theory framework without diminishing sensitivity. The authors show that while this type of naïveté does lead to dynamic inconsistency, the deviation in ex-post behavior is outcome-independent.

general, the closer an agent is to the white-colored area in Figure A1, 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.

Similarly, the assumption of a finite planning horizon is important as well. In contrast to the prediction of early exit in the gain domain, theoretical work on the dynamics of CPT under an infinite planning horizon predicts that agents continue gambling ad infinitum, or until all their wealth is spent, independent of the outcome (Ebert and Strack, 2015).³⁸ This result is driven by the fact that if the planning horizon is long enough, agents can always generate a strategy with enough skewness to justify taking on more risk, irrespective of whether she is in the gain or loss domain. A finite planning horizon, in contrast, restricts the potential skewness of a dynamic “loss-exit” strategy. As the number of remaining investment decisions decreases, so does the CPT agent’s willingness to continue gambling if she is in the gain domain. It is worth noting that traders in our brokerage data behave according to the dynamic predictions of CPT with finite horizon, even though there are no binding restrictions on how long they can hold open positions. The reason for this is likely the fact that the relevant horizon in dynamic models of CPT corresponds to the period before the reference point resets. Even in contexts where objective time horizon is long, psychological factors that lead to reference point resetting result in substantially shorter time horizons in practice. The prevalence of such psychological factors—e.g. the realization of gains and losses (Imas, 2016; Barberis and Xiong, 2012), temporal markers (e.g., end of the week, Dai, Milkman, and Riis, 2014), and attention-based narrow bracketing (Koszegi and Matejka, 2018; Evers and Imas, 2019)—suggests that a finite time horizon may be appropriate in many real world settings.

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:

Namely, people stop taking on risk earlier than anticipated, but contrary to our empirical findings, do so to the same extent after gains and losses.

³⁸ Ebert and Strack (2015) show that the same result holds with a finite horizon if at any given point the outcome of a gamble can be arbitrarily large.

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

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

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

and

$$\pi_i^{RDU} = w(p_{-m} + \dots + p_i) - w(p_{-m} + \dots + p_{i-1}) \quad (\text{A7})$$

For consistency, we assume the probability weighting function w is the same as in A2.

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 A2, 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 A2 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.

Figure A2, 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 A2). 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; Camerer and Ho, 1994).³⁹ ‘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 A2. 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.⁴⁰

Quasi-Hyperbolic Discounting

We follow Laibson (1997) and assume a quasi-hyperbolic discount function. The expected

³⁹ This result is dependent on the assumption about initial wealth. Assuming initial wealth \$1,000, all agents in Figure A2 who accept the initial gamble will continue until the end, independent of the outcome.

⁴⁰ 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 theoretical and empirical settings.

utility of the agent is assumed to be:

$$U_t = \mathbb{E}_t \left[u(c_t) + \beta \sum_{\tau=1}^{T-t} \delta^\tau u(c_{t+\tau}) \right] \quad (\text{A8})$$

where t denotes the round number and the utility function $u(c_t)$ is given by Equation A6. We assume that the agent consumes her entire wealth as soon as she stops investing. If the agent decides not to start investing at all, consumption takes place before the first round.⁴¹ Other than that, the simulations for quasi-hyperbolic discounting are done in the same way as those for CPT and RDU. Similar to RDU, we need to make an assumption about initial wealth before the agent enters the experiment. We obtain results for an initial wealth equal to the agent's endowment of \$2.6 and an initial wealth of \$1,000. We assume a discount factor $\delta = 0.97$ (Laibson, 1997) and simulate the investment decisions for all parameter combinations of $\gamma \in [0, 3.5]$ and $\beta \in [0.6, 1]$.

We find that an agents with quasi-hyperbolic discounting would not accept the first gamble in a sequence. Furthermore, even if the agent was forced to enter, the ex-ante plan for all parameter combinations implies exiting as soon as possible, independent of the outcome. It follows trivially that quasi-hyperbolic discounting does not predict the type of dynamic inconsistency observed in our experiment.

To understand why, note that an agent with a discount rate of 1 (i.e., $\beta = 1$ & $\delta = 1$) would not accept the first gamble in a sequence because the prospective endogenous skewness is not enough to compensate for the risk, as discussed above for EUT. Setting $\beta < 1$ or $\delta < 1$ introduces further reasons to reject the first gamble in a sequence. In a dynamic environment, extreme outcomes take time to accumulate but $\delta < 1$ makes these extreme outcomes less attractive. Consequently, an agent with skewness preferences and $\delta < 1$ would require a higher prospective skewness than an agent with $\delta = 1$ in order to start taking risk. In addition, $\beta < 1$ makes the risk-free option to reject the sequence of gambles more attractive because it implies immediate consumption.

⁴¹ Please note that under an alternative assumption that all wealth is consumed in the final round, the problem is reduced to the case of EUT, which is outlined above.

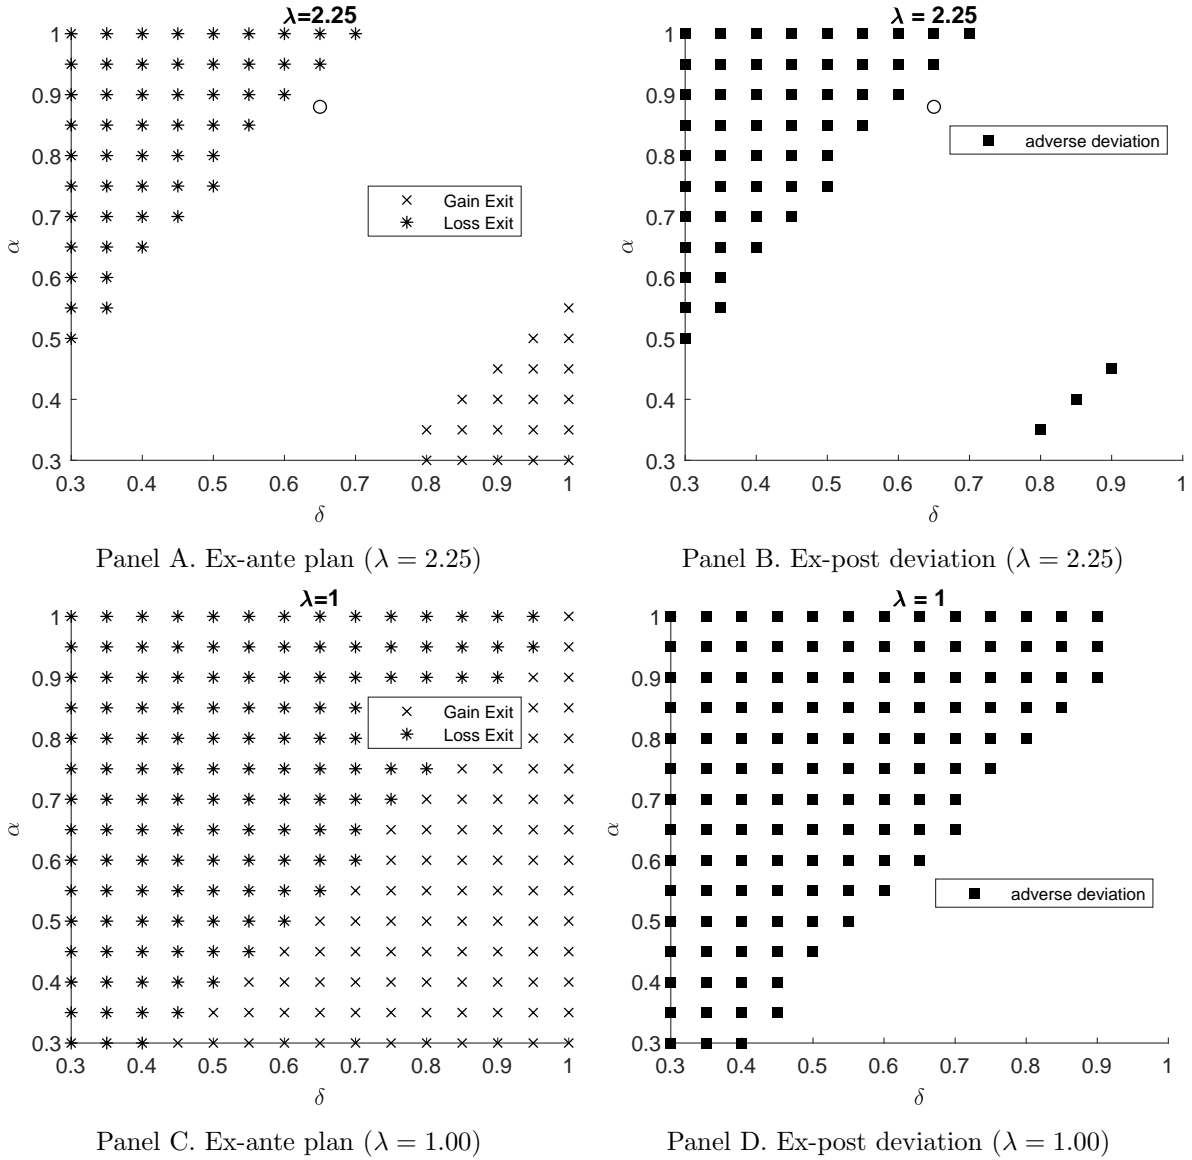

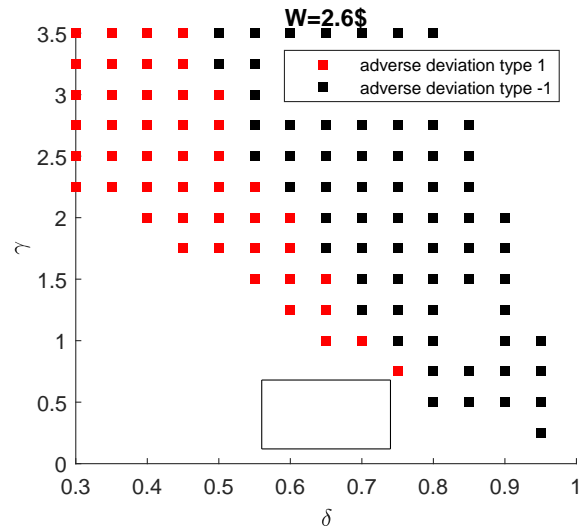


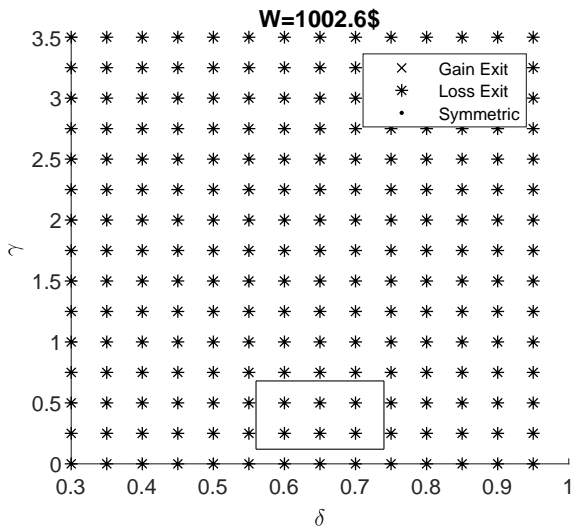
Figure A1. Theoretical predictions of Cumulative Prospect Theory. This figure illustrates the ex-ante plan (Panel A and C) and the ex-post deviation (Panel B and D) of agents with CPT preferences given different levels of probability weighting (δ) and diminishing sensitivity (α), i.e., the extent to 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 valuation than rejecting the gamble at the beginning, in $t = 0$. The representative agent is marked by a circle in Panels A and B.



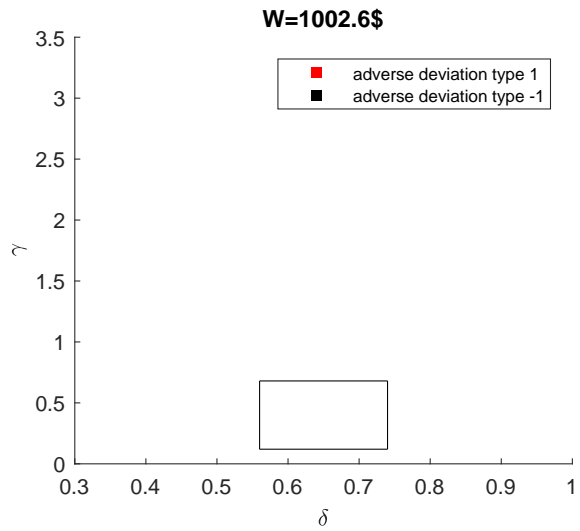
Panel A. Ex-ante plan ($W = 2.6\$$)



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 A2. 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 denoted as ‘type -1’. W indicates the assumption about initial wealth. The parameter-region estimated from previous experimental studies is marked by the box.

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 Hake (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 not. 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.⁴²

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]

⁴² In a further experiment, which was designed to test for the role of biased beliefs on the investment decision, we included a modified Hard Plan treatment. This treatment contained the following note here: “Imagine you could repeat this investment decision using the same gain and loss limits over and over again. The average outcome from all of your decisions will be zero regardless of how many times you repeat the investment. An average of zero does not mean that you won’t sometimes win and sometimes lose.”

Appendix CIV. Outcome Frame

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. After finding out the outcome of the first investment, 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. After either your loss or gain limit is reached, you will stop investing automatically and this part of the experiment is over.

- 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. 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.

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.

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 invest? [Yes]/[No]

Screen: Round 1 of 26: Results

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

Screen: Your gain and loss limits

Please indicate your loss limit and your gain limit. You will keep investing 10 cents in each subsequent round until your total gain or loss (including round 1) reaches your gain limit or loss limit respectively. After either your loss or gain limit is reached, you will stop investing automatically.

Loss limit (in cents):...

Gain limit (in cents):...