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

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

This draft: January 23, 2023

ABSTRACT

We document a robust dynamic inconsistency in risky choice. Using a unique brokerage dataset and a series of experiments, we compare people's initial risk-taking plans to their subsequent decisions. Across settings, people accept risk as part of a "loss-exit" strategy—planning to continue taking risk after gains and stopping after losses. Actual behavior deviates from initial strategies by cutting gains early and chasing losses. More people 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 explore implications for theory and welfare.

Rawley Heimer is at the W.P. Carey School of Business, Arizona State University (e-mail: rheimer@asu.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, Michaela Pagel, Collin Raymond, Alex Rees-Jones, Hersh Shefrin, Kelly Shue, Dmitry Taubinsky, George Wu, and Florian Zimmermann for helpful suggestions and illuminating discussion. We also thank seminar participants at Arizona State University, Babson, Cornell, Florida State University, Middlebury College, the Radboud University Nijmegen, Simon Fraser University, the University of Colorado, Boulder, University of Bonn, University of Miami, the University of Washington, and Wharton, as well as participants at American Finance Annual Meetings, 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, Yanghua Shi and Christy Kang 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

This paper studies risk-taking in dynamic environments. Unlike the case of static settings, peoples' behavior in dynamic settings is determined both by their planned and actual choices in response to experiencing gains and losses. For example, a person may purchase a stock 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 depends on whether the expected outcomes of a given strategy are more attractive than a safer alternative. The person's actual choices after experiencing gains and losses may follow her plan—or they may not. The same investor who planned to sell the stock after a loss may do so, or she may deviate from her strategy and double down. Documenting a discrepancy between planned and actual choices has implications for theory, welfare, and the interpretation of prior empirical findings on risk-taking in dynamic environments such as the disposition effect (Shefrin and Statman, 1985; Odean, 1998; Weber and Camerer, 1998).

We study what strategies motivate people to start taking risk in dynamic settings and whether their actual behavior follows these strategies. Using evidence from both the lab and field, we identify a robust dynamic inconsistency in risk-taking. People start taking risk—what we term the 'entry decision'—as part of a "loss-exit" strategy that involves continuing to take risk after gains and stopping early after losses. Notably, "loss-exit" strategies generate a more positively-skewed outcome distribution than any individual gamble in isolation. People's actual choices deviate dramatically from their plans, exhibiting "gain-exit" behavior that cuts gains early and chases losses. Notably, some appear aware of their dynamic inconsistency: more people begin to take on risk when they can commit to their initial strategy than when they cannot.

Interpreting the behavioral data through the lens of theory helps rationalize a discrepancy in risk-taking between static and dynamic environments. Prior research shows seemingly anomalous levels of risk aversion in static settings (Kahneman and Tversky, 1979). This contrasts with the high levels risk-seeking observed in otherwise similar dynamic settings (Imas, 2016; Barber and Odean, 2000). Indeed, we find that people are 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 then proceed to formally compare different models of risky choice. We find that a dynamic framework featuring probability weighting, reference dependence, and diminishing sensitivity such as Cumulative Prospect Theory (Tversky and Kahneman, 1992) is most consistent with the observed behavioral patterns. As Barberis (2012) demonstrates theoretically, the discrepancy between static and dynamic environments is not driven by the intrinsic nature of the environments per se, but by the fact that people's "loss-exit" strategies in the latter can generate a level of positive skew that is unavailable in the static case. Probability weighting leads to the overweighting of low probability outcomes, which leads a greater willingness to accept a risky bet as part of a "loss-exit" strategy than the same bet in isolation. Our experiments explicitly test this prediction. Finally, we explore the welfare consequences of dynamic inconsistency in risk-taking and show that the deviations in ex-post behavior have significant utility costs.

We begin by examining the dynamics of risk-taking in the field 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.¹

We document a significant discrepancy between people's planned and actual risk-taking 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. Using historical price series in the same settings, we use simulations to show that traders' initial strategies imply a return distribution with positive skew. However, traders' subsequent choices

¹ As a general example, take an individual who places a loss limit of 10% and a gain limit of 20%. This corresponds to a risk-taking strategy that pairs a willingness to lose 10% for the chance of gaining 20%. The individual 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 choosing to abstain from risk in that position before the limits are hit, e.g., selling the asset after a 5% gain.

follow a "gain-exit" strategy—the opposite of their intended plans. Experiencing losses leads traders to revise their limits further downward, while experiencing gains leads them to sell the assets before the limit is hit; by the time a position is closed, the vast majority of gain limits are smaller than the loss limits. This behavior leads to a realized return distribution that is *negatively* skewed. Importantly, though traders' initial strategies suggest that they would prefer to hold winners longer than losers, their ex-post behavior is consistent with the disposition effect, whereby they hold losers longer than winners (Shefrin and Statman, 1985; Odean, 1998).

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. 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 an experiment with the following features: (i) the ability to elicit incentivized ex-ante strategies and compare them to ex-post behavior, (ii) the elicitation of 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 a series of experiments (N = 2621), participants are offered the choice to accept or reject a *fair* symmetric gamble either in isolation or as part of a sequence. They are provided feedback after every decision and have the choice to stop anytime. The experiment ends either when the next gamble in the sequence is rejected or the end of the sequence is reached. In the One-Shot treatment, participants decide to accept or reject a single gamble. In the Sequential treatment, the same gamble is offered as part of a long finite sequence. In the Hard Plan treatment, participants face the same sequence of gambles but report their risk-taking strategies before deciding whether or not to accept the first gamble. Similar to the field setting, we elicit strategies in the form of loss (gain) limits that correspond to the most a participant is willing to lose (gain) before rejecting further risk. In addition to being intuitive and easy to explain, these limits are sufficient to characterize participants' risk-taking plans under mild assumptions. Importantly, the gain and loss limits are binding: the participants cannot revise their strategy after accepting risk. This allows us to ensure that participants are incentivized to report their preferred ex-ante strategies. We find that people 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.² More than 80% of participants' exante strategies can be classified as "loss-exit" plans; strikingly, the average participant initially accepts risk with a gain limit that is more than 3 times higher than her loss limit. Only 7% of strategies can be classified as "gain-exit." In contrast, participants' actual choices follow the reverse pattern: they are significantly more likely to stop after winning than after losing, replicating the behavioral pattern we observe in the field. Looking at the implied and realized outcome distributions, we find that participants' ex-ante strategies generate positive skew but their actual choices generate negative skew, with the difference in skew being both statistically and economically significant.

The experimental treatments also allow us to examine whether people are aware of their dynamic inconsistency.³ We find that a significant proportion of participants are indeed sophisticated: people are more likely to begin taking risk when they are provided with a commitment opportunity. However, Section VI presents results that suggest limits to this sophistication. In a separate experimental treatment, participants report their gain and loss limits similar to the Hard Plan treatment but these limits are non-binding—people can deviate from their initial strategies once their limits are reached. Most participants indeed deviate from their strategies in a manner similar to when no commitment is available. At the same time, they are equally likely to take on risk as when the strategies are binding. This suggests that people may overestimate the efficacy of 'soft,' non-binding commitment opportunities in disciplining their behavior.

To better understand the mechanism driving our empirical findings, Appendix A derives the dynamic predictions of several models of risk-taking. We show that the observed behavioral data cannot be rationalized by models that incorporate only diminishing sensitivity (Expected Utility Theory, EUT thereafter) or only probability weighting without reference dependence

² 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).

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'.

(Rank Dependent Utility, RDU thereafter; Quiggin, 1982; Yaari, 1987).⁴ As also shown in Barberis (2012), a dynamic version of Cumulative Prospect Theory (CPT, thereafter) which incorporates these features generates both the pattern of planned choices and deviation in actual choices, as well as the demand for commitment that we observe in our data. The framework also predicts that the same person will reject a single fair gamble while accepting the same gamble as part of a dynamic sequence. Importantly, the predicted discrepancy in risk-taking is not due to the dynamic environment per se; rather, people are more likely to accept a gamble as part of a sequence because the dynamic setting allows them to generate a level of positive skew over the outcome distribution that is unavailable in the single gamble.

To test this proposition directly, we ran a separate study that took out the dynamic element but allowed people to match the skew generated by participants' strategies in the main study. Specifically, we created two one-shot gambles that either matched the skew generated by participants' initial strategies (positive skew) or ex-post behavior (negative skew). We then presented a separate group of participants with a choice between keeping a certain monetary payoff or selecting one of two one-shot gambles. Consistent with the theoretical framework, participants overwhelmingly chose the positively skewed gamble over both the negatively-skewed gamble and the certain payoff.

Assessing the welfare consequences of dynamic inconsistency requires establishing a normative benchmark. This is not straightforward in the case of risky choice because both planned and actual decisions appear to be driven by non-normative factors. Such normative ambiguity requires additional data to characterize the welfare-relevant domain. Drawing on the behavioral welfare framework outlined in Bernheim and Taubinsky (2018), we ran 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, which suggests that ex-post decisions correspond to costly deviations. In Appendix B, we use simulations to quantify the welfare consequences of dynamic inconsistency for a broad range of parameter combinations, finding them to be substantial. Abstracting away from our attempt to assess wel-

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).

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.

fare 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 an important 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 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 the hypothesis introduced by Barberis (2012) that the disposition effect is inconsistent with traders' ex-ante preferences and therefore has welfare consequences beyond financial costs.⁶

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 nonbinding. As discussed further in Section VI, soft commitment can lead a substantial fraction of individuals to accept risk that they would have otherwise avoided. The observed dynamic inconsistency implies non-binding limits may potentially be welfare reducing in some settings.

Related Literature Our paper contributes to two strands of literature on the dynamics of risky choice. On the theoretical side, prior work has shown that while the standard expected

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.

utility framework predicts that planned and actual choices will be the same, dynamic inconsistency can potentially emerge from many models of non-expected utility (e.g., Karni and Safra (1990). In his influential survey of the decision-theory literature, Machina (1989) 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). Another approach to addressing dynamic inconsistency in non-expected utility models includes theories of *resolute choice*, where the agent sticks to her ex-ante plan after gains and losses despite preferring a different choice if the prospective choices were presented to her *de novo* (McClennen, 1988).⁷

In work most closely related to our own, Barberis (2012) and Ebert and Strack (2015) present the predictions of CPT in dynamic environments. Barberis (2012) studies a discrete time setting with symmetric gambles and a finite number of rounds. He derives two main sets of predictions. First, the ability to take risk as part of a "loss-exit" plan leads a CPT agent to accept a gamble as part of a dynamic sequence that she would otherwise reject in isolation. This follows because the strategy increases the positive skew relative to the single gamble, which increases the prospect's attractiveness due to probability weighting. Second, the CPT agent exhibits an outcome-dependent dynamic inconsistency—stopping earlier (later) than planned after gains (losses). Ebert and Strack (2015) investigate the case with an infinite horizon in continuous time and predict that CPT agents will take on risk until bankruptcy regardless of whether they see gains or losses. Ebert (2020) shows that agents will accept gambles even if they are negatively skewed as long as the horizon is long enough for them to generate positive skew through their plans.

On the empirical side, our findings contribute to the literature on dynamic choice under uncertainty more broadly. In one of the first empirical papers on the topic, Thaler and Johnson (1990) find that people take on more risk after a loss, but only if the upside of the gamble allows them to recover from it and get back to the reference point (the "break-even" effect).⁸

⁷ McClennen (1988) refers to resolute choice as an *instrumental* theory, where the agent decides to stick to her ex-ante plan because of personal rules and norms.

⁸ A line of work on the description-experience gap in risky choice presents participants with a set of prospects that have unknown payoffs and probabilities. Participants select between prospects and are given feedback based on their choices. A series of papers finds that people are less sensitive to low probability events when the relevant information needs to be learned than when likelihoods are given (see Hertwig and Erev (2009) for review). This work does not examine dynamic (in)consistency nor how prior outcomes per se affect risky choice.

Some follow-up work has confirmed the increased risk-taking after losses (Langer and Weber, 2008), while other papers have found the opposite—that people take on less risk after losses (Shiv et al., 2005). Imas (2016) suggests that these seemingly-contradictory findings can be reconciled by distinguishing between paper and realized outcomes. In his framework, paper losses are bracketed together with prospects and lead to greater risk-taking; realized losses reset the reference point and lead to less risk-taking. A series of experiments provides support for these predictions.⁹

Looking at dynamic (in)consistency, Cubitt, Starmer, and Sugden (1998) outline four principles—separability, timing independence, frame independence, and reduction of compound *lotteries*—which can be used to characterize the properties of dynamic choice under EUT. They present experimental tests of these principles in the context of the common ratio effect (Allais, 1953). Of the four, only timing independence is violated, which suggests a dynamic inconsistency in decision-making. Notably, however, the outcomes of prior gambles are not realized before participants make their choices. Looking at outcome-dependent deviations, Ploner (2017) and Andrade and Iyer (2009) find that people plan to bet more after a gain than a loss, while Barkan and Busemeyer (2003) find that people plan to bet more after a loss than a gain. In these papers, gambling is either enforced (both Ploner (2017) and Barkan and Busemeyer (2003) require participants to take risk in the first round), or planned choices are non-binding and can be revised (Andrade and Iyer, 2009). These experiments feature two rounds of decisions; after observing the outcome of the lottery, participants end up taking more risk after a loss than after a gain. Imas (2016) uses a within-subject hypothetical planning stage to compare planned and actual behavior in the domain of losses across four rounds. In contemporaneous work, Dertwinkel-Kalt, Frey, and Köster (2020) study strategies in a dynamic optimal stopping problem through the lens of Salience Theory (Bordalo, Gennaioli, and Shleifer, 2012). They elicit hypothetical, within-subject gain and loss limits that can be revised once they are hit. Participants end up taking on more risk after losses than their non-binding plans imply, but the study design does not permit full characterization of the ex-post strategy (e.g., that actual behavior is "gain-exit").

⁹ In Appendix A we discuss the importance of bracketing for the dynamic predictions of CPT and how it relates to both our theoretical and empirical framework.

We contribute to this literature by providing lab and field evidence for "loss-exit" strategies and compare planned behavior with ex-post choices in the same setting. Our between-subject design allows us to elicit both ex-ante strategies and ex-post choices in a fully incentivized manner. We also directly examine risk-taking as a function of the time horizon. Together, our empirical setting allows us to identify and fully characterize dynamic inconsistency in risky choice and interpret it through a theoretical lens.

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 (Laibson, 1997) and empirically (DellaVigna and Malmendier, 2006). Hyperbolic or quasi-hyperbolic discounting is the primary mechanism proposed to explain dynamic inconsistency in the intertemporal choice literature (O'Donoghue and Rabin, 1999). Note that this mechanism is conceptually distinct from the ones we argue can best rationalize the findings in dynamic risky choice.¹⁰

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 IV.C discusses them in light of theory. 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 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

¹⁰ Unlike a dynamic framework with probability weighting and diminishing sensitivity, models of hyperbolic or quasi-hyperbolic discounting 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, these models would not predict the discrepancy between "loss-exit" strategies and "gain-exit" behavior observed in our data. As demonstrated in Appendix A, quasi-hyperbolic discounting predicts that people would not accept risk in our dynamic setting, and even conditional on entry, the majority would not have outcome-dependent ex-post deviation from the ex-ante plan.

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 (CFDs), 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 currency pairs (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), has shown that traders in CFD 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 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.¹¹ 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 open a position 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 the position before the limits are hit.

¹¹ Linnainmaa (2010) also studies how limit orders affect the trading outcomes of individual investors. However, he studies the *choice* between market and limit orders using a data set of Finnish individual investors.

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. The short holding period on individual trades also implies that traders are unlikely to hold many open positions at the same time. Therefore, portfolio effects are unlikely to drive the behavior that we observe on each individual transaction. Second, the foreign exchange CFD market has 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 and myopic loss aversion as in Benartzi and Thaler (1995)—a motive that requires risk to have a positive expected value.¹²

A. Results

We start by defining a trader's risk-taking strategy as the combination of the gain and loss limits 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 a relative majority (46%) of traders use "loss-exit" as their modal strategy when deciding to open a new position (see Figure 1). Substantially fewer adopt "gain-exit"

¹² Here, return aggregation refers to the property that a positive expected-value risky asset is more likely to generate a loss over a shorter time horizon than a longer one. For example, a gamble with equal chances of yielding +2 and -1 has a 1/2 chance of generating a loss in one play. But the probability of a loss goes down to 1/4 after two plays, 1/8 after three plays, etc. Myopic loss aversion states that a loss-averse trader will be more willing to accept positive expected-value risk if she checks performance infrequently, since returns will be aggregated for longer and she will be less likely to see a loss.

(32%) and "neutral" (22%) as their modal strategies.¹³ Moreover, traders' preference for "lossexit" strategies is robust to the characteristics of the trade and trader attributes (see Appendix D Table DI). For instance, traders are most likely to use "loss-exit" strategies even after controlling for trade-specific characteristics such as the position's leverage, direction, capital, and instrument. Despite the negative expected returns in this setting, traders may be overconfident and thus take on risk because they *believe* the returns to be positive in expectation. The fact that trading experience does not appear to impact behavior is suggestive evidence against this explanation. Because more experienced traders should have learned about the non-positive expected returns of CFD trading, the combination of overconfidence and myopic loss aversion is unlikely to drive our results. That being said, overconfidence cannot be ruled out without access to traders' beliefs; this explanation is accounted for in our experimental settings because people are given complete information about the return distributions and experience them directly.

Next, we examine how traders' ex-post behavior 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 2 presents traders' behavior in response to paper (i.e., unrealized) gains and losses.¹⁴ 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.¹⁵ 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 such revisions of loss limits nearly 20 percentage points more often than the next most

 $^{^{13}}$ All differences are significant at the 1% level.

¹⁴ 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.

¹⁵ Forty-six percent of positions are closed without any revisions, which can either correspond to the limits being hit or a manual exit.



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. An "equidistant" strategy implies equidistant limits.

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; they 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 nearly identical to the full sample, suggesting that actual choices follow the opposite pattern from planned choices for the same position.

Figure 3 highlights the sizable magnitude of the reversal from "loss-exit" to "gain-exit" strategies. Panel A presents the difference between the gain and loss limits when traders open their positions, and Panel B when traders close their positions. Specifically, we calculate the percentage distance between the price of the asset when opening the position and the unlevered initial gain limit (loss limit). We then bin the average percentages into centiles, and take the



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" strategy (i.e., initial loss limit closer to opening price than initial gain limit).

difference between the gain limit and loss limit at each centile. A value greater than zero indicates that traders' orders are tilted toward a "loss-exit" strategy. We use a bin-scatter plot

14

to present the data because of the large number of trades, as well as to standardize the overall magnitude of risk-taking on each position (the bin-scatter helps us compare, for example, a position with a 50% gain limit and 100% loss limit to one with a 1% gain limit and 2% loss limit, both of which have loss-limits that are twice as large as the gain limits.).

Panel A shows that traders choose initial "loss-exit" strategies at nearly all points in the distribution; the difference between limits is positive at nearly every centile. The median position has a gain limit that is six basis points larger than the size of the loss limit. To gauge the impact of these limits on returns it is necessary to also consider the amount of leverage used. Taking leverage into account, the median position's gain limit is 394 basis points larger than the loss limit. However, traders reverse their strategies while the position is active. Panel B shows that the difference between the revised gain limit and the revised loss limit is *negative* by the time traders close their positions. Notably, traders push the loss limit orders substantially downward. Taking leverage into account, the median position has a revised loss limit that is 546 basis points larger than the gain limit.

The differences in ex-ante and ex-post strategies should be reflected in the implied return distributions of those strategies. The ex-ante expected distribution of returns can be estimated for each trader from her initial strategy under additional assumptions about the stochastic process of the underlying contracts. To this end, we use the sample distribution of intraday returns of each underlying asset over the sample period and run simulations to estimate the expected outcome distribution of every position. In every iteration, we draw randomly and independently with replacement from the sample return distribution of the underlying CFD until either of the limits are triggered. We aggregate the resulting binary position-level return distributions by equally-weighting all positions with non-equidistant limits. The resulting aggregate expected return distribution has a positive skew of 0.3 (p < 0.01).

The distribution of returns that traders ultimately realize is impacted in the reversal of their strategies. Figure 4 presents a bin-scatter of traders' realized returns (excluding leverage) on each position. The distribution of returns has a negative skew of -3.29 (p < 0.01).¹⁶ These results show that the reversal in traders' ex-ante versus ex-post strategies leads to a reversal in the skewness of expected versus realized returns.

 $^{^{16}}$ $\,$ Although less than 40% of positions are closed at a loss, the losses are much larger in magnitude than the gains.



Panel A. Limit orders when *opening* positions



Panel B. Orders when *closing* positions

Figure 3. Limit order modifications while positions are open. This figure presents the unlevered difference in magnitudes for gain and loss limits. We calculate the percentage distance between the price of the asset when opening the position and the initial gain limit or loss limit. We then bin the percentages, taking the average, into centiles, and take the difference between the gain and loss limits at each centile. Panel A presents the difference between the gain and loss limits when traders *open* their positions, and Panel B presents the difference when traders *close* their positions. Both figures exclude positions with equidistant limits and the largest 5% in magnitude of both types of orders. Though the figure uses unlevered returns, the median position is opened with 100x leverage.



Figure 4. Realized returns. This figure plots the cumulative distribution of traders' realized returns on each position. The returns are calculated as percentages excluding leverage.

In sum, we find that traders allow larger losses to accumulate and realize gains 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. While this suggests a dynamic inconsistency, several aspects of the field setting such as traders' potential uncertainty about the data generating process and the (known) ability to revise limits precludes a clean test. We developed an experimental design that allows us to identify a dynamic inconsistency between planned versus actual behavior and to formally distinguish between different models of risky choice. A total of 2621 participants took part in the study. This section describes the main experiment; a full write-up of the robustness treatments and extensions can be found in Appendix C.

Participants (N = 295) were recruited from Prolific and received approximately \$60 per hour for their participation in the experiment.¹⁷ We designed the experiment to: (i) elicit

¹⁷ Online platforms, such as Prolific, are increasingly used in economics and the broader social sciences to recruit subjects for experiments. Studies have shown that laboratory results broadly replicate on these

incentivized initial strategies and compare them to subsequent behavior, (ii) examine 'entry' decisions—the initial choice to take on risk or not—as a function of the number of rounds and access to commitment opportunities, and (iii) study subjects' sophistication about dynamic inconsistency given the availability of commitment opportunities.

People faced binary choices of whether or not to invest 50 cents in fair symmetric gambles that had an expected value of zero. Each gamble featured a simple 50/50 chance that the investment either doubled or was lost. Participants could not invest more than the assigned per-round amount of 50 cents; those who chose not to invest would get to keep the amount. Participants gained experience about the data generating process by drawing ten observations from a stratified sample before making a choice.¹⁸

The experiment consisted of three between-subject treatments that are outlined in Figure 5 below.¹⁹ In the One-Shot treatment, participants received an endowment of 50 cents and decided whether or not to invest in a single gamble. In the Sequential treatment, participants received a total endowment of \$13. They then had to decide whether or not to invest 50 cents in the first round; we refer to this as the 'entry' decision. If the participant chose to invest, the first round outcome was revealed (gain or loss) and she decided whether to invest in the next round, and so on, for a maximum of 26 rounds.²⁰ Once a participant chose to stop investing, the experiment would end and she would keep her remaining endowment. Feedback on previous outcomes and the total accumulated gains and losses were provided between each round.

In contrast to the Sequential treatment, participants in the Hard Plan treatment first chose a risk-taking strategy over a maximum of 26 rounds. Specifically, we elicited participants' plans 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,

online platforms (e.g., Horton, Rand, and Zeckhauser, 2011; Snowberg and Yariv, 2021; Gupta, Rigott, and Wilson, 2021). The hourly compensation was well above the the typical amount offered by experiments on Prolific. Appendix C shows that the results replicate when the more standard average remuneration of approximately \$12 per hour is used.

¹⁸ Christine, Weber, and Haisley (2013) and Hogarth and Soyer (2015) provide evidence on the benefits of sampling for the understanding of probabilities.

¹⁹ As highlighted in Cubitt, Starmer, and Sugden (1998), a between-subject design allows for identification of dynamic inconsistency under a minimal set of assumptions compared to a within-subject designs.

²⁰ We chose a finite sequence of 26 rounds for theoretical considerations because it allows us to differentiate between different models of dynamic decision-making under risk.

respectively, correspond to the outcome "equidistant" 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 decided whether or not to accept the first gamble as part of their risk-taking strategy. We informed participants that they would automatically stop investing if either of their limits were triggered. This ensured that the limits were fully incentivized and binding.²¹



Figure 5. Experimental design and sample size.

The experiment was programmed in Qualtrics and administered it on Prolific in July 2021. We use sample selection criteria based on location (US, UK or Ireland) and approval rate (> 97%). Before the beginning of the experiment, subjects underwent a bot check, as well as attention and comprehension checks (see Appendix E). Only subjects who passed all checks were allowed to proceed with the study. Participants also filled out a questionnaire about demographic characteristics (age, gender, study field, highest level of education) and statistical skills (self-reported) before beginning the investment task. Fewer than 5% of participants exited the experiment after being randomized into treatments, which assuages concerns about selective attrition.

²¹ Fischbacher, Hoffmann, and Schudy (2017) propose optional and amendable gain and loss limits as an intervention to reduce the disposition effect. In contrast to their study, the limits in our experiments are mandatory and binding to prevent any selection effects and ensure that ex-ante strategies were incentivized.

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.

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 6 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—by at least 20 percentage points—in the One-Shot treatment compared to any of the multi-round treatments. The differences in entry rates between the One-Shot and the Sequential treatment as well as between the One-Shot and the Hard Plan treatment are significant (Mann-Whitney test, p < 0.01 in each test). Participants are more likely to start taking risk as a part of a dynamic sequence than in isolation.²²

Next, we compare the entry rates between the Sequential and Hard Plan treatments. As shown in Figure 6, participants are more likely to accept the multi-round gamble if they can commit to an ex-ante strategy. We find that participants are 13 percentage points more likely to initially accept risk if they can commit to an ex-ante strategy (Mann-Whitney test, p < 0.01), which suggests a significant demand for commitment.

A.1. Robustness

This section explores a series of potential explanations for the results described above. We first examine the possibility that the higher entry rates in the multi-round treatments are 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 outcome distribution and communicating it in a simple, straightforward way—both by description and through experi-

²² In the Appendix C, we report the results from separate studies where endowments between the One-Shot and multi-round treatments are equalized. The difference in entry rates is robust to this modification, suggesting that wealth effects cannot explain it.



Figure 6. Entry decision. This figure shows the percentage of participants in each treatment who accept risk in the first round.

enced 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.²³ 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.

Next, we look at whether participants' perceptions of task complexity can explain the higher levels of risk-taking in the Hard Plan treatment compared to the Sequential treatment. In particular, one may argue that because the Hard Plan treatment restricts the strategy-choice set

²³ See Doob's optional sampling theorem.

to a subset of strategies that can be described by a pair of limits, the perceived task-complexity is lower than the task-complexity in the Sequential treatment. Under this assumption, taskcomplexity may explain why participants in the Hard Plan treatment have a higher propensity to enter than participants in the Sequential treatment. To test this claim, in one experiment we elicited the perceived complexity of the main task. Results show that perceived complexity is unlikely to drive the higher entry rates in the Hard Plan treatment—if anything, perceived complexity was higher in the Hard Plan treatment than in the Sequential treatment.²⁴

A related argument is that the higher entry rate in the Hard Plan treatment may be caused by decision aversion as the number of choices in the Hard Plan treatment (3) is on average smaller than in the Sequential treatment. However, this implies that the entry rates should be highest in the treatment with the lowest number of decisions—the One-Shot treatment—which is inconsistent with our results.

It may be the case that a sunk cost fallacy with respect to the cognitive effort required to make an explicit ex-ante plan drove the differences in entry rates. In a separate study, which is described in detail in Section V, we reversed 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 found no significant difference in the entry rates between the baseline Hard Plan treatment (95.39%) and the treatment with the reverse order (95.63%).

Finally, our analysis has thus far focused on treatment differences since these comparative statics can differentiate between dynamic models of risky choice. However, because participants received the endowment at the beginning of the experiment, the investment rate *levels* (as opposed to differences) may partly be driven by a house money effect that raised overall risk taking (Thaler and Johnson, 1990). We conducted additional studies to examine the possibility that the treatment differences may be affected by risk-taking levels. In these studies, which are described more fully in Appendix C, the investment decision removed the endowment frame by following the format used in Kahneman and Tversky (1979). Specifically, rather than being given an endowment that they could either invest or keep, participants faced a choice between two options: one option offered the equivalent of the one-round investment amount with certainty, while the other offered a 50/50 chance of either doubling that amount or losing it. In the case

²⁴ We elicited perceived complexity within the scope of the Traditional Stakes experiment using the four-item score of Maynard and Hakel (1997).

of multiple rounds, a choice of the safe option resulted in the participant keeping that amount for the remaining rounds. Note that the structure of the decision problem is equivalent to our main experimental specification. Nonetheless, changing the endowment frame decreased entry rates across the board; e.g., the One-Shot entry rate fell to 24%. Importantly, however, the differences between treatments remained significant and of similar magnitudes to the main experiment reported above.

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 and then compare them to ex-post behavior. The specific pattern of deviations, e.g., whether they are outcome-dependent or not, allows us to distinguish between different theoretical explanations.

B.1. Ex-Ante Strategies

Figure 7 illustrates the cumulative distribution of gain and loss limits in our sample. We find striking differences in the loss and gain limits: the distribution of gain limits first-order stochastically dominates the distribution of loss limits. Table III further shows that the average participant who accepts risk sets a gain limit that is more than 3 times higher than her loss limit. This corresponds to an average difference that is around 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 (84.0%) begin taking risk as part of a "loss-exit" strategy, whereas only 7.4% do so as part of a "gain-exit" strategy.

[INSERT TABLE III 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 they differ after gains versus losses), and whether this pattern is consistent with their ex-ante "loss-exit" plans.

23



Figure 7. Ex-ante strategy. This figure illustrates participants' strategies in the Hard Plan treatment. It shows the cumulative distribution of loss and gain limits. A gain limit is defined as the minimum gain at which the subject would prefer to stop taking risk rather than continue. A loss limit denotes the maximum loss at which she would prefer to stop taking rather than continue.

The cleanest test for dynamic inconsistency comes from comparing ex-ante strategies to behavior after the feedback in the first round—before endogenous exit decisions in the subsequent rounds had 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. Any differences in subsequent behavior in response to first-round gains and losses can thus be causally attributed to a differential reaction to losses versus gains. In contrast, because participants can choose to exit in any round, behavior in response to cumulative gains and losses may increasingly be a function of unobservable differences in exit preferences in previous rounds as opposed to a causal effect of cumulative gains and losses. To solve this identification challenge in later rounds, the firstround gains and losses serve as a proxy for the cumulative gains and losses, and the exogenous variation allows for causal inference.

Figure 8 presents the probability differences of continuing to invest after a first-round gain versus a first-round loss. A positive (negative) value corresponds to a greater (lower) propensity



Figure 8. Ex-ante plan versus ex-post choice after experiencing first-round gains and losses. This figure illustrates the differences in the probability to continue investing in the lottery after experiencing first-round gains versus first-round losses for participants in the Hard Plan and the Sequential treatment. The bars represent the average marginal effects of Logit regressions as outlined in Table IV and averaged across five-rounds intervals.

to take on risk after a gain than a loss.²⁵ Table IV reports the corresponding regression results. We find a striking difference between ex-ante plans and ex-post behavior. Participants' planned behavior in the Hard Plan treatment results in a significantly higher chance of taking on risk after a gain than a loss. This is the result of following "loss-exit" strategies—a plan with a more binding loss limit than a gain limit. As shown in Figure 8, the greater propensity to take on risk after a gain persists across the rounds. Participants' ex-post behavior in the Sequential treatment follows the opposite pattern: they are more than 20 percentage points more likely to take on risk after a first-round loss than a first-round gain. This negative difference in continuation probabilities persists across rounds.

Table IV presents the difference-in-difference in exit rates between planned (Hard Plan) and actual (Sequential) behavior. The coefficients are positive and sizable in the majority of

²⁵ These probability differences correspond to the marginal effects of Logit regressions on the subsample of participants who had entered the lottery.

rounds.²⁶ Together, the data on exit rates paint a stark contrast between ex-ante strategies and ex-post choices—a hallmark of a dynamic inconsistency in sequential risk-taking.

[INSERT TABLE IV ABOUT HERE]

Another way of testing for a dynamic inconsistency is comparing the cumulative outcome distributions generated by ex-ante plans versus ex-post behavior. The single-round lottery is symmetric and any outcome-independent strategy would also result in a symmetric outcome distribution. In contrast, a "loss-exit" strategy would result in a positively-skewed outcome distribution with higher probabilities of small losses and lower probabilities of large gains; "gain-exit" behavior would result in a negatively-skewed distribution. Figure 9 compares the ex-ante and ex-post outcome distributions of the Hard Plan and Sequential treatments, respectively. The outcome distribution in the Hard Plan treatment is based on 1000 simulated independent outcome sequences for the strategy of each participant who has entered the lottery. The outcome distribution in the Sequential treatment is based participants' actual behavior in the study.

The figure shows that the expected outcome distribution from ex-ante planned behavior has a positive skew consistent with a "loss-exit" strategy. In contrast, the ex-post sample outcome distribution has a negative skew, which is consistent with "gain-exit" behavior.²⁷ In particular, the ex-ante outcome distribution has a higher probability of loss (48.6%) than gain (39.5%). The reverse is true in the ex-post outcome distribution, where the probability of a loss (35.8%) is smaller than the probability of a gain (48.1%). A Kolmogorov-Smirnov (KS) test reveals that the difference between the ex-ante expected outcome distribution and the ex-post sample outcome distribution is statistically significant (p = 0.018). An alternative way of comparing the skewness of planned versus actual behavior is to calculate the probability that the strategies reported in the Hard Plan treatment would generate the negative skewness observed in the Sequential treatment. To do this, we run 1000 simulations, each time varying the outcome sequences of the participants in the Hard Plan treatment and calculating the resulting sample skewness of the outcome distribution. These simulations reveal that the outcome skewness from actual behavior corresponds to the 0.2nd percentile of the outcome skewness distribution

²⁶ In the later rounds, the outcome dependence in the Sequential treatment and the difference-in-difference between the treatment both decrease. This would be expected given that the choices in the later rounds are further removed from experiencing a gain or loss in first round.

²⁷ The extent of negative skewness is marginally significant (p = 0.066, D'Agostino-Pearson test).

in planned behavior. This suggests that the observed difference in skewness is highly unlikely without a systematic shift from "loss-exit" to "gain-exit" behavior.



Figure 9. Cumulative outcome distribution ex-ante plans versus ex-post behavior. This figure illustrates the differences in the cumulative outcome distribution in the Hard Plan treatment (cdf ex-ante) and the Sequential treatment (cdf ex-post), and the corresponding skewness estimates, for participants who accepted initial risk. The expected cumulative outcome distribution in the Hard Plan treatment is calculated based on 1000 independent simulations for each participant. The outcome distribution in the Sequential treatment is based on actual behavior.

Having identified a dynamic inconsistency in sequential risk-taking, we now consider whether people are sophisticated about it. To test for such awareness, we return to the initial decision to start taking on risk (see Figure 6). The higher entry rate in the Hard Plan treatment compared to the Sequential treatment suggests that a statistically significant portion of participants are sophisticated about their dynamic (Mann-Whitney test, p < 0.01). Namely, they are more willing to take on risk if they can do so while committing to follow through on their ex-ante strategy. We examine the extent of this sophistication in Section VI.

One potential concern for our analysis of dynamic inconsistency is that differences in entry decisions between treatments may bias any inference about ex-post choices. We address this issue by estimating 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 Hard Plan treatment assuming the same type of conservative assumption on selection, and (iii) replicate the comparisons between the Sequential and the Hard Plan treatment using these estimates. In particular, we assume that those who did not enter in the Sequential treatment would have followed a "loss-exit" strategy in the Hard Plan treatment. We thus adjust our estimates of the outcome distribution in the Hard Plan treatment to randomly omit a subset of people with "loss-exit" strategies until the number of investors is the same as in the Sequential treatment. Our results hold under this more conservative test. The skewness of the expected outcome distribution in the Hard Plan treatment decreases only slightly to 0.233. The Hard Plan outcome distribution remains significantly different from the Sequential treatment (KS test, p = 0.020). Moreover, the skewness of the Sequential distribution is in the 0.5th percentile of the Hard Plan treatment, providing further evidence that the observed differences are due to a dynamic inconsistency in strategies. Finally, Section VI and Appendix C.CII reports results of an additional experiment that compares participants' ex-ante plans and ex-post behavior within-subject, where selection issues are not present.

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 evidence that some people are sophisticated about this dynamic inconsistency, exhibiting a greater initial willingness to accept risk when offered an opportunity to commit to their strategy. Finally, we document a greater willingness to accept risk when it is part of a sequence of choices than when the same one-round gamble is offered in isolation.

Notably, 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).

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 for our experimental setting. The combination of findings outlined above is most consistent with CPT.²⁸ 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 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, once the agent starts taking on risk and experiences gains and losses, the probabilities of obtaining particular outcomes become less extreme and converge to the likelihoods of the single gamble. Since probability weighting has less bite, the prospect of taking on risk becomes less attractive and the agent deviates from her initial strategy.

While probability weighting generates the dynamic inconsistency, the shape of the value function contributes to its specific form. Diminishing sensitivity predicts risk-seeking behavior in the loss domain and risk-aversion in the gain domain. This leads the agent to display "gain-exit" behavior: compared to her initial strategy, she stops taking on risk earlier after gains and continues for longer after losses.²⁹ 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. In Appendix A we also discuss the importance of bracketing for the dynamic predictions of both CPT and related models. We follow prior work in considering the case where prior outcomes and prospective choices are bracketed together until either the agent stops taking on risk or the finite sequence ends (see, e.g., Barberis (2012) and Weber and Camerer (1998)).³⁰

To illustrate the intuition for the predictions of CPT, consider an agent facing a sequence of two 50:50 gambles, each with an upside of G and downside of -G. Let the agent be indifferent

²⁸ Despite notable differences between the settings, the predictions and experimental results are consistent with those outlined in Barberis (2012), who was the first to derive the dynamic behavioral predictions of CPT.

²⁹ As noted in the Appendix A, 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).

³⁰ Notably, the predictions of the case where outcomes and prospects are bracketed in isolation do not match the empirical patterns documented in this paper.

between accepting or rejecting the gamble in isolation. A "loss-exit" strategy generates a lottery of (2G, 1/4; 0, 1/4; -G, 1/2) over final outcomes, which has substantially more positive skew than the single gamble or an outcome-independent strategy, (2G, 1/4; 0, 1/2; -2G, 1/4). Probability weighting prompts the agent to overweight the small probability of the larger gain, leading her to accept the first gamble. After a gain, the agent faces a prospect of (2G, 1/2; 0, 1/2); after a loss, the prospect is (0, 1/2; -2G, 1/2). Given the even odds, the probability-weighting motive to take on risk is muted. Instead, risk-aversion (risk-seeking) in the gain (loss) domain leads the agent to reject the gamble after a win and accept it after a loss over a wide range of parameters. The same agent who accepts the first bet as part of a "loss-exit" strategy systematically deviates from it by stopping earlier after winning and continuing later after losing.

Importantly, the model implies that the discrepancy in risk-taking is not due to the dynamic environment per se; rather, people are more likely to accept a gamble as part of a sequence because they can generate a level of positive skew over the outcome distribution that is unavailable when the gamble is offered in isolation. To test this proposition, we ran two alternative One-Shot experiments that used gambles featuring return distributions matching the estimated skew generated by participants' initial strategies in the Hard Plan treatment (PS-lottery) and their actual decisions in the Sequential treatment (NS-lottery) of the main study (see Appendix C for details). In the first study, a separate group of participants (N = 148) were presented with a choice between the PS-lottery, the NS-lottery, or keeping an endowment with certainty. This study conceptually replicates our main experimental setting, where the choice to take on risk over the endowment represents the entry decision, while the choice of the PS-lottery over the NS-lottery represents a preference for the distribution generated by the ex-ante strategy over ex-post behavior. We find that 76.4% of participants took on risk, which was close to the entry rates in the Hard Plan treatment. Of those who took on risk, 77.0% preferred the PS-lottery to the NS-lottery. These results suggest that the prospect to construct more skewed risk—rather than the dynamic nature of the environment per se—drives initial choices in our study.³¹

³¹ We also ran another study (N = 49) featuring a direct choice between the PS- and NS-lottery without the option to keep the endowment. A similar proportion of participants chose the PS-lottery over the NS-lottery (80%).

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 and exhibit "gain-exit" behavior. While peoples' strategies generate a positively-skewed distribution over final earnings, their actual choices generate a negativelyskewed 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 key step for behavioral welfare analysis is identifying which choices merit deference, i.e., should be included in the welfare-relevant domain. To do so in our context, we ran a separate study (N = 160) 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.³² It begins in the same way as the Sequential treatment. However, after learning whether she lost or won the lottery in the first round, each participant was asked 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 choice set 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

³² The study was pre-registered here: https://aspredicted.org/blind.php?x=vw3uu8.

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 focus participants on the final outcomes of their choices and evaluate whether their prior gains and losses have reached the maximum amount they are willing to gain or lose.³³ 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 an ex-post "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 < 0.01). Importantly, the outcome frame appeared to successfully focus participants on evaluating their prior gains and losses with respect to the most that they are willing to gain or lose. There were no 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 = 0.60). Furthermore, participants' actual choices in the outcome frame generate a positively-skewed distribution over final outcomes (skew = 0.61, p < 0.01). This is consistent with the positively-skewed distribution in the Hard Plan treatment and opposite to the negatively-skewed distribution in the Sequential treatment (see Figure 9).

Based on this evidence, we can assess the welfare costs of dynamic inconsistency using a best-fit model. We use the CPT framework for this exercise because it is most consistent with our empirical finding, though one could of course use an alternative model of risky choice. As outlined in Barberis (2012), being dynamically inconsistent leads to two types of potential welfare losses, 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

³³ This design was motivated by the work in neuroscience and psychology which argues that negative urgency and compulsivity direct excess attention to prior outcomes in sequential risk-taking (for review, see Zhang and Clark, 2020).

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 naïveté*. 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*.³⁴

Appendix B reports results from simulations that assess the welfare costs of dynamic inconsistency in our experimental setting. Several takeaways are obtained. First, sophisticated agents would be willing to pay substantial amounts for commitment to execute their ex-ante strategy—up to 166% of the one-round endowment; the value of commitment is higher for those with greater skewness preferences. Second, both naïve and sophisticated agents incur a welfare loss in the absence of a commitment device. Naïve agents incur an additional loss because they are unaware of their inability to implement their ex-ante strategy and accept the initial gamble rather than rejecting it; the cost of naïveté leads to total welfare losses up to 200% of the one-round endowment. Third, 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. We should stress that this simulation exercise is meant to be illustrative, using one specific model of risky choice. Future work should extend this approach to more formal structural analysis and broaden its scope to other domains with normative ambiguity.

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

³⁴ As we show in Appendix A, with Expected Utility preferences, even the optimal and dynamically consistent strategy generates negative utility for the agent. In turn, the agent would not experience a welfare loss with or without commitment—she would not enter the lottery in either case. The agent would experience a welfare loss if she were coerced into entering the lottery, even if she followed the type of "loss-exit" strategy that was optimal for a sophisticated CPT agent.

behavior exhibits the opposite pattern—people cut gains early and chase their losses. 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 provide suggestive evidence that naïveté about deviations from these plans imply significant welfare costs.

Although our findings suggest the prevalence of naïveté, we also find evidence for sophistication about this dynamic inconsistency. Specifically, people are more likely to begin taking on risk when they are provided with a commitment opportunity. However, there appear to be limits to sophistication about the efficacy of non-binding commitment in disciplining behavior. Using a separate study, we examine behavior in the presence of 'soft' commitment devices where the gain and loss limits can be overwritten (see Appendix C.CII for experimental details). These non-binding commitment devices are widespread in the real world (see, e.g., our field setting). In a design similar to our Hard Plan treatment, participants (N = 149) were incentivized to report their gain and loss limits before they found out whether they were binding or not. Each was then randomized to either a setting similar to the Hard Plan treatment or one where they could *deviate* from their strategies if the limits were hit. Participants made their entry decision to take on risk after finding out whether the limits were binding or not. Similar to the field setting, we find a substantial dynamic inconsistency in this 'Soft Plan' treatment: of the participants whose limits are triggered, 80.7% deviate to take on more risk. At the same time, they display a similar demand for 'soft' commitment as hard commitment, despite the former being largely ineffective. These results point to an "illusion of commitment," suggesting that people may overestimate the efficacy of non-binding limits more generally.

The evidence on dynamic inconsistency and the "illusion of commitment" is important in light of the discretion that firms (e.g., financial brokers, casinos, etc.) have in designing and softening the commitment opportunities available to their clients. It also highlights the need for regulation to be evidence-based. For example, 'soft' commitment opportunities may have unintended adverse effects when bundled with regulation of consumer products, as in the case of the recently introduced "depreciation reporting rule" that is part of the revised European market in financial instruments regulation (MiFID II). The rule requires all European wealth managers, brokers, and financial advisers to immediately notify their clients when their portfolio looses at least 10% of its value relative to the beginning of the quarter.³⁵ This corresponds to an nonbinding loss limit at the level of 10%. Although the rule was intended to protect retail investors, our results suggest that the majority of investors will likely ignore the notification and continue with their current positions—or even double down. More importantly, however, the rule might change investors' ex-ante choices regarding the type and amount of risk to take. Specifically, the "illusion of commitment" may lead them to seek out types of risk that they would otherwise avoid—such as volatile assets with a zero or negative risk premium. As a result, 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.

On the theoretical side, our results suggest that although dynamic inconsistency in the domain of risk may not be normatively appealing (Machina, 1989), it is an important feature of the data. While 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; McClennen, 1988), our paper offers a starting point for the further development of models that capture dynamic inconsistency as an important feature of behavior.

Future research should examine individual-level heterogeneity in dynamic inconsistency. The goal of the current paper is to identify dynamic inconsistency and collect rich enough data to differentiate between different models of dynamic choice under uncertainty. To do so, we faced an inherent tension between clean identification of dynamic inconsistency, which, as noted in Cubitt, Starmer, and Sugden (1998), requires a between-subject design, and measuring individual-level parameters, which requires within-subject data and repeated measures. Such individual-level data would allow for the structural estimation of parameters and, thus, for making more refined predictions and generating welfare counterfactuals. Follow-up work should expand on our empirical design to further explore individual-level determinants of dynamic inconsistency.

Finally, our results relate to the work on self-control, impulsivity, and financial decisionmaking. 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 Previtero, 2016) to increased trade frequency and inferior financial performance. The form of

³⁵ 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/.
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.

Panel A: Trader characteristics								
	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				
Panel B: Position-level statistics								
	N_{total}	Mean	Median	Std dev.				

	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 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}) and Hard Plan treatment $(D^{hardplan})$. Column (1) shows results including the One-Shot treatment as a reference group. Column (2) displays results excluding the One-Shot treatment, hence the reference group is the Sequential treatment. zstatistics are in parentheses. *, ** and *** indicate statistically significant at the 10%, 5%, and 1% level, respectively.

	(1) mfx	(2) mfx
D^{seq}	0.154***	
$D^{hardplan}$	$(3.417) \\ 0.360^{***} \\ (5.425)$	$\begin{array}{c} 0.141^{***} \\ (2.642) \end{array}$
Pseudo R^2 N	$.124 \\ 295$	$.062 \\ 198$

Table III Ex-Ante Strategies

This table illustrates the ex-ante strategies in the Hard plan treatment. Panel A reports the share of participants who have a loss-exit, gain-exit, or equidistant neutral strategies. A loss-exit (gain-exit) strategy is defined as lower (greater) loss limit than gain limit. Column (1) reports the results for all participants. Column (2) reports the results only for those who initially choose to take on risk. Panel B reports the average gain and loss limits in USD. Panel C 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 ($\frac{qain}{loss}$) and "Difference" refers to their difference in USD (gain - loss). t-statistics of Wald tests for H0: Ratio = 1 and H0: Diff = 0, respectively, are in parentheses. Panel D reports the expected skewness of the aggregate sample outcome distribution that results from the gain and loss limits.

	All Subjects	Entered Lottery
	(1)	(2)
N	99	94
Panel A. Strategy Cate	gorization	
Loss-exit	84.8%	84.0%
Equidistant	8.1%	8.5%
Gain-exit	7.1%	7.4%
Panel B. Average I	Limits	
Gain Limit (\$)	7.879	7.989
Loss Limit (\$)	4.061	4.202
Panel C. Aggregate S	statistics	
Ratio	3.386***	3.187***
	(5.354)	(5.295)
Difference (\$)	3.818^{***}	3.787***
	(9.566)	(9.229)
Panel D. Expected Outcom	e Distributions	
Expected Skewness	0.269	0.254

Table IV

Ex-Ante Strategies vs. Ex-Post Behavior After an Immediate Gain vs. Loss

This table reports the marginal effects (mfx) of logit regressions of the probability of continuing to invest beyond round t. The independent variable in each column, which is denoted by a round number (t), equal 1 if a subject, who has entered the lottery at the beginning, continues to invest beyond round t. Panel A reports the results for the early rounds 2 to 10, Panel B reports the results for rounds 11-12, and Panel C reports the results for rounds 20 to 26. Row (I) reports results of logit regressions for the Hard Plan Treatment (ex-ante choice). The main independent variable is a dummy variable for immediate gain (i.e. gain in round 1). A positive coefficient indicates higher probability for continuing beyond round t after immediate gains than after immediate losses. Since subjects in the Hard Plan Treatment provide a fully contingent plan, we have information on their ex-ante planned choices both after immediate gains and immediate losses (strategy method). We simulate the case of immediate gains and immediate losses for each subject (N=94), resulting in 188 observations. We cluster the standard errors at the subject level. Row (II) reports marginal effects of analogous logit regressions (N=81) for the Sequential Treatment (ex-post choice). The main independent variable is a dummy variable for an immediate gain in round 1. Row (III) reports the interaction effects of a logit model including a dummy for immediate gain (vs. loss), a treatment dummy for Sequential treatment (vs. Hard Plan treatment) and their interaction. As the choices of each subject in the Hard Plan treatment are simulated for both cases of immediate gain and immediate loss, we cluster the standard errors at the subject level. z-statistics are in parentheses. *, ** and *** indicate statistically significant at the 10%, 5%, and 1% level, respectively.

Panel A. Rounds 2-10 (mfx)									
Round (t)	2	3	4	5	6	7	8	9	10
(I) G-L ex-ante	0.032	0.122**	0.109**	0.086^{*}	0.075	0.086	0.086	0.118**	0.118**
	(0.793)	(2.318)	(2.104)	(1.660)	(1.439)	(1.581)	(1.581)	(2.063)	(2.063)
(II) G-L ex-post	-0.189**	-0.269***	-0.273***	-0.205^{*}	-0.256**	-0.114	-0.117	-0.121	-0.097
	(-2.115)	(-2.746)	(-2.674)	(-1.899)	(-2.392)	(-1.030)	(-1.077)	(-1.137)	(-0.917)
(III) DID	0.221**	0.386***	0.379***	0.290**	0.331***	0.199^{*}	0.202*	0.238**	0.214*
× /	(2.356)	(3.552)	(3.400)	(2.491)	(2.849)	(1.661)	(1.715)	(2.065)	(1.867)
			Panel B. R	ounds 11	-19 (mfx)				
Round (t)	11	12	13	14	15	16	17	18	19
(I) G-L ex-ante	0.128**	0.138**	0.117**	0.106*	0.096	0.106*	0.085	0.085	0.074
< / <	(2.196)	(2.365)	(1.963)	(1.767)	(1.551)	(1.716)	(1.348)	(1.348)	(1.169)
(II) G-L ex-post	-0.097	-0.097	-0.026	-0.051	-0.027	-0.079	-0.081	-0.081	-0.081
	(-0.917)	(-0.917)	(-0.247)	(-0.501)	(-0.271)	(-0.802)	(-0.848)	(-0.848)	(-0.848)
(III) DID	0.225*	0.235**	0.143	0.158	0.123	0.185*	0.166	0.166	0.155
()	(1.955)	(2.060)	(1.282)	(1.439)	(1.140)	(1.760)	(1.637)	(1.637)	(1.539)
			Panel C. R	ounds 20	-26 (mfx)				
Round (t)	20	21	22	23	24	25	26		
(I) G-L ex-ante	0.074	0.085	0.085	0.116*	0.116*	0.116*	0.116*		
()	(1.169)	(1.331)	(1.331)	(1.816)	(1.816)	(1.816)	(1.816)		
(II) G-L ex-post	-0.033	-0.009	-0.011	-0.011	-0.011	-0.011	-0.011		
	(-0.358)	(-0.101)	(-0.127)	(-0.127)	(-0.127)	(-0.127)	(-0.127)		
(III) DID	0.107	0.094	0.096	0.128	0.128	0.128	0.128		
× /	(1.098)	(0.976)	(1.040)	(1.366)	(1.366)	(1.366)	(1.366)		

REFERENCES

- Allais, Maurice, 1953, Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école américaine, <u>Econometrica</u>: Journal of the Econometric Society 503–546.
- Allcott, Hunt, and Dmitry Taubinsky, 2015, Evaluating behaviorally motivated policy: Experimental evidence from the lightbulb market, American Economic Review 105, 2501–38.
- Andrade, Eduardo B., and Ganesh Iyer, 2009, Planned versus actual betting in sequential gambles, Journal of Marketing Research 46, 372–383.
- Arditti, Fred D., 1967, Risk and the required return on equity, <u>The Journal of Finance</u> 22, 19–36.
- Barber, Brad M., and Terrance Odean, 2000, Trading is hazardous to your wealth: The common stock investment performance of individual investors, <u>The Journal of Finance</u> 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, <u>Journal of Financial Economics</u> 104, 251–271.
- 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, <u>Review of Financial Studies</u> 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, 2009, Beyond revealed preference: choice-theoretic foundations for behavioral welfare economics, <u>The Quarterly Journal of Economics</u> 124, 51–104.
- Bernheim, B. Douglas, and Dmitry Taubinsky, 2018, Behavioral public economics, in <u>Handbook</u> of Behavioral Economics: Applications and Foundations 1, volume 1, 381–516 (Elsevier).
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2012, Salience theory of choice under risk, The Quarterly Journal of Economics 127, 1243–1285.
- Brown, Jeffrey R., and Alessandro Previtero, 2016, Saving for retirement, annuities and procrastination, working paper.
- Chen, Daniel, Martin Schonger, and Chris Wickens, 2016, oTree An open-source platform for laboratory, online, and field experiments, <u>Journal of Behavioral and Experimental Finance</u> 9, 88–97.
- Chetty, Raj, Adam Looney, and Kory Kroft, 2009, Salience and taxation: Theory and evidence, American Economic Review 99, 1145–77.
- Christine, Kaufmann, Martin Weber, and Emily Haisley, 2013, The role of experience sampling and graphical displays on one's investment risk appetite, Management Science 59, 323–340.
- Cubitt, Robin, Chris Starmer, and Robert Sugden, 1998, Dynamic choice and the common ratio effect: An experimental investigation, The Economic Journal 108, 1362–1380.
- Dai, Hengchen, Katherine L. Milkman, and Jason Riis, 2014, The fresh start effect: Temporal landmarks motivate aspirational behavior, Management Science 60, 2563–2582.
- DellaVigna, Stefano, and Ulrike Malmendier, 2006, Paying not to go to the gym, <u>American</u> Economic Review 96, 694–719.
- Dertwinkel-Kalt, Markus, Jonas Frey, and Mats Köster, 2020, Optimal stopping in a dynamic salience model, <u>mimeo</u>.
- Ebert, Sebastian, 2020, On taking a skewed risk more than once, Available at SSRN 3731565 .

- Ebert, Sebastian, and Paul Karehnke, 2021, Skewness preferences in choice under risk, <u>Available</u> at SSRN .
- 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 R.K., Alex Imas, and Christy Kang, 2022, On the role of similarity in mental accounting and hedonic editing., Psychological Review 129, 777.
- 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, Shane, George Loewenstein, and Ted 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, <u>The</u> Quarterly Journal of Economics 112, 631–645.
- Gonzalez, Richard, and George Wu, 1999, On the Shape of the Probability Weighting Function, Cognitive Psychology 38, 126–166.
- Gupta, Neeraja, Luca Rigott, and Alistair Wilson, 2021, The experimenters' dilemma: Inferential preferences over populations, working paper, Cornell University.
- Heimer, Rawley, and Alp Simsek, 2019, Should retail investors' leverage be limited?, <u>Journal of</u> <u>Financial Economics</u> 132, 1–21.
- Heimer, Rawley Z., 2016, Peer Pressure: Social Interaction and the Disposition Effect, <u>The</u> Review of Financial Studies 29, 3177–3209.
- Hertwig, Ralph, and Ido Erev, 2009, The description–experience gap in risky choice, <u>Trends in</u> cognitive sciences 13, 517–523.
- Hogarth, Robin, and Emre Soyer, 2015, Communicating forecasts: The simplicity of simulated experience, Journal of Business Research 68, 1800–1809.
- Horton, John J., David G. Rand, and Richard J. Zeckhauser, 2011, The online laboratory: Conducting experiments in a real labor market, Experimental Economics 14, 399–425.

- Imas, Alex, 2016, The realization effect: Risk-taking after realized versus paper losses, <u>American</u> Economic Review 106, 2086–2109.
- Kahneman, Daniel, and Amos Tversky, 1979, Prospect theory: An analysis of decision under risk, Econometrica 47, 263–291.
- Karni, Edi, and Zvi Safra, 1990, Behaviorally consistent optimal stopping rules, <u>Journal of</u> Economic Theory 51, 391–402.
- Karni, Edi, and David Schmeidler, 1991, Atemporal dynamic consistency and expected utility theory, Journal of Economic Theory 54, 401–408.
- Kaustia, Markku, 2010, Disposition effect, in <u>Behavioral Finance</u>, 169–189 (Wiley Online Library).
- 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, <u>Quarterly Journal of Economics</u> 112, 443–477.
- Langer, Thomas, and Martin Weber, 2008, Does commitment or feedback influence myopic loss aversion?: An experimental analysis, <u>Journal of Economic Behavior & Organization</u> 67, 810–819.
- Linnainmaa, Juhani T., 2010, Do limit orders alter inferences about investor performance and behavior?, The Journal of Finance 65, 1473–1506.
- Machina, Mark J., 1989, Dynamic consistency and non-expected utility models of choice under uncertainty, Journal of Economic Literature 27, 1622–1668.
- Maynard, Douglas C., and Milton D. Hakel, 1997, Effects of objective and subjective task complexity on performance, Human Performance 10, 303–330.

- McClennen, Edward F, 1988, Constrained maximization and resolute choice, <u>Social Philosophy</u> and Policy 5, 95–118.
- 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, <u>Journal of</u> Economic Behavior & Organization 161, 68–78.
- Odean, Terrance, 1998, Are investors reluctant to realize their losses?, <u>The Journal of finance</u> 53, 1775–1798.
- O'Donoghue, Ted, and Matthew Rabin, 1999, Doing it now or later, <u>American Economic Review</u> 89, 103–124.
- Peköz, Erol A., 2002, Samuelson's fallacy of large numbers and optional stopping, <u>Journal of</u> Risk and Insurance 69, 1–7.
- Ploner, Matteo, 2017, Hold on to it? an experimental analysis of the disposition effect., Judgment & Decision Making 12, 118–127.
- 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, <u>Journal of Economic Behavior &</u> Organization 3, 225–243.
- Samuelson, Paul A., 1963, Risk and uncertainty: A fallacy of large numbers, Scientia.
- Shefrin, Hersh, 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.
- Shiv, Baba, George Loewenstein, Antoine Bechara, Hanna Damasio, and Antonio R Damasio, 2005, Investment Behavior and the Negative Side of Emotion., <u>Psychological Science</u> 16, 435–9.
- Snowberg, Erik, and Leeat Yariv, 2021, Testing the waters: Behavior across participant pools, American Economic Review 111, 687–719.

- 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, working paper.
- Volij, Oscar, 1994, Dynamic consistency, consequentialism and reduction of compound lotteries, Economics Letters 46, 121–129.
- Weber, Martin, and Colin Camerer, 1998, The disposition effect in securities trading: an experimental analysis, Journal of Economic Behavior & Organization 32, 167–184.
- 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.³⁶

Appendix AI. 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.³⁷ A decision-maker considers the gamble $L = (x_{-m}, p_{-m}; ...; x_{-1}, p_{-1}; x_0, p_0; x_1, p_1; ...x_n, p_n)$, where p_i corresponds to the likelihood of attaining outcome x_i . Outcomes are ordered such that -m...-1 correspond to those below the reference point x_0 , here assumed to be the status quo, and 1...n correspond to those above the reference point. We follow Tversky and Kahneman (1992) in assuming that utility derived from this gamble can be represented by:

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

where

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

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

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

and value function

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

³⁶ 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.

³⁷ 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.

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

In our setting, let L = (-50, 1/2; 50, 1/2). 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 plans available after the decision-maker accepted the first gamble and lost. Each plan $s \in S_{t,j}$ is a mapping from every potential outcome of the sequence of gambles from round t onward to actions $a \in \{continue, exit\}$.

Each s generates a random variable \tilde{G}_s , which corresponds to the accumulated gains or losses conditional on s being carried out. For example, take $s \in S_{0,1}$ where the decision maker plans to accept the first gamble, continue if she wins in t = 1 and then exits regardless of the next outcome in t = 2, and exiting in t = 1 if she loses the first gamble. This plan corresponds to $\tilde{G}_s \sim (-50, 1/2; 0, 1/4; 100, 1/4)$. 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, as estimated by Tversky and Kahneman (1992) (i.e. $\lambda = 2.25$, $\alpha = 0.88$, $\delta = 0.65$), would require in order to accept the first gamble in our empirical setting. We find that this number is 26 rounds.³⁸ Thus

³⁸ 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.

we examine the behavioral predictions given a dynamic sequence of 26 potential gambles, such that $t \in \{0, ..., 26\}$.

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 \$13, hence the agent does not update the reference point until the final period when the outcome is paid out.³⁹ This assumption is consistent with prior experimental evidence (see Imas, 2016).

To examine the sensitivity of the CPT predictions to the exact shape of the probability weighting function, we further consider a two-parameter weighting function proposed by Gonzalez and Wu (1999):

$$\tilde{w}(p) = \frac{\gamma p^{\beta}}{\gamma p^{\beta} + (1-p)^{\beta}}$$
(A5)

The parameters γ and β control the elevation and curvature, respectively. We run simulations over a wide range of parameter values to capture a large variety of possible shapes of the weighting function.

³⁹ We discuss an alternative assumption about the reference point below.

Results (CPT)

In the following we present the simulation results of CPT with the probability weighting function originally proposed by 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.

Figure A1, Panels A and C present the skewness and the standard deviation of the ex-ante expected outcome distribution for each combination of the CPT parameters α and δ . The level of loss aversion is fixed at $\lambda = 2.25$, which corresponds to the parameter estimate by Tversky and Kahneman (1992). The skewness of the outcome distribution is an indicator of the type of outcome dependence of the agent's optimal plan s^* . As a benchmark, an equidistant neutral plan would result in zero skewness. Positive skewness indicates a "loss-exit" plan, which is characterized by a lower loss limit than the gain limit. Negative skewness indicates a "gain-exit" plan, which is characterized by a lower gain limit than the loss limit. The standard deviation of the expected outcome distribution is indicative of the duration of investment. The unit of investment used in the simulations is normalized to 1, hence the standard deviation of the single-round lottery equals 1. For comparison, the maximum standard deviation of investing in the lottery unconditionally for the maximum duration of 26 rounds equals to 5.1 (i.e., $\sqrt{26}$).

Several findings are obtained. First, a broad range of CPT agents would accept the first gamble in a sequence with endogenous exit even though they 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" for agents with pronounced non-linear probability weighting and moderate diminishing sensitivity. Interestingly, the optimal plan s^* of all CPT agents, who start investing with a "loss-exit" plan, is homogeneous and consists of the smallest possible loss limit and a

non-binding gain limit. 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.⁴⁰

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. For lower levels of loss aversion (i.e. $\lambda < 2.25$), the area of parameter combinations of agents who would not accept the sequential gamble shrinks.

Figure A1, Panels B and D 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. 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 homogeneous in the "loss-exit" parameter region, the actual behavior is heterogeneous and jointly driven by probability weighting and diminishing sensitivity, i.e., the extent to which agents are risk-averse in the gain domain and risk-seeking in the loss domain. Importantly, probability weighting alone predicts a reversal from a "loss-exit" ex-ante plan, characterized by a positive skewness of the outcome distribution, to a "gain-exit" ex-post behavior, which is characterized by a negative skewness of the outcome distribution. The level of diminishing sensitivity determines the extent of "gain-exit" behavior. Compared to agents with $\alpha = 1$, agents with smaller α exit sooner in the gain domain, thus realizing a more negative skew and a lower standard deviation of the ex-post outcome distribution. Notably, our experimental results outlined in Section IV suggest early exit in the gain domain, which is characteristic for CPT agents with both pronounced probability weighting and pronounced diminishing sensitivity. Critically, all agents who accept risk as part of a "loss-exit" strategy end up with an outcome distribution that has a *lower* utility value than rejecting the first gamble in t = 0. This represents substantial dynamic inconsistency between planned and actual behavior.⁴¹

⁴⁰ 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.

⁴¹ They also exit after breaking even if the remaining number of rounds is below an agent-specific minimum 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.

The 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 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.

It is important to highlight that our assumption that agents do not update their reference point until the final round is critical for predictions on ex-post behavior. This amounts to assuming that agents "bracket" their choices together until earnings are realized, which takes place either 1) when the agent stops taking on risk or 2) the final found. This assumption is consistent with the findings of Imas (2016), who builds on the realization utility framework of Barberis and Xiong (2012). There, he develops a theoretical model to show that the inconsistency in how people respond to prior losses can be explained by whether those losses are realized or not. Realized losses reset the reference point and are not bracketed with prospective choices; paper losses do not reset the reference point and are evaluated in the same bracket as prospective choices. A series of experimental studies provide direct support for this framework. In our framework, all outcomes before the last round of risk-taking are treated as paper gains and losses; the reference point does not update and the agent integrates them with prospective choices both when forming her strategy and in her ex-post behavior.⁴² Our experimental procedures draw on Imas (2016) to ensure that outcomes within a dynamic sequence were not realized before risk-taking was terminated (either the participant or in the final round), and that this was clearly communicated to participants beforehand.

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 value of a "loss-exit" strategy is

⁴² Note that this assumption is implicit in essentially all models of dynamic risky choice. For example, Weber and Camerer (1998) assume that prior gains and losses are bracketed together with the prospect when deriving predictions for the disposition effect.

lower than it was in the beginning of the sequence. Once the number of rounds falls below the agent-specific threshold, she exits. This leads to the prediction that the agent is just as likely to exit after a loss as after a gain.⁴³ In general, the closer an agent is to the white-colored area in Figure A1, Panel A, the higher is her agent-specific threshold 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).⁴⁴ 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—including 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, Imas, and Kang, 2022)—suggests that a finite time horizon may be appropriate in many real world settings.

⁴³ 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. 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.

⁴⁴ 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.

Next, we replace the probability weighting function originally proposed by Tversky and Kahneman (1992) with the two-parameter weighting function proposed by Gonzalez and Wu (1999). Figure A2 reports the simulation results for the ex-ante plan (Panels A and C) and the actual behavior (Panels B and D), respectively. We report the skewness and the standard deviation of the expected ex-ante outcome distribution and the ex-post outcome distribution as indicators of the type of outcome dependence (i.e., "loss-exit" versus "gain-exit") and the duration of investment in the lottery. The simulations are conducted analogously for a broad range of combinations of *elevation* (γ) and *curvature* (β) of the weighting function. The CPT parameters are fixed at $\lambda = 2.25$ and $\alpha = 0.88$ consistent with the estimates by Tversky and Kahneman (1992).

The simulation results suggest that the two-parameter weighting function predicts similar ex-ante plans but more types of ex-post deviations. First, in the parameter region of $\gamma < 1 \cap \beta < 1$ the predictions are largely consistent with the model considered above. In particular, most parameter combinations are linked to a "loss-exit" ex-ante plan that is largely homogeneous across agents and "gain-exit" actual behavior that varies in scope depending on the exact parameter combination. In contrast, the opposite is predicted for agents in the parameter region of $\gamma > 1 \cap \beta > 1$. If agents in this region start investing, they begin with a "gain-exit" plan and deviate to "loss-exit" behavior ex-post. In the remaining parameter regions, agents also mostly deviate from their ex-ante plan but their ex-post behavior is in the same category as the ex-ante plan. Notably, the elevation parameter, γ , determines the form of the ex-post deviation. This is a generalization of the one-parameter probability weighting function by Tversky and Kahneman (1992), which only allows for "gain-exit" ex-post behavior.

Our experimental results outlined in Section IV show that this is enough to capture the behavior of the average participant in our experiment. However, a two-parameter weighting function may better capture the heterogeneity among participants as it allows for a broader range of possible ex-post deviations. Importantly, the predicted ex-post behavior of all agents based on both the one-parameter and the two-parameter probability weighting functions displays some form of outcome-dependence. This is a distinctive feature of CPT compared to RDU, as we show in the next section.

Rank-Dependent Utility

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

$$\sum_{-m}^{n} \pi_i^{RDU} u(W + x_i), \tag{A6}$$

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

$$u(W+x) = \begin{cases} \frac{(W+x)^{1-\gamma}}{1-\gamma} & \text{for } \gamma \ge 0 \& \gamma \ne 1, \\ ln(W+x) & \text{for } \gamma = 1, \end{cases}$$
(A7)

and

$$\pi_i^{RDU} = w(p_{-m} + \dots + p_i) - w(p_{-m} + \dots + p_{i-1})$$
(A8)

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 equals 26 times the one-round investment amount.

Results (RDU)

Figure A3, Panels A and C present our findings on the skewness (as an indicator of strategy type) and the standard deviation (as an indicator of duration of investment) of the expected outcome distribution of 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. Only agents with low levels of risk aversion and strongly non-linear probability weighting would accept the gamble for a single round. Note that this result depends critically on the wealth assumption. Alternatively, assuming a wealth level of \$1,000 leads to the prediction that all agents 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 A3, Panels B and D present the skewness (as an indicator of strategy type) and the standard deviation (as an indicator of duration of investment) of the realized outcome distributions of RDU agents. The main result is that while RDU does predict deviations from the "loss-exit" strategy for some agents, the deviations are almost always outcome-independent in response to gains and losses. In particular, the most likely deviation is to keep investing unconditionally until the end. This behavior results in zero skewness and maximum standard deviation of the outcome distribution. This is in contrast to the results from CPT which predict outcome-dependent deviations.

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 A3. 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. This result is consistent with Ebert and Karehnke (2021) who formally show that skewness preferences cannot be first order in driving behavior in EUT and that a behavioral model is needed. It also follows trivially that EUT does not predict any dynamic inconsistency.⁴⁵

Quasi-Hyperbolic Discounting

We follow Laibson (1997) and assume a quasi-hyperbolic discount function. The expected utility of the agent is assumed to be:

$$U_t = \mathbb{E}_t \left[u(c_t) + \beta \sum_{\tau=1}^{T-t} \delta^{\tau} u(c_{t+\tau}) \right]$$
(A9)

where t denotes the round number and the utility function $u(c_t)$ is given by Equation A7. 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.⁴⁶

⁴⁵ 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.

⁴⁶ 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.

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.



Panel C. Standard deviation of ex-ante plan

Panel D. Standard deviation of ex-post behavior

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). Panels A and B report the skewness of the outcome distribution of the agent's ex-ante plan and her actual (ex-post) behavior. Positive skewness indicates a "loss-exit" strategy (i.e. loss limit lower than the gain limit), whereas negative skewness indicates a "gain-exit" strategy (i.e. loss limit greater than the gain limit). Equidistant strategies result in neutral skewness. Panels C and D report the corresponding standard deviation of the ex-ante and ex-post outcome distributions, respectively. The standard deviation indicates the duration of investment. The investment unit for a single round of investment is normalized to 1 for the purposes of the simulations. For comparison, a standard deviation of 5.1 (i.e., $\sqrt{26}$) indicates that the agents invests in the lottery for the maximum number of rounds unconditionally. The loss aversion is set to 2.25 in all simulations, which corresponds to the loss aversion of the representative agent in Tversky and Kahneman (1992).



Panel C. Standard deviation of ex-ante plan

Panel D. Standard deviation of ex-post behavior

Figure A2. Theoretical predictions of Cumulative Prospect Theory with two-parameter probability weighting function. 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 and a Gonzalez and Wu (1999) probability weighting function given different levels of curvature (β) and elevation (γ). Panels A and B report the skewness of the outcome distribution of the agent's ex-ante plan and her actual (ex-post) behavior. Positive skewness indicates a "loss-exit" strategy (i.e. loss limit lower than the gain limit), whereas negative skewness indicates a "gain-exit" strategy (i.e. loss limit greater than the gain limit). Equidistant strategies result in neutral skewness. Panels C and D report the standard deviation of the ex-ante and ex-post outcome distributions, respectively. The standard deviation indicates the duration of investment. The investment unit for a single round of investment is normalized to 1 for the purposes of the simulations. For comparison, a standard deviation of 1 indicates that the agent invests in the lottery for one round. Analogously, a standard deviation of 5.1 (i.e., $\sqrt{26}$) indicates that the agents invests in the lottery for the maximum number of rounds unconditionally. We set the value function parameters $\lambda = 2.25$ and $\alpha = 0.88$, corresponding to the representative agent in Tversky and Kahneman (1992).



Panel C. Standard deviation of ex-ante plan

Panel D. Standard deviation of ex-post behavior

Figure A3. 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). Panels A and B report the skewness of the outcome distribution of the agent's ex-ante plan and her ex-post behavior. Positive skewness indicates a "loss-exit" strategy (i.e. loss limit lower than the gain limit), whereas negative skewness indicates a "gain-exit" strategy (i.e. loss limit greater than the gain limit). Equidistant strategies result in neutral skewness. Panels C and D report the standard deviation of the ex-ante and ex-post outcome distribution, respectively. The standard deviation indicates the duration of investment. In the simulations the investment unit for a single round of investment is normalized to 1. For comparison, a standard deviation of 5.1 (i.e., $\sqrt{26}$) indicates the agents invests in all 26 rounds of the lottery unconditionally. We assume zero initial wealth, the agent receives an investment endowment which is just enough for all rounds.

Appendix B. Cost of Naïveté and Value of Commitment

In the following we describe 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 exante 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.

Figure B1 presents the certainty equivalents of the ex-ante strategy (i.e., the value of commitment) and the ex-post behavior (i.e., the 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. Agents would be willing to pay up to 166% of the one-round endowment depending on their parameter combination. The "value of commitment" is higher for agents with stronger skewness preferences (i.e., those with a lower δ and higher α).

In the absence of a commitment device, both naïve and sophisticated agents incur a welfare loss that corresponds to the value of commitment reported in Figure B1, Panel A. Naïve 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 B1, 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. For most agents in this area, the "costs of naïveté" range between 1 and 2. This means that these agents would prefer sacrificing between one and two times the one-round investment amount (i.e., between 50 cents and \$1 in our main experiment) to avoid the option of investing in the sequential lottery. 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.



Figure B1. 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.

Appendix C. Additional Experiments

In addition to our baseline experiment, we conducted a series of online experiment to test for potential alternative explanations. Below we describe the modifications and extensions of the baseline design, which we introduce in the additional experiments. An overview of all experiments is provided in Table CI. The main findings are displayed in Table CII.

Appendix CI. Traditional Stakes

We run two pre-registered experiments with a total of 940 subjects in February and June 2019 with traditional stakes. These studies largely replicate the design of the baseline experiment.⁴⁷ The latter experiment further extends the design by an additional treatment that is further described in Section C.CII. In the One-Shot treatment, participants receive an endowment of 10 cents, instead of 50 cents, and decide whether or not to invest in a single gamble. In the multi-round treatments, participant receives an initial endowment of \$2.60 at the beginning of the experiment, instead of \$13, and decides whether to invest 10 cents in each round or to keep it over a maximum of 26 rounds. The experimental stakes, albeit smaller than the stakes in our baseline experiment, were on the higher end of those typically offered on the platform and worked out to approximately \$12 an hour.

Table CII outlines the pooled results of both experiments. First, consistent with the results of our baseline experiment, we show that more participants enter the lottery if presented in a sequence as opposed to in isolation. Second, we find that the majority of participants start taking risk with a dynamic "loss-exit" strategy. Third, results of the cleanest test of ex-post deviation, which is a comparison of the continuation probability after an immediate gain or loss in the first round, reveals an opposite ex-post behavior – subjects are more likely to continue taking risk if they have endured a loss. This suggests a dynamic inconsistency similarly to the findings of our baseline experiment. Fourth, we find evidence of a demand for commitment consistent with sophistication, as significantly more participants choose to start taking risk in the Hard Plan treatment (with commitment device) than in the Sequential treatment (without commitment device).

⁴⁷ The pre-registration reports can be found at https://aspredicted.org/RAE_OWZ and https://aspredicted.org/KWF_XVL, respectively.

Appendix CII. Illusion of Commitment and Policy Implications

In our baseline experiment, we find a significant demand for commitment, which suggests that a significant part of the participants are sophisticated about the dynamic inconsistency. We explore sophistication about the efficacy of non-binding commitment in an additional, fourth, between-subject treatment with 149 subjects, which we added to the baseline design in an experiment conducted in June 2019. The *Soft Plan treatment* 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 examine whether setting a non-binding plan affects the decision whether or not to start taking on risk and the ex-post behavior.

The fact that the plans are non-binding constitutes a challenge for their incentivization. To ensure that the plans are fully incentivized, we elicited the plans before randomly distributing the participant in either the Hard Plan (binding plans) or the Soft Plan treatment (non-binding plans).⁴⁸ The instructions of the baseline Hard Plan treatment were modified by the following statement in both the Hard Plan and the Soft Plan treatments:

In case either your loss or gain limit is reached, you will either stop investing automatically or we will inform you immediately. On the next page we will randomly determine whether your limits are executed automatically or we inform you. The probability for each condition is 50%.

After submitting their limits, participants were informed whether or not they were assigned to the Soft Plan treatment. For participants in the Soft Plan treatment, the information was as follows⁴⁹:

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 X cents or your loss limit of Y cents is reached before the final round. You will then choose whether to continue or stop investing.

⁴⁸ 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.

⁴⁹ The instructions of the corresponding Hard Plan treatment in the experiment conducted in June 2019 were altered correspondingly. Analogously to the Soft Plan treatment, subjects, who were randomly assigned to the Hard Plan treatment, were informed that their limits will be automatically executed after the limits were elicited and before subjects chose whether or not to start 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.

The participants then proceed to the investment choice in the first round, which is analogous to the investment choice in the baseline Hard Plan treatment. Conditional on entering the lottery, participants are then notified in subsequent rounds in case either of the limits are triggered and are asked to decide whether or not they continue investing. Such notifications are issued every time a limit is triggered.

Table CII shows that a non-binding commitment device significantly increases the likelihood that participants enter and accept risk. Moreover, the increase is comparable to the increase due to the availability of a binding commitment device. The difference between the entry rates with a binding and a non-binding commitment device is not significant at the 10% level. Our results thus suggest that subjects presume that non-binding commitment devices are just as effective against dynamic inconsistency as binding commitment devices.

Examining the ex-post choices in the Soft Plan treatment, however, we find that a large majority of the subjects override their ex-ante loss limits when they are non-binding. 80.7% of the participants whose limits are triggered decide to continue investing and thus deviate from their ex-ante strategies. The majority (70.2%) of the triggered limits are loss limits, which should be expected given that most participants set loss limits closer than the gain limits. 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. Such all-the-way deviation is observed in 56% of the cases in which a loss limit is triggered. As a result of the specific pattern of deviation, the skew of the outcome distribution in the Soft Plan treatment drops from 0.471 ex-ante (p = 0.023) to 0.280 (p = 0.166) ex-post.⁵⁰

⁵⁰ A larger difference could be expected in a setting where (i) subjects are also allowed to stop investing manually before their limit is hit, as is the case in our field data, or (ii) the time lag between initial investment choice and deviation is longer.

Taken together, we find that a non-binding commitment device significantly increases the likelihood that participants accept initial risk. Yet, such devices are ineffective in mitigating dynamic inconsistency. Importantly, non-binding commitment devices are closer to commitment devices used in the real world. For instance, a gambler who has committed to limiting her losses by withdrawing a certain amount of cash may use credit or withdraw additional cash to deviate from his non-binding commitment. Traders who set stop-loss and take-profit limits can revise the orders to prevent them from being triggered, which is the case in our field data. In Europe, advisors and portfolio managers are required by the recently introduced MiFID II regulation to inform their clients immediately in case their portfolio depreciates by more than 10% from the beginning of the quarter (i.e., depreciation reporting). Such reporting rules, just like any reminders, can be viewed as non-binding commitment devices. They highlight a particular dynamic strategy, namely to take action in case of a loss, that the investor may ex-ante intend to follow but fail to execute ex-post. As we discuss in Section VI, such realistic commitment opportunities can have unintended adverse consequences for consumer welfare.

Appendix CIII. Robustness: House Money Effect

Our main analysis focuses on differences between the level of risk-taking in between-subject treatments. The exact levels of risk taking should be interpreted with caution, as they can be driven by a house money effect by design. We conduct two additional experiments to demonstrate the impact of a house money effect on the *level* of risk taking. First, focusing on the One-Shot treatment alone, we conduct an online experiment to eliminate the initial endowment (i.e., house money) in an *option-framed One-Shot treatment*, and to compare it in a counterbalanced within-subject design (N=59) to our baseline One-Shot treatment. Second, we eliminate the initial endowment in all treatments and replicate our baseline experiment with 314 subjects using only the option-framed versions of the treatments. Specifically, instead of assigning the participants an initial endowment at the beginning and asking them to choose whether or not to invest in a fair gamble, we ask them to choose between two options as in Kahneman and Tversky (1979):

You are given the following options. Which one would you choose?

Option A: 50% chance to win 20 cents, 50 chance to win nothing.

Option B: 10 cents for sure.

Participants in the multi-round treatment *option-framed Hard Plan* and *option-framed Sequential* were then instructed to think of investing in the lottery as choosing Option A and to think of keeping the money as choosing Option B. The rest of the instructions, importantly the elicitation of the gain and loss limits, followed the baseline design.

Table CII shows that (i) eliminating the initial endowment reduces the level of risk taking in the One-Shot treatment by 32 percentage points, and (ii) the treatment differences documented in the baseline experiment are robust to the house money effect. Importantly, the level of risk taking in the *option-framed One-Shot treatment* was 24% and thus comparable to the level documented in Kahneman and Tversky (1979) for a similar fair symmetric lottery (which was 16%).

Appendix CIV. Robustness: Expectations

In a further experiment, which was designed to test for the role of biased beliefs on the investment decision, we included a *Hard Plan treatment with beliefs induction* and tested it against the baseline Hard Plan treatment and the One-Shot treatment using a between-subject design. This treatment was designed analogously to the the baseline Hard Plan treatment and extended by a note on Doob's Optional Stopping Theorem. The instructions were as follows:

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.

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.

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

Table CII shows that the entry rate in the treatment with beliefs induction does not differ significantly from the entry rate without beliefs induction. This finding suggests that the high entry rate in the baseline Hard Plan treatment and the difference to the entry rate in the One-Shot treatment cannot be explained by overoptimism in the expected return of one's own dynamic strategy.

Appendix CV. Robustness: Wealth Effects

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 higher participation fee to equalize the wealth level of the participants in all treatments. Subjects in the *One-Shot treatment with high participation fee* received \$3 participation fee and 10 cents endowment to invest in the main task, which was a single fair symmetric lottery. The experiment was conducted with 122 participants and tested the new treatment against the baseline One-Shot treatment (50 cents participation fee and 10 cents endowment) and the baseline Hard Plan treatment (50 cents participation fee and \$2.60 endowment) using a between-subject design.

Table CII shows that the difference between the One-Shot and the multi-round treatment (Hard Plan) is robust to equalizing the participants' wealth level.

Appendix CVI. Outcome Frame

Finally, we conduct a study with 160, in which we provide subjects feedback of their outcome in the first round before eliciting their limits. This *Outcome Frame* allows us to test for a sunk cost fallacy as a potential explanation of the high entry rates in the baseline Hard Plan treatment. The participants were informed of the sequence of events before making an investment decision as follows:

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.

The rest of the instructions prior to the investment decision were analogous to the Hard Plan treatment. After the outcome of the first round has been revealed analogous to the Sequential Treatment, we elicited the limits as follows: 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.

The entry rate in the Outcome Frame is comparable to the entry rate in the baseline Hard Plan treatment. This indicates that a sunk cost fallacy is unlikely to explain the high entry rates in the Hard Plan treatment and, in particular, the significant demand for commitment devices.

Another purpose of the Outcome Frame study is to identify the welfare-relevant benchmark based on participants' limits after the outcome of the first round is revealed. The results of this analysis are discussed in detail in Section V.

Appendix CVII. Static Skewness Preferences

In Section IV.C we discuss the role of skewness preferences resulting from probability weighting as a driver of the documented dynamic inconsistency. We conduct an additional experiment with 149 participants to compare the pattern of risk taking in dynamic environments with skewness preferences in a static environment. In particular, we aim to understand whether, in a static context, subjects prefer the positively skewed lotteries they construct in the dynamic environment ex-ante, the negatively skewed lotteries that they end up with ex-post, or neither. We conduct an additional experiment to elicit participants' *static skewness preferences* by giving them a choice between three options—a fair positively skewed lottery, a fair negatively skewed lottery, and not investing. Both lotteries were selected to closely match the expected skewness of the average ex-ante plan and the skewness of the ex-post outcome distribution in the "Large Stakes" experiment. The two lottery options were as follows:

Lottery A: 40% chance to receive \$6, 60% chance to receive \$1.

Lottery B: 60% chance to receive \$5, 40% chance to receive \$0.

The lotteries have an identical mean and variance and opposite skew of 0.41 in Lottery A and -0.41 Lottery B. If the participant chooses not to invest in the lottery she would receive \$3 for sure, corresponding to the mean of the lotteries.

Our findings can be summarized as follows. First, an investment rate of 76.4% in either one of the lotteries indicates that participants prefer risk over certainty in a static environment, similar to entry rates in the dynamic environment. Second, a majority of 77.0% of the participants who choose to invest selected the positively skewed lottery. This is consistent with the preference for ex-ante "loss-exit" strategies that we document in dynamic environments. Taken together, these results suggest that participants begin to take on risk in dynamic environments because this allows them to construct positive skew over final outcomes that is not available when a gamble is offered in isolation; specifically, the difference in risk-taking between static versus dynamic environments is due to skewness preferences rather than temporal aspects of the context.

Appendix CVIII. Execution

In total, we ran experiments with 2621 participants. Thereof, 1085 subjects were recruited from the online platform Prolific and 1536 were recruited from Amazon Mechanical Turk via Cloud Research. In both cases, we used standard sample selection criteria. For experiments on Prolific we included subjects with an approval rating of at least 97% with location in the US, UK or Ireland. For experiments on MTurk via Cloud Research, we included subjects with approval rate of 95 or higher, less than 10k HITs and location in the US. We further apply the suspicious geocode block provided by Cloud Research. Furthermore, within each online platform, we excluded subjects who had already participated in previous experiments. Last but not least, experiments conducted on Prolific included a set of attention checks, a bot check and a language check. Subjects, who did not pass either one of these checks, were not allowed to proceed with the experiment. In all experiments, fewer than 5% of the participants exited the experiment on their own after being randomized into treatments. Hence, there is little room for selective attrition.

All experiments consist of an entry-level questionnaire that elicits demographic characteristics, in particular, age, gender, and self-reported statistical skills. Table CI outlines the demographic composition in all experiments. Experiments conducted on MTurk were programmed in oTree (Chen, Schonger, and Wickens, 2016) and ran in multiple batches to isolate potential day-of-week effects. Experiments conducted on Prolific were programmed in Qualtrics.

EXPERIMENT	Date	N	TREATMENTS	Stakes &	PLATFORM	SAMPLE	AVG	%
				Earnings		Selection	Age	Male
Large stakes	07/2021	295	One-Shot, Sequential, Hard Plan	50(x26) cents for investment, 50 cents for participation; average \$60 per hour	Prolific	Location US, UK or IE; approval rate > 97%; bot, attention and comprehension check passed	31	37%
Traditional stakes	02&06/ 2019	940	One-Shot, Sequential, Hard Plan, Soft Plan	10(x26) cents for investment, 50 cents for participation; average \$12 per hour	Mturk via Cloud Research	HITs approved <10,000; ap- proval rate > 95; Location US; sus- picious geocode block	35	58%
Robustness house money effect	07/2021	583	option- framed One-Shot, option- framed Sequential, option- framed Hard Plan	10(x26) cents for investment, 50 cents for participation	Prolific	Location US, UK or IE; approval rate > 97%; bot, attention and comprehension check passed	30	32%
Robustness house money effect, One-Shot only	07/2021	59	One-Shot, option- framed One-Shot (within- subject design)	10 cents for investment, 50 cents for participation	Prolific	Location US, UK or IE; approval rate > 97%; bot, attention and comprehension check passed	38	44%
Robustness expecta- tions	07/2020	314	Hard Plan, Hard Plan with beliefs induction	10x26 cents for investment, 50 cents for partic- ipation	Mturk via Cloud Research	HITs approved <10,000; ap- proval rate > 95; Location US; sus- picious geocode block	37	67%
Robustness wealth effect	05/2020	122	One-Shot, high-fee One-Shot, Hard Plan	10(x26)centsforinvestment,50centsforparticipationinOne-Shotand Hard Plan,300centsforparticipationinhigh-feeOne-ShotOne-ShotNot	Mturk via Cloud Research	HITs approved <10,000; ap- proval rate > 95; Location US; sus- picious geocode block	38	57%
Outcome frame	09/2020	160	Sequential with elic- itation of limits after feedback on the out- come of first round	10x26 cents for investment, 50 cents for partic- ipation	Mturk via Cloud Research	HITs approved <10,000; ap- proval rate > 95; Location US; sus- picious geocode block	36	66%

 Table CI

 Overview of Experimental Design and Execution

 Table CI

 Overview of Experimental Design and Execution

Experiment	Date	N	TREATMENTS	Stakes & Earnings	Platform	Sample Selection	Avg Age	% Male
Static skewness preferences	12/2022	148	One-Shot choice between investing in a positively skewed fair lottery, a negatively skewed fair lot- tery, or not investing	\$3 for invest- ment, 50 cents for participa- tion	Prolific	Location US, UK or IE; approval rate > 97%; bot, attention and comprehension check passed	39	51
Table CIIOverview of Experimental Results

This table provides an overview of the main findings of all experiments described in Appendix C. The experimental design is summarized in Table CI. Investment Rate refers to the probability of starting to invest in the first round. The probabilities are displayed separately for each treatment. For comparisons between treatments, we conduct Mann-Whitney tests and corresponding p-values are displayed in parentheses. Ex ante Plan refers to the share of participants who have a loss-exit, gain-exit, or equidistant neutral strategies among all participants in all Plan treatments (Hard Plan or Soft Plan). A loss-exit (gain-exit) strategy is defined as lower (greater) loss limit than gain limit. Continuation Probability (Immediate Gain minus Loss) reports the marginal effects of logit regressions of the probability of continuing to invest beyond the first round upon receiving feedback about the outcome in the first round. This is the cleanest test of deviation in behavior as cumulative gains and losses in subsequent rounds depend on participants' earlier choices. A positive coefficient indicates a higher probability for continuing beyond the first round after a gain than a loss (immediate gain/loss refers to gain/loss in round 1). Ex ante reports the regression results for the Plan treatments (Hard or Soft Plan), where subjects provide a fully contingent plan, hence we have information on their ex-ante planned choices after both immediate gains and immediate losses (strategy method). We simulate the cases of immediate gains and losses for each subject, resulting in two observations per subject. We cluster the standard errors of the corresponding logit regressions on subject level. Ex post reports the regression results for the Sequential treatment, where we only have one observation (either a gain or a loss) per subject. DID reports the difference in difference - the interaction effect of a logit model including a dummy for immediate gain (vs. loss), a treatment dummy for Sequential treatment (vs. Hard and Soft Plan treatment) and their interaction. As the choices of each subject in the Plan treatments are simulated for both cases of immediate gain and immediate loss, we cluster the standard errors at the subject level. *p*-values are in parentheses.

Experiment	Investment Rate		Ex ante Plan		Δ Continuation Probability	
					(Immediate Gain minus Loss)	
Large stakes	One-Shot	60.82%	Loss-exit	84.8%	ex-ante	+3.2pp
	Sequential	81.82%	Gain-exit	7.1%		(p = 0.428)
		\neq One-Shot	Equidistant	8.1%	ex-post	-18.9pp
		(p < 0.01)				(p = 0.034)
	Hard Plan	94.95%			DID	+22.1pp
		\neq One-Shot				(p = 0.018)
		(p < 0.01);				
		\neq Sequential				
		(p < 0.01)				

Experiment	INVESTMENT RATE		EX-ANTE PLAN		Δ Continuation Probability		
					(IMMEDIATE	GAIN MINUS LOSS)	
Traditional	One-Shot	67.68%	Loss-exit	80.35%	ex-ante	+9.5pp	
stakes	Sequential	85.00%	Gain-exit	6.99%		(p < 0.01)	
		\neq One-Shot	Equidistant	12.66%	ex-post	-20.5pp	
		(p < 0.01)				(p < 0.01)	
	Hard Plan	93.53%			DID	+30.0pp	
		\neq One-Shot				(p < 0.01)	
		(p < 0.01);					
		\neq Sequential					
		(p < 0.01)					
	Soft Plan	92.62%					
		\neq One-Shot					
		(p < 0.01);					
		\neq Sequential					
		(p = 0.044)					
		(F 0.011)					
Robustness	One-Shot	38.54%	Loss-exit	80.61%	ex-ante	+17.4pp	
house money	Sequential	67.18%	Gain-exit	8.67%		(p < 0.01)	
effect		\neq One-Shot	Equidistant	10.72%	ex-post	-11.2pp	
		(p < 0.01)	-			(p = 0.144)	
	Hard Plan	79.08%			DID	+27.5pp	
		\neq One-Shot				(p < 0.01)	
		(p < 0.01)				(4)	
		\neq Sequential					
		(n < 0.01)					
		(p < 0.01)					
Robustness	One-Shot	55.93%					
house money,	One-Shot	23.72%					
One-Shot only	(option-	\neq One-Shot					
	frame)	(p < 0.01)					
		(r (0.01)					
Robustness	Hard Plan	95.39%	Loss-exit	62.74%			
expectations	Hard Plan	95.67%	Gain-exit	11.46%			
	(w beliefs)	\approx Hard Plan	Equidistant	25.80%			
		(p = 0.903)					
		(z)					
1							

Experiment	Investment Rate		Ex-ante Plan		Δ Continuation Probability
					(Immediate Gain minus Loss)
Robustness	One-Shot	67.50%	Loss-exit	74.36%	_
wealth effect	One-Shot	55.81%	Gain-exit	12.82%	
	(high fee)	\approx One-Shot	Equidistant	12.82%	
		(p = 0.277)			
	Hard Plan	94.87%			
		\neq One-Shot			
		(p < 0.01);			
		\neq high fee			
		One-Shot			
		(p < 0.01)			
Outcome frame		95.63%			
Static skewness	Invest,	76.40%			_
preferences	thereof				
	in (+) skew	77.00%			
		\neq in (-) skew			
		(p < 0.01)			

Appendix D. Additional Analyses

Table DI

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.

	Panel A					
	Ex-ante strategy (loss-limit $= 1$)					
	(1)	(2)	(3)	(4)		
Constant	0.533^{***}	0.534^{***}	0.521^{***}	0.535^{***}		
female	(0.0050)	(0.0056) -0.00480 (0.012)	(0.0042)	(0.0074)		
trader location (Africa omitted) Asia			0.00158			
Europe			(0.0070) 0.00440 (0.0054)			
N Amer			(0.0054) 0.0225^{***}			
Oceania			(0.0058) 0.0206^{**} (0.0008)			
S Amer			(0.0038) 0.0920^{*} (0.048)			
trading experience ² (years)			(0.040)	0.00324^{**}		
trading experience (years)				(0.0014) -0.0102 (0.0069)		
R2 N	-8.1e-11 11,465,145	$\begin{array}{r} 0.000012 \\ 11,322,186 \end{array}$	$0.0022 \\ 11,437,756$	$0.00023 \\ 11,465,145$		

Panel B							
	Ex-ante strategy (loss-limit $= 1$)						
	(5)	(6)	(7)	(8)			
Constant	0.540***	0.531***	0.569***	0.582***			
long position	(0.0052) -0.0140*** (0.0039)	(0.014)	(0.0026)	(0.024)			
currency pair groups (EUR/USD omitted)	(0.0000)	0.00004					
USDpairs		-0.00224					
EURpairs		(0.014) 0.00874					
IDVnoire		(0.014) 0.0242*					
Jr I pairs		(0.0242) (0.014)					
position leverage (400:1 omitted)		()	0.000				
2:1			-0.222 (0.21)				
5:1			0.161***				
10:1			(0.058) -0.0608*** (0.0068)				
25:1			-0.0886***				
50:1			(0.0049) - 0.0566^{***}				
100:1			(0.0034) - 0.0408^{***}				
200:1			(0.013) -0.0163**				
log(position capital)			(0.0082)	-0.0158^{**} (0.0062)			
R2 N	0.00020	0.00019	0.0042	0.0021			
± '	11,100,110	11,100,110	11,100,110	11,100,110			

Table DITraders' Ex-Ante Strategies

Appendix E. Experimental Instructions

Appendix EI. Attention Checks (All Treatments)



Figure E1. Screenshots of bot, attention and comprehension checks To participate in the experiment subjects need to pass all checks.

Appendix EII. Welcome Screen (All Treatments)

Dear Participant,

Thank you for taking the time to participate in our study. The aim of this study is to gain insights about decision making under risk. The study takes approximately 9 minutes to complete.

You will receive financial remuneration for participating. The exact amount depends on your investment decisions. Further details will be explained at the beginning of the investment task. To receive your remuneration you need to complete the entire study. At the end of the study, you will receive a completion code you need to submit on the Prolific platform. Click "Next" to proceed.

Appendix EIII. Demographics (All Treatments)

Before we start with the experiment please answer the following questions. [Note: Your participation does not depend on your answers.]

• Your age

77

- Your gender
- Your highest level of education
- Your field of study
- How would you rate your statistical knowledge? Please choose a category between 1 ("very bad") and 6 ("very good").

Appendix EIV. One-shot Treatment

Screen: Instructions

You have 50 cents. You can choose to invest 50 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 50 cents, for a total of 100 cents. With a chance of 1/2 (50%) the lottery will "fail" and you will lose the 50 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 50 cents in the lottery.

Do you want to invest? $[\mathrm{Yes}]/[\mathrm{No}]$

Appendix EV. Sequential Treatment

Screen: Instructions (1)

You can choose to invest 50 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 50 cents, for a total of 100 cents. With a chance of 1/2 (50%) the lottery will "fail" and you will lose the 50 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 1300 cents in total to invest with. You can invest 50 cents per round in the lottery for up to 26 rounds. At the beginning you will choose whether or not to invest in the first round. After learning the outcome of your investment (whether you won or lost), you will choose whether to invest again or nor. You can stop investing at any time. Once you decide to stop investing, this part of the experiment will end.

Your earnings for this part of the experiment are as follows: At the end, we will count the number of rounds you have won (n) and the number of rounds you have lost (m). Your total gain or loss is given by the difference between these numbers multiplied by your investment per round which is 50 cents.

- If n ≥ m, you have earned a total gain of (n m)× 50 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)× 50 cents. In this case, you will receive the rest of your initial endowment after deducting the amount of your total loss.

Screen: Round 1 of 26: Investment decision

You can now choose whether or not to invest 50 cents in the lottery in round 1.

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

. . .

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.]

Appendix EVI. Hard Plan Treatment

Screen: Instructions (1)

You can choose to invest 50 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 50 cents, for a total of 100 cents. With a chance of 1/2 (50%) the lottery will "fail" and you will lose the 50 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 1300. [Note: Your earnings do not depend on your responses to this question.]

My first number:...

My second number:...

Screen: Instructions (2)

You have 1300 cents in total to invest with. You will choose whether or not to invest 50 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 50 cents, you will keep investing 50 cents in each subsequent round until your total gain or loss reaches your gain limit or loss limit respectively. In case either your loss or gain limit is reached, you will stop investing automatically.

- 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 $[y_1/50]$ 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 $[y_2/50]$ rounds more often than lost.

79

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 1300 cents plus the amount of your total gain of (n m)*times*50 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)×50 cents.

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.

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: Investment decision

You can now choose whether or not to start investing 50 cents in the lottery over a series of up to 26 rounds until you automatically stop.

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

80