

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

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

where

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

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

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

and value function

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

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 $\lambda \geq 1$.

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 $a \in \{\textit{continue}, \textit{exit}\}$.

Each s generates a random variable \tilde{G}_s , which corresponds to the accumulated gains or losses conditional on s being carried out. For example, take $s \in S_{0,1}$ where the decision maker plans to accept the first gamble, continue if she wins in $t = 1$ and then exits regardless of the next outcome in $t = 2$, and exiting in $t = 1$ if she loses the first gamble. This plan corresponds to $\tilde{G}_s \sim (-50, 1/2; 0, 1/4; 100, 1/4)$. 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 $t \in \{0, \dots, 26\}$.

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

The optimal plan s^* for each agent is the one that is connected to the outcome distribution with the highest expected value among all possible strategies, as given by the objective function in Equation A1. The agent accepts the sequential gamble if the expected value of the optimal strategy is higher than the value of exiting, which is normalized to zero. If the agent accepts the gamble in the first round, she revisits her decision in every subsequent round. For this purpose, the agent compares the expected value of continuing to accept the gamble, assuming that she will adhere to the ex-ante optimal strategy, with the value of exiting. The value of continuing to take on risk is determined by running 10,000 new simulations to determine the updated outcome distribution. The value of exiting is given by the value of the accumulated gains or losses since the beginning. We assume that the reference point is the initial endowment of \$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):

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

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 round. 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:

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

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

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

and

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

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:

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

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

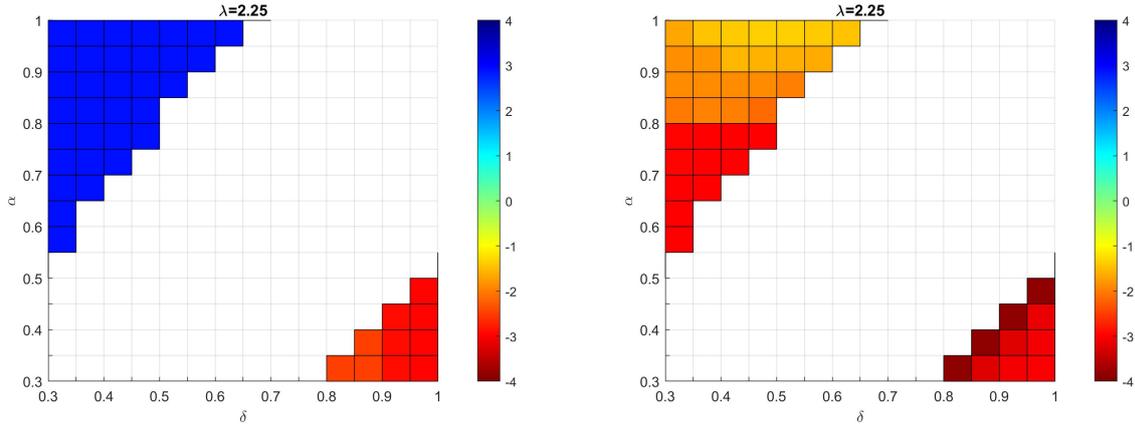
⁴⁵ For the case of gambles with positive expected value, Peköz (2002) shows that skewness preferences in combination with endogenous exit can explain the Samuelson paradox (Samuelson, 1963) as subjects will follow a “loss-exit” strategy to generate positive skewness. This is not the case for the fair 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 $\beta \in [0.6, 1]$.

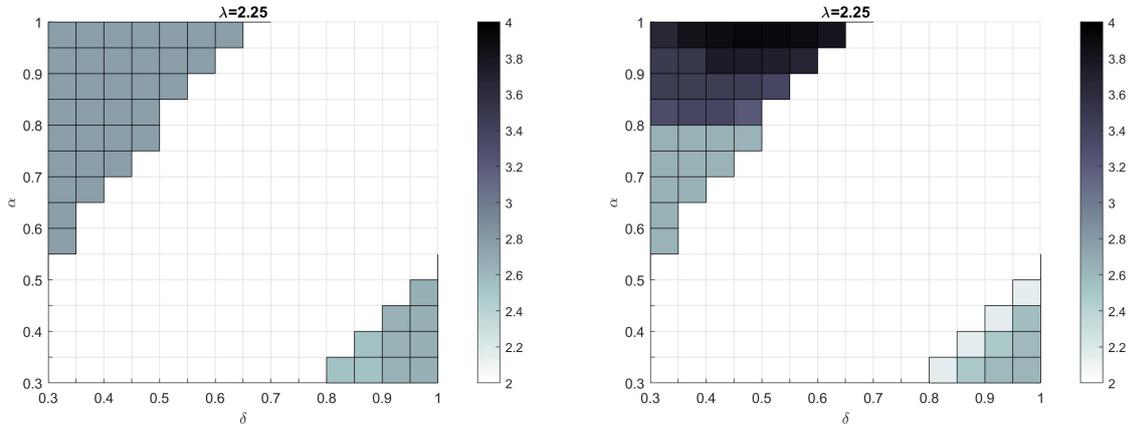
We find that an agents with quasi-hyperbolic discounting would not accept the first gamble in a sequence. Furthermore, even if the agent was forced to enter, the ex-ante plan for all parameter combinations implies exiting as soon as possible, independent of the outcome. It follows trivially that quasi-hyperbolic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

Panel B. Skew of ex-post behavior



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 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. 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).

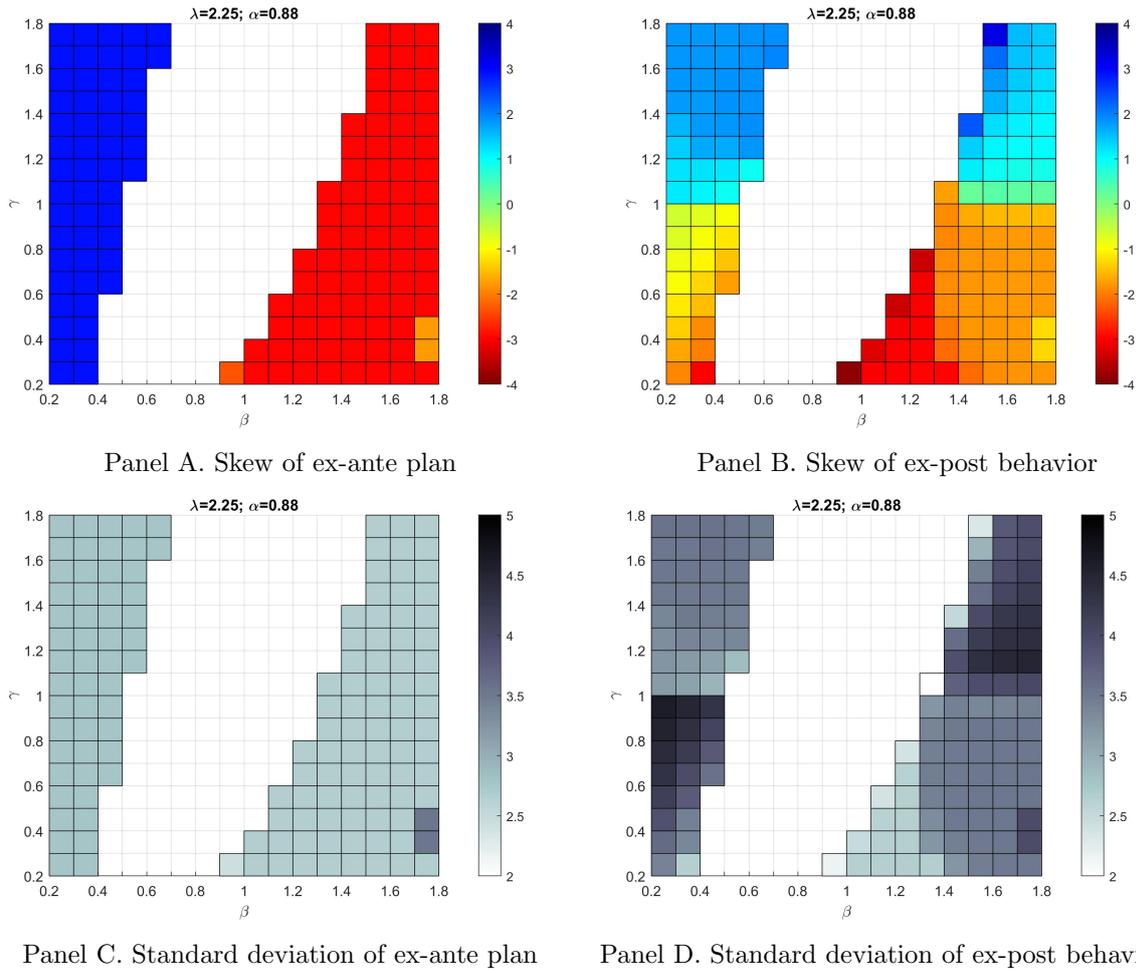
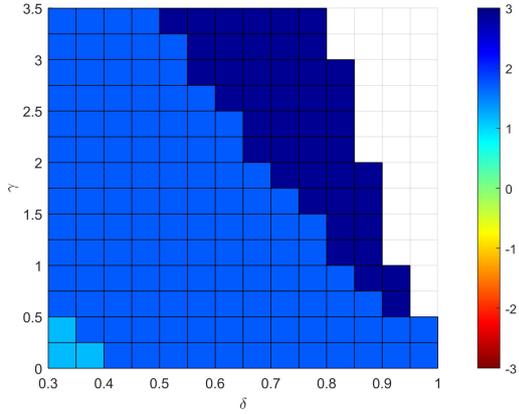
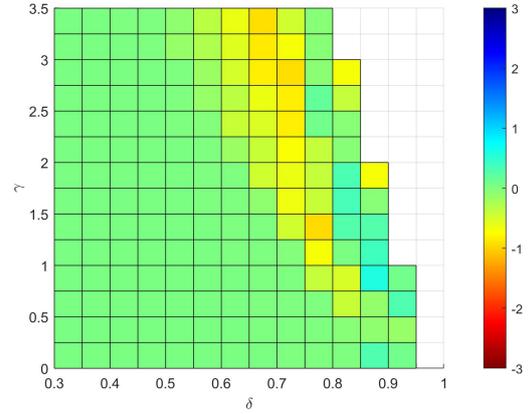


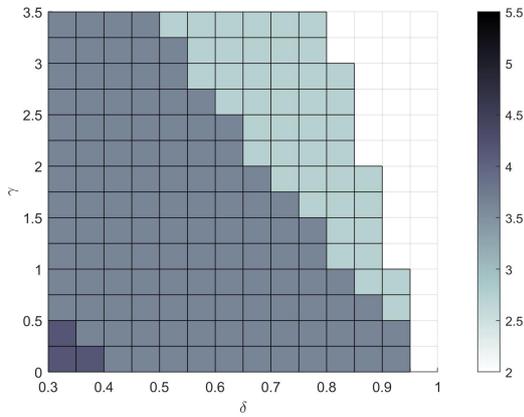
Figure A2. Theoretical predictions of Cumulative Prospect Theory with two-parameter probability weighting function. This figure illustrates the ex-ante plan (Panel A and C) and the ex-post deviation (Panel B and D) of agents with CPT preferences and a Gonzalez and Wu (1999) probability weighting function 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).



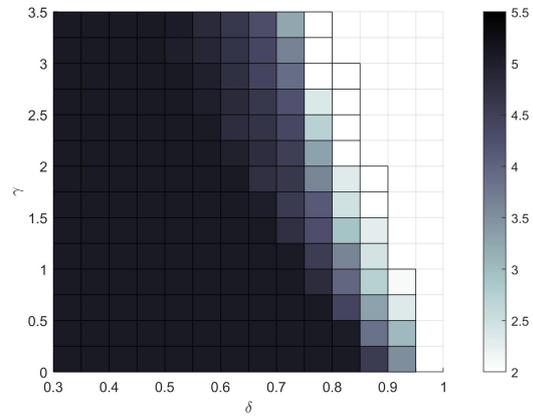
Panel A. Skew of ex-ante plan



Panel B. Skew of ex-post behavior



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 1 indicates that the agent invest in the lottery for one round, whereas 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 ex-ante strategy and the ex-post behavior, respectively. We obtain the outcome distributions from simulations in which the one-round investment amount is a numeraire (see Appendix A for further details about the simulations). Hence, certainty equivalents are measured as multiples of the one-round investment amount.

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.

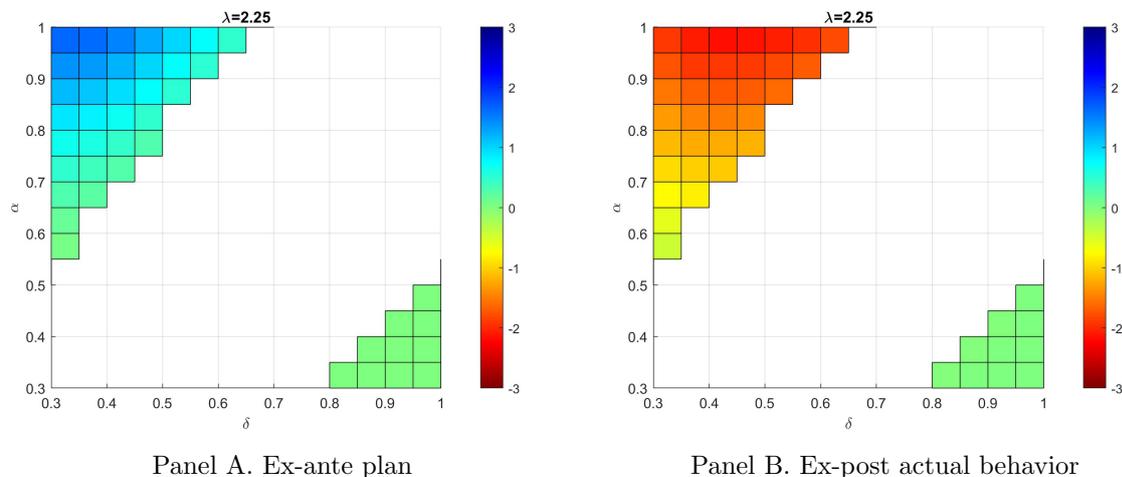


Figure B1. Welfare implications of dynamic inconsistency: value of commitment and costs of naïveté. This figure illustrates the certainty equivalent of agents with CPT preferences with different levels of probability weighting (δ) and diminishing sensitivity (α), i.e., the extent to which agents are 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.

Table CI
Overview of Experimental Design and Execution

EXPERIMENT	DATE	N	TREATMENTS	STAKES & EARNINGS	PLATFORM	SAMPLE SELECTION	AVG 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; approval rate > 95; Location US; suspicious 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 expectations	07/2020	314	Hard Plan, Hard Plan with beliefs induction	10x26 cents for investment, 50 cents for participation	Mturk via Cloud Research	HITs approved <10,000; approval rate > 95; Location US; suspicious geocode block	37	67%
Robustness wealth effect	05/2020	122	One-Shot, high-fee One-Shot, Hard Plan	10(x26) cents for investment, 50 cents for participation in One-Shot and Hard Plan, 300 cents for participation in high-fee One-Shot	Mturk via Cloud Research	HITs approved <10,000; approval rate > 95; Location US; suspicious geocode block	38	57%
Outcome frame	09/2020	160	Sequential with elicitation of limits after feedback on the outcome of first round	10x26 cents for investment, 50 cents for participation	Mturk via Cloud Research	HITs approved <10,000; approval rate > 95; Location US; suspicious geocode block	36	66%

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 lottery, or not investing	\$3 for investment, 50 cents for participation	Prolific	Location US, UK or IE; approval rate > 97%; bot, attention and comprehension check passed	39	51

Table CII
Overview 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%	ex-ante +3.2pp (<i>p</i> = 0.428) ex-post -18.9pp (<i>p</i> = 0.034) DID +22.1pp (<i>p</i> = 0.018)
	Sequential	81.82%	
	\neq One-Shot (<i>p</i> < 0.01)	Loss-exit 84.8% Gain-exit 7.1% Equidistant 8.1%	
	Hard Plan	94.95%	
	\neq One-Shot (<i>p</i> < 0.01); \neq Sequential (<i>p</i> < 0.01)		

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

EXPERIMENT	INVESTMENT RATE	EX-ANTE PLAN	Δ CONTINUATION PROBABILITY (IMMEDIATE GAIN MINUS LOSS)
Robustness wealth effect	One-Shot 67.50% One-Shot 55.81% (high fee) \approx One-Shot $(p = 0.277)$ Hard Plan 94.87% \neq One-Shot $(p < 0.01)$; \neq high fee One-Shot $(p < 0.01)$	Loss-exit 74.36% Gain-exit 12.82% Equidistant 12.82%	—
Outcome frame	95.63%	—	—
Static skewness preferences	Invest, thereof 76.40% 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.

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

Table DI
Traders' Ex-Ante Strategies

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

Appendix E. Experimental Instructions

Appendix EI. Attention Checks (All Treatments)

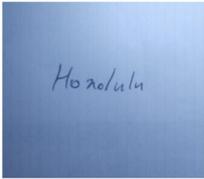
<p>Bot Screening</p> <p>Please answer the following questions to confirm that you are not a bot.</p> <p>Please click the box below.</p> <p><input type="checkbox"/> I'm not a robot </p>	 <p>Please type the word above into the space below.</p> <input type="text"/>	 <p>Please type the letter that is written on the picture of the strawberry:</p> <input type="text"/>
<p>Which of the following sentences most accurately describes you?</p> <p><input type="radio"/> I am @ lion</p> <p><input type="radio"/> I am @ rabbit</p> <p><input type="radio"/> I am @ human</p>	<p>Please read the passage below and answer the question that follows.</p> <p>Tatyana knew that Grandma always enjoyed serving an abundance of food to her guests. Now Tatyana watched as Grandma gathered Tatyana's small mother into a wide, scrawny embrace and then propelled her to the table, lifting her shawl from her shoulders, seating her in the place of honor, and saying simply: "There's plenty."</p> <p>Who does the 'her' in the above passage refer to?</p> <p><input type="radio"/> Tatyana's mother</p> <p><input type="radio"/> Grandma</p> <p><input type="radio"/> Tatyana</p>	

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

- 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? [Yes]/[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 not. You can stop investing at any time. Once you decide to stop investing, this part of the experiment will end.

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

- If $n \geq m$, you have earned a total gain of $(n - m) \times 50$ 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 $(m - n) \times 50$ 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.

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 $(m - n) \times 50$ 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]