

Eliciting Stopping Times

Sebastian Ebert

Maximilian Voigt[†]

September 7, 2026

ABSTRACT

We study how individuals plan dynamic risk-taking by eliciting stopping times—complete contingent plans of when to continue or stop taking a given risk. Despite substantial heterogeneity in stopping times, a machine-learning algorithm identifies a small number of economically meaningful strategy clusters. Many subjects use path-dependent and randomized rules, and trailing stop-loss strategies are more common than threshold strategies. Most subjects plan using forward rather than the normatively appropriate backward induction. Comparing planned stopping times with sequential risk-taking, we decompose dynamic inconsistency into the effects of memory constraints, defaults, mechanical inconsistencies induced by planning constraints, and planning as such.

Keywords: Stopping Times, Repeated Risk-Taking, Elicitation, Dynamic Consistency, Mechanical Inconsistency


JEL Classification: D01, D81, D91

[†] Sebastian Ebert is at the Department of Economics at Heidelberg University. E-mail: sebastian.ebert@awi.uni-heidelberg.de. Maximilian Voigt is at the Department of Finance at HEC Montréal. E-mail: maximilian.voigt@hec.ca. The experiments were pre-registered at the AEA RCT registry (AEARCTR-0009684 and AEARCTR-0013653).

I. Introduction

We propose a first experimental method to elicit unconstrained stopping times—the complete (time- and outcome-) contingent plan for taking a risk repeatedly. A large literature has studied risk-taking in static contexts and identified various aspects that affect the attractiveness of static risks (e.g., Kahneman and Tversky 1979). In this paper, we take the dual stance and consider a very simple static risk while focusing on dynamic aspects of risk-taking. A risk that is unattractive on its own may be attractive if taken repeatedly because the plan when to stop taking the risk can systematically affect outcomes. Stopping times and their characteristics have thus received significant attention in economic models and probability theory more generally, and an extensive recent literature studies “behavioral stopping.”¹

FIGURE 1.—Elicitation of unconstrained stopping times



Notes: The left panel of Figure 1 shows our representation of the repeated risk as an initially colored (green) ball triangle. The leftmost ball represents the first of five risks. In the second column, the upper (lower) ball corresponds to the second risk conditional on winning (losing) the first risk. Balls further to the right represent subsequent risks conditional on earlier outcomes. For example, the second-highest ball in the fourth column corresponds to the risk after two wins and one loss. Subjects *decolor* the ball triangle by clicking each ball in an order of their choice and thereby design their stopping time τ . The right panel provides an example. Risk-taking stops at black balls and continues at white balls.

We experimentally elicit unconstrained (i.e., time- and outcome-contingent) stopping times for taking a risk repeatedly for up to five times. We visualize the five-times repeated binary, fair, and symmetric risk as a ball triangle that represents a recombining binomial tree. The left

¹ Shiryaev (2007) provides a classical textbook treatment of optimal stopping under standard preferences. Björk, Khapko, and Murgoci (2021) provide a comprehensive overview of the recent, though already quite extensive, literature on “behavioral stopping” under non-standard preferences.

panel of Figure 1 showcases our experimental approach. Subjects design their preferred stopping time τ by removing the coloration from all initially colored balls of the ball triangle during the *Decoloring Task*. The right panel of Figure 1 shows an exemplary *decoloring*—a stopping time. Risk-taking stops at black balls and continues at white balls. If the Decoloring Task determines the subject’s payment, risk-taking is exogenously stopped according to the subjects’ decoloring. The Decoloring Task constitutes the first incentivized method for the elicitation of unconstrained stopping times. Methodologically, the incentive-compatibility of the Decoloring Task roots in the strategy method (Selten 1967) and our implementation is inspired by the behavioral stopping model of Barberis (2012) and He et al. (2017). After the Decoloring Task, subjects take the risk sequentially without precommitment.

In addition to time- and outcome-contingent risk-taking, we allow some subjects to design path-dependent or randomized risk-taking strategies, as motivated by theory (Shiryaev 2007; He et al. 2017; Hu, Oblój, and Zhou 2022). Although our experimental design allows for hundreds of behaviorally different risk-taking strategies, the graphical interface enables a parsimonious elicitation of even complex path-dependent or randomized risk-taking strategies. The order in which subjects make their choice for each contingency is endogenous and each action is visualized as part of the whole strategy, supporting subjects’ perception of the stopping time as a fully contingent risk-taking strategy beyond sequential risk-taking considerations.

The observed decolorings display known salient properties of repeated risk-taking, suggesting that the Decoloring Task elicits economically meaningful stopping times. Most subjects start taking a five-times repeated binary, fair, and symmetric risk, and do so using continue-when-winning and stop-when-losing (or stop-loss) strategies. While such behavior is inconsistent with risk-averse expected utility (EU) preferences, it is expected under non-EU preferences (e.g., Barberis 2012; Ebert and Strack 2015), and has been documented previously (Fischbacher, Hoffmann, and Schudy 2017; Strack and Viefers 2021; Heimer et al. 2025).

Facts about stopping times. Stopping times are complex, high-dimensional data objects. We use machine-learning methods and structural estimation techniques to analyze stopping times, and document three new facts about the chosen stopping times.

First, there is significant heterogeneity between subjects. We use an unsupervised machine-learning algorithm—that can identify arbitrary patterns in the data—to cluster the several hundred stopping times we observe. We find that stopping times are comprehensively described as stop-loss, take-profit, buy-and-hold, take-the-risk-once, and never-start strategies. Neither did the algorithm—by definition—learn these strategy types from training data, nor

did the experimental design restrict the strategies of subjects to certain types. Yet, we find economically interpretable clusters of unconstrained risk-taking strategies that have been predicted in theoretical work.

Second, we extend the non-parametric analysis of risk-taking strategies by proposing a method to structurally estimate prospect theory parameters from stopping times. To our knowledge, this is the first attempt to infer preference parameters from stopping time data. On average, the estimated parameters yield a prospect theory formulation with diminishing sensitivity, probability weighting, and loss aversion. We also estimate a mixture model to capture the heterogeneity in stopping times. The results map well to the strategy clusters identified through the unsupervised machine-learning algorithm, implying that prospect theory might be a good starting point for describing and modelling stopping behavior.

Third, path-dependence and randomization are both used often. About 55% of subjects use path-dependence at least once, while 80% of subjects use randomization at least once. We do not find a simple systematic pattern in the usage of path-dependence, which is more consistent with a revealed demand for randomization than with a direct preference for path-dependence. Randomized stopping probabilities are mostly below 50%, shifting risk-taking strategies closer to buy-and-hold strategies than without randomization. Neither the popularity nor the usage of path-dependence or randomization is systematically affected by making them costly, indicating a robust demand for these non-standard risk-taking strategies.

Applications. In the second part of the paper, we use the Decoloring Task to obtain new insights on several economic applications. First, we inform about subjects' mental models when planning repeated risk-taking. Having subjects choose the order of decoloring the ball triangle yields insights into their decision-making process. We find that about 60% of subjects choose their stopping time by forward induction as opposed to the normatively prescribed backward induction (27% of subjects).

Second, eliciting unconstrained risk-taking strategies allows us to analyze what kind of stop-loss strategies subjects use: We find that trailing stop-loss strategies—strategies with an increasing stopping barrier like the one in the right panel of Figure 1—are a better description of the observed stopping times and are used about 1.5 times as often as threshold stop-loss strategies (i.e., strategies with a constant stopping barrier). Using a dedicated experimental study, we find no significant change in aggregate measures of stopping times when subjects must choose a threshold strategy. The focus on threshold risk-taking strategies in theoretical,

empirical, and experimental work, therefore, seems to be a reasonable simplification in the aggregate, but constitutes a binding restriction for individual-level stopping times.

Third, comparing each subject's risk-taking plan to her sequential risk-taking actions during the second part of the experiment, we find the *standard shift in dynamic risk-taking* (Barberis 2012; Heimer et al. 2025): Subjects systematically deviate from their stop-loss plans by continuing longer when losing and stopping earlier when winning (the Disposition Effect; cf. Shefrin and Statman 1985; Odean 1998). We document that the standard shift prevails with unconstrained risk-taking plans and make another observation about dynamic inconsistency that may be unexpected: Subjects follow their plan for 89% of their risk-taking actions. This result exploits that our data contain all planned actions for all possible contingencies. Therefore, the standard shift (and thus the disposition effect) arises from a small number of systematic deviations rather than from an overall inconsistency. As the disposition effect lowers investors' returns due to transaction costs, taxes, or foregone time-series momentum, our findings should be informative to regulators aiming at reducing the disposition effect.

Fourth, we quantify the inconsistency between planned and sequential decisions that arises from constrained choice sets rather than from preferences. Subjects forced to choose a plan from a constrained set may deviate from it simply because the desired plan was not available. For example, a subject that must choose a threshold strategy deviates from the constrained-optimal choice if her unconstrained-optimal strategy is not a threshold strategy. Such a deviation does not constitute a preference reversal and “mechanically” increases dynamic inconsistency estimates. This fact – which we refer to as *mechanical inconsistency* – has been ignored in the literature thus far. Ex-ante, mechanical inconsistency could increase or decrease the standard shift. Accounting for mechanical inconsistency, we find that the standard shift remains significant although dynamic inconsistency estimates decrease by about 50%. The size of this bias shows that mechanical inconsistency effects can be severe, and thus our corroborating evidence of the standard shift is strong. In general, mechanical inconsistency (introduced and quantified here in a behavioral stopping paradigm) may not only cause severe overestimation of dynamic inconsistency but also overturn results in any experiment or field setting where subjects' choices among different choice sets, situations, or contexts are compared.

Fifth, we study the causes of dynamic inconsistency using additional between-subject studies and obtain a behavioral decomposition of the standard shift. The largest difference in stop-loss intensity arises when comparing the plans of subjects restricted to using threshold strategies with the risk-taking actions of subjects who are not required to submit a plan (as in

Cubitt, Starmer, and Sugden 1998; Heimer et al. 2025). We decompose this difference into five behavioral components. The principal cause of the standard shift is a lack of commitment, accounting for about 38% of the total effect. Memory constraints explain about 34%, while explicit planning accounts for about 15%. Mechanical inconsistency induced by a restriction to threshold strategies contributes the remaining 13%. Nudges through a plan-consistent default have no measurable impact.

By eliciting unconstrained stopping times, we hope to inspire experimental and empirical work on applied research questions. Many, if not most, risks that people face can be taken repeatedly and stopping problems abound. Examples include the sale of an asset, job search, real options, or the decision to conclude data collection. The experimental method introduced in this paper can be adapted for such purposes (see the Discussion section). We also hope to inform the fundamental and applied theoretical stopping literature. Given this large literature (summarized, for example, in the textbooks of Shiryaev 2007 or Björk, Khapko, and Murgoci 2021), it may be surprising that unconstrained stopping times have not been elicited yet.

Related literature. On a general level, we contribute to the literature on repeated risk-taking and stopping decisions. Classical experiments in this literature are Thaler and Johnson (1990), Redelmeier and Tversky (1992), Tversky and Shafir (1992), Gneezy and Potters (1997), or Barkan and Busemeyer (1999). A classical theory textbook on stopping is Shiryaev (2007) or, with a focus on economic applications, Dixit and Pindyck (1994). More recently, theoretical research analyzed stopping problems under non-standard preferences (Karni and Safra 1990; Barberis 2012; Björk, Khapko, and Murgoci 2021; Heidhues and Strack 2021; He and Zhou 2022). We seek to inform this literature by proposing and implementing an experimental method to elicit unconstrained stopping times, and we provide the first experimental evidence on the unconstrained stopping times that subjects use.

The experimental stopping literature mostly focuses on sequential risk-taking (e.g., Langer and Weber 2008; Magnani 2015; Imas 2016; Nielsen 2019; Strack and Viefers 2021). Barkan and Busemeyer (1999) elicit incentivized plans for gambling once, after a gain or loss, respectively, and find that subjects gamble more than planned after losses. More recently, researchers elicited stopping times for longer horizons and under constraints suitable for the application under study.² Fischbacher, Hoffmann, and Schudy (2017) show that commitment

² Previous research that elicited risk-taking strategies typically focused on threshold strategies due to their analytical tractability (Dixit 1993; Xu and Zhou 2013; Ebert and Strack 2015; Magnani 2015; Henderson, Hobson, and Tse 2018) and suitability for experiments even with continuous-time processes.

to threshold stop-loss and take-profit strategies reduces the disposition effect. Heimer et al. (2025) documented the above-described standard shift for threshold strategies in experimental and observational data. To test the implications of salience theory (Bordalo, Gennaioli, and Shleifer 2012), Dertwinkel-Kalt and Frey (2024) show that subjects start taking a risk repeatedly with a stop-loss plan even if the stochastic process has a negative drift. In contrast, subjects in the experiment by Antler and Arad (2023) often select a stop-gains strategy among five pre-selected threshold strategies. Magnani et al. (2022) elicit time- and outcome-dependent as well as path-dependent stochastic control (rather than stopping) plans to study dynamic portfolio choices.³ We contribute to this literature by proposing the Decoloring Task as a method to elicit unconstrained stopping times and by documenting properties of unconstrained stopping times in our data. In addition to designing path-dependent risk-taking strategies, we allow some subjects to use a randomization device, thereby extending previous work on randomized risk-taking in static contexts (Agranov and Ortoleva 2022) to a dynamic setting. The extension of our method to other risks, like skewed risks (Ebert and Köster 2026) or risks with time-changing moments (Nielsen 2019), is well possible.

We also contribute to the growing literature in microeconometrics that uses machine-learning methods to analyze and interpret data from experiments (see, for example, Fudenberg and Liang 2020; Fudenberg et al. 2022). We show how supervised and unsupervised machine-learning methods can be used to analyze stopping times and to structure between-subject heterogeneity. We combine machine-learning methods with a structural estimation of cumulative prospect theory preference parameters (Tversky and Kahneman 1992; Wakker 2010) by extending the method in Bruhin, Fehr-Duda, and Epper (2010).

We further contribute to the literature on dynamic inconsistency in stopping decisions, which is theoretically predicted under non-standard preferences (Machina 1989; Cubitt, Starmer, and Sugden 1998; Barberis 2012; Ebert and Strack 2015; He and Zhou 2022). Our first contribution here concerns the comparison of unconstrained risk-taking strategies (possibly allowing for path-dependence or randomization) to sequential risk-taking actions when assessing dynamic consistency. We document that subjects in our study follow their risk-taking strategies at most balls. Relatedly, Nielsen and Rehbeck (2022) find that subjects follow canonical (static) choice axioms if they previously expressed consent with the axiom. The emerging pattern suggests that individuals are more likely to adhere to a plan, rule, or norm if they could choose the rule or express agreement therewith. Second, we are, to our knowledge,

³ Andrade and Iyer (2009) and Imas (2016) ask for a similar plan without incentivizing it.

the first to discuss and investigate the possible overestimation of dynamic consistency that arises from confounding dynamic and mechanical inconsistency. A third contribution to the dynamic inconsistency literature is the behavioral decomposition of the standard shift, with ensuing policy implications. Consistent with prior work that highlights the importance of memory for individual-level risk-taking (Gödker, Jiao, and Smeets 2025; Jiang et al. 2025; Gödker, Odean, and Smeets 2026), we find that a memory-aid reduces dynamic inconsistency, as does planning as such.

II. Overview of experimental design

Eliciting unconstrained stopping times requires a controlled experiment since complete contingent plans are not observable otherwise. Our method enables a parsimonious and intuitive elicitation of unconstrained risk-taking strategies through the *Decoloring Task* and allows for comprehensive comparisons with sequential stopping behavior through the *Walking Task*. An overview video (<https://youtu.be/ParT-P9gX04>) and an interactive version of the experiment are available.⁴

FIGURE 2.—Schematic storyboard of the experiment



Notes: Subjects first receive general and task-specific instructions, accompanied by comprehension questions. We then elicit the precommitted stopping time in Task 1, the Decoloring Task. Subjects then receive another short instruction for Tasks 2 – 11, the Walking Tasks, and complete these ten tasks. The experiment ends with a questionnaire and subjects learn their payment. We include one attention check during each set of instructions.

II.A Structure of the experiment

The experiment consists of ten between-subject *studies* with different versions of the experimental tasks. In each task, subjects are endowed with £1.50 (~\$1.85) and decide about taking a 50:50 risk of winning or losing £0.25 for up to five times. Following Barberis (2012),

⁴ The interactive version of the experiment is available at <https://rtsexpdemo.onrender.com> and further videos of all experimental studies in Table 1 can be found at https://www.youtube.com/playlist?list=PLr_xX-ToY8MNLtLLa3176gLOlw8rcNLmy. Appendix SA.5 shows further details. The experiment was programmed using oTree (Chen, Schonger, and Wickens 2016).

we represent *stopping times*—the complete (time- and outcome-) contingent plan for taking the risk for up to five times—as a recombining binomial tree, visualized and described as a ball triangle to subjects (see Figure 1).

Figure 2 illustrates the course of the experiment. For nine of the ten between-subject studies, the first task of the experiment is a *Decoloring Task*. In the Decoloring Task, each subject uses an interactive tool to design her precommitted risk-taking strategy by decoloring the balls of an initially colored ball triangle.⁵ The second part of the experiment consists of ten *Walking Tasks*. In each Walking Task, subjects sequentially decide whether to take (another) risk depending on how the previous risk(s) turned out. One study consists solely of ten Walking Tasks, without a Decoloring Task. All subjects go through interactive instructions, practice elements, and comprehension questions before each task. They must pass attention checks and complete a short questionnaire before receiving their payment, which is determined by the subject’s behavior in one randomly selected task (see Appendix SA.5).

TABLE 1.—Overview over the ten experimental studies

#	Name	Decoloring Task	Walking Task	Session date	<i>N</i>
1	Basic I	Basic	Basic	Nov 16, 2022	73
2	Path-dependence	Path-dependent	Basic	Nov 16, 2022	75
3	Path-dependence ^C	Path-dependent ^C	Basic	Nov 16, 2022	67
4	Randomization	Randomized	Randomized	Nov 16, 2022	88
5	Randomization ^C	Randomized ^C	Randomized ^C	Nov 16, 2022	97
6	Basic II	Basic	Basic	May 23, 2024	74
7	No default	Basic	No default	May 23, 2024	85
8	No reminder	Basic	No reminder	May 23, 2024	81
9	Threshold	Threshold	Basic	May 23, 2024	78
10	No planning	-	No reminder	May 23, 2024	163

Notes: Table 1 shows the different between-subject studies. The second column shows the name that we use to refer to a specific study, while the third and fourth columns show the version of the Decoloring and Walking Task used in the study. We ran the studies in two sessions, and column five shows when the respective study took place. Column six shows the number of subjects for each study. See Appendix SA.5.1 for details on the subject pool.

II.B Studies

Table 1 provides an overview of the ten between-subject studies. In a *Basic* task, subjects can decide to stop or continue risk-taking at any ball of the ball triangle. In a *Path-dependent* task,

⁵ We refer to the task and the subject’s interaction with the tool as *decoloring*: the ball triangle initially has a color (green, yellow, or red), and the subject must transform it into a greyscale ball triangle.

subjects can additionally condition on the path taken through the ball triangle; for example, a subject’s risk-taking strategy could be to stop risk-taking if the first risk is won and the second risk is lost but continue risk-taking if the first risk is lost and the second risk is won (up-down vs. down-up). Note that sequential risk-taking during the Walking Tasks is necessarily path-dependent and thus there is no dedicated path-dependent Walking Task (details in Section IV.C).⁶ Subjects in a *Randomized* task have access to an independent randomization device that allows them to stop risk-taking with a probability of their choice.⁷ The *Path-dependent^C* and *Randomized^C* tasks are analogous to their counterpart described above, except that path-dependence and randomization costs £0.01 per ball, respectively. Subjects in a *Threshold* task can only condition on the profit level but not on the risk repetition. They can, for example, stop risk-taking for all balls that yield a profit of £0.25 and continue otherwise. Finally, we vary the availability of a default risk-taking action and of a strategy reminder in the *No default* and *No reminder* Walking Tasks (details in Section IV.E).

II.C Eliciting stopping times with the Decoloring Task

Figure 3 illustrates the interactive tool that subjects use to design their stopping times during the Decoloring Task. The data we obtain from the Decoloring Task for each subject—a decoloring of a ball triangle—specifies a complete contingent plan of when to continue or stop taking the risk for up to five times, also called risk-taking strategy or stopping time.

The first (second, third) row of Figure 3 is for the *Basic (Path-dependent, Randomized)* Decoloring Task. Following Johnson et al. (2021), subjects make their decisions by giving instructions to the computer.⁸ Each subject in a given study clicks a ball in the right half of the instruction panel—shown in panel (a) of Figure 3—and gives risk-taking instructions for that ball in the left half of the instruction panel. She proceeds like this until all balls are decolored and can revise the decoloring of any ball as often as desired. If the Decoloring Task is the payoff-relevant task, a subject’s bonus payment is determined by taking the risk up to five


⁶ All *Basic* Walking Tasks are identical in that subjects can continue or stop risk-taking at any ball that is reached. During the Walking Tasks, the subjects sequentially resolve the risk and thereby realize a specific path to any given ball. The risk-taking action at a given ball is thus inherently conditional on the path taken to that ball. Consequently, subjects cannot use a costly option during the *Basic* Walking Tasks in study *Path-dependence^C*.

⁷ We simplify the task by discretizing the grid of stopping probabilities to multiples of 10 percent, such that subjects can stop risk-taking with 0, 10, ..., 90, 100 percent probability at any ball.

⁸ Subjects here give instructions to the computer and not to the experimenter. First, the computer is the recipient of the instructions and implements them, such that we do not deceive subjects. Second, as computers follow instructions given to them, this frame highlights the decisiveness of the instructions.

times according to the decoloring. The Decoloring Task is thus an incentivized method to elicit stopping times under precommitment.

FIGURE 3.—Decision screens and median decolorings elicited in the Decoloring Task



Notes: Subjects in a *Basic* (*Path-dependent*, *Randomized*) Decoloring Task must decolor all 15 balls of an initially green (yellow, red) ball triangle. Subjects click a ball in the right half of the instruction panel (a) and give the risk-taking instruction for that ball in the left half of the instruction panel (a). The selected ball turns blue and pulses. The selected ball is decolored white (black) if the subject unconditionally continues (stops) risk-taking at that ball. In *Path-dependent* and *Randomized* tasks, subjects can also decolor balls to become black-white gradient (indicating path-dependence) or numbered (the number indicating the stopping probability). Each subject proceeds like this for each ball in an order of her choice. Subjects can go back and change their decoloring. Panel (b) illustrates the median decolorings observed in the experiment, whereby the median is taken over the effective decoloring at each ball for all subjects.

The decoloring scheme is as follows: If a ball is white, the subject takes the 50:50 risk at that ball. If a ball is black, the subject stops risk-taking. Balls with a black-white gradient indicate that stopping at that ball is path-dependent in that the subject continues risk-taking for at least one path leading to that ball and stops risk-taking for at least one other path. A numbered

ball shows that the subject uses an independent randomization device that yields stopping with the probability shown on the ball.

In the analysis, we focus on incentivized decisions by considering “effective decolorings,” for which we treat decisions at balls that cannot be reached when following the risk-taking strategy as ‘stop risk-taking,’ consistent with Barberis (2012). We pool subjects from different studies with the same Decoloring Task.⁹ Details on the *Threshold* study and on the Walking Tasks will be given in Section IV.

III. Evidence on stopping times

We obtain the first dataset of unconstrained stopping times through the Decoloring Task. This section analyzes the stopping times.

III.A Basic Decoloring Tasks

Panel (b) of Figure 3 shows the aggregate precommitted risk-taking strategies for the *Basic*, *Path-dependent*, and *Randomized* Decoloring Task, respectively. We first focus on the median *Basic* decoloring obtained from $N = 313$ subjects. The median is taken ballwise, such that a ball is black if more than 50% of the subjects stop risk-taking at that ball and white otherwise.


Frequent risk-taking. The median decoloring indicates risk-taking (white first ball), continue-when-winning (the upper part of the triangle being white), and stop-when-losing (the lower balls being black). 96% of subjects take the risk at least once; 3.82 times on average. The large share of subjects that begin taking a mean-zero risk in our experiment is at odds with risk-averse expected utility (EU) preferences, but expected under non-EU preferences (e.g., Barberis 2012; Ebert and Strack 2015). The fact that the risk can be taken more than once is crucial for this theoretical prediction, and our evidence for it is in line with other experiments on repeated risk-taking (Fischbacher, Hoffmann, and Schudy 2017; Strack and Viefers 2021; Heimer et al. 2025).

Trailing stop-loss. The median strategy is of the loss-exit type defined by Barberis (2012): A strategy is *loss-exit* if the expected number of risks taken conditional on stopping at a loss ball is less than the expected number of risks taken conditional on stopping at a gain ball. Similarly, a strategy is *gain-exit (symmetric)* if this number is larger (the same). More

⁹ Whenever present, the Decoloring Task is each subject’s first experimental task and not affected by differences in subsequent Walking Tasks. We found no statistically significant differences in the decolorings of subjects in different studies with the same Decoloring Task, highlighting the robustness of our findings especially when considering the different execution dates of the studies (see Table 1).

specifically, the median strategy is a level-3 *trailing stop-loss*.¹⁰ Stopping occurs after losing thrice (level-3; ignoring that there is a final period) so that the stopping barrier trails profits. We can reject that aggregate risk-taking is consistent with a *threshold stop-loss*: The median risk-taking strategy would be a threshold stop-loss strategy if the bottom ball in the third column was black. 58% of the subjects decolor the bottom ball in the third column white, which is significantly different from 50% ($t(312) = 2.68$).

FIGURE 4.—Summary statistics of the precommitted risk-taking strategies



Notes: Panel (a) shows the distribution of the correlation between the stopping time τ and the induced outcome distribution S_τ . Panel (b) shows the share of loss-exit, gain-exit, and symmetric-exit strategies.

Continue-when-winning. On an individual level, we also find that most strategies are stop-loss strategies. Panel (a) of Figure 4 shows the correlation ρ between the stopping time τ and its induced outcome distribution S_τ for the different Decoloring Tasks. A positive correlation indicates continue-when-winning (or, likewise, stop-when-losing; i.e., a stop-loss strategy). For the *Basic* Decoloring Task, we find a positive average correlation ($\rho = 0.35, t(312) = 13.71$), which is non-negative for 75% of all individual risk-taking strategies. The average expected stopping time conditional on stopping with a gain is 4.41 and significantly higher than the expected stopping time conditional on stopping with a loss (3.54, $t(312) = 12.54$). Panel (b) of Figure 4 shows that a majority of 61% of the subjects in the *Basic* Decoloring Task use a loss-exit strategy. Heimer et al. (2025) document a popularity of loss-exit plans among threshold strategies. We find that loss-exit strategies even emerge endogenously when

¹⁰ Stop-loss strategies are a subset of loss-exit strategies, though we use the terms interchangeably when no confusion arises. Threshold and trailing stop-loss strategies are specific forms. Similarly, stop-gain or take-profit strategies are a subclass of gain-exit strategies.

allowing for unconstrained risk-taking strategies, and that many of them are of the trailing type (see Section IV.B).

By the optimal stopping theorem, the expected earnings of any risk-taking strategy are zero, and thus, risk-averse EU preferences are incompatible with the results of our experiment. The induced outcome distribution S_τ is, on average, positively skewed (skew = 0.36, $t(312) = 12.24$), in line with Ebert and Köster's (2026) result that relates a positive ρ to a positively skewed S_τ . As observed by Barberis (2012), taking risks repeatedly in a way that skews the profit distribution to the right is consistent with a wide range of cumulative prospect theory (CPT) preference parameters (see Section III.D).

Stopping times elicited from the *Basic* Decoloring Task are qualitatively consistent with predictions from non-EU models—foremost CPT—and with prior experimental findings. They are also internally consistent (cf. footnote 9 and section V.B), suggesting that the graphical interface of the Decoloring Task allows for a parsimonious elicitation of economically meaningful risk-taking strategies.

III.B Path-dependent and Randomized Decoloring Tasks

In addition to time- and outcome-contingent risk-taking, subjects in the *Path-dependence* study ($N = 75$) can costlessly condition their risk-taking strategy on the path taken to the ball and subjects in *Randomization* ($N = 88$) can condition their risk-taking strategy on an independent randomization device.

We elicit risk-taking strategies for these strategy spaces for two main reasons: (1) Path-dependent as well as randomized strategies are theoretically predicted. A decision-maker with law-invariant preferences that are not quasi-convex—as is the case for rank-dependent utility (Quiggin 1982), prospect theory (Kahneman and Tversky 1979; Tversky and Kahneman 1992), or the dual theory of choice (Yaari 1987)—strictly prefers randomized strategies over path-dependent strategies, which she in turn prefers over *Basic* strategies (He et al. 2017; Henderson, Hobson, and Tse 2017; Hu, Obłój, and Zhou 2022).¹¹ (2) Path-dependence and randomization are empirically relevant. The house money effect (Thaler and Johnson 1990) is a classic example of path-dependent risk-taking. Stopping behavior as well as the evaluation of risks often depend on the path of outcome realizations (Rabin 2002; Grosshans and Zeisberger 2018; Strack and Viefers 2021), and price paths affect stock returns (He and Liang 2024).

¹¹ Note that path-dependent and randomized risk-taking strategies do not improve upon the path-independent and randomized risk-taking strategies for which we allow in *Randomization*, see He et al. (2017).

Randomized risk-taking, in addition, is a longstanding finding in experimental research (Mosteller and Noguee 1951; Machina 1985) that has recently been characterized as a deliberate choice rather than a mistake (Agranov and Ortoleva 2017, 2022). Optimal risk-sharing and competitive equilibria require exogenous randomization when some individuals are risk-seeking (Lauzier et al. 2024). High-stakes decisions such as study programs (Dwenger, Kübler, and Weizsäcker 2018), charitable donations (Zhang and Zhong 2020) and important life choices (Levitt 2021) are often randomized. We give a first analysis of path-dependence and randomization in stopping times, whether subjects use them, and (if so) how they affect risk-taking strategies.

Path-dependence and randomization are both used often. Path-dependence can be used at the six inner balls of the ball triangle, as all outer balls can only be reached via one path (always up or down). 55% of subjects use path-dependence at least once and 78% of those subjects use it more than once.

No systematic patterns in the usage of path-dependence. We do not observe a simple systematic pattern in the usage of path-dependence. Subjects are neither more likely to stop risk-taking for paths that reach a ball from above or below (see Figure SA.1.2), nor to stop more frequently when reaching the highest or lowest earnings along a path. Thus, path-dependence might be used to achieve randomized stopping decisions (we elaborate on this in Section V.D). Consistently, we find that 81% of subjects in *Randomization* use at least one randomization, and more subjects randomize at all 15 balls than at none of the balls (30% vs. 19%, respectively). Making path-dependence or randomization costly does not affect their utilization (see Figure SA.1.1).¹²

Stop-loss strategies continue to be popular when allowing for *Path-dependence*. The second row of Panel (b) in Figure 3 shows the median risk-taking strategy for $N = 75$ subjects in *Path-dependence*. The median *Path-dependent* decoloring resembles the median *Basic* decoloring as it shows risk-taking (first ball being white for 92% of subjects), and a stop-loss strategy. The median strategy is (roughly) a trailing stop-loss strategy, in which subjects stop risk-taking after two to three losses. On an individual level, most subjects in *Path-dependence* use a loss-exit strategy (panel (b) of Figure 4), with a higher share of gain-exit strategies than in *Basic*. Overall, path-dependent strategies are behaviorally similar to *Basic* strategies.

¹² The frequency of path-dependent balls is statistically identical ($\chi^2(6) = 3.898$) in *Path-dependence* and *Path-dependence*^C, while the frequency of randomized balls differs at the 10%-level ($\chi^2(15) = 22.487$) for *Randomization* and *Randomization*^C.

Randomization strongly affects risk-taking toward buy-and-hold. The bottom row of Panel (b) in Figure 3 shows the median risk-taking strategy for $N = 88$ subjects in *Randomization*. Most subjects start risk-taking (white first ball) and design risk-taking strategies with a higher stopping probability in the loss domain than in the gain domain. 93% of subjects start risk-taking with a positive probability and we find an average probability of taking at least one risk of 83%. The median decoloring is of a stop-loss type in that the balls in the bottom of the ball triangle tend to be darker than in the upper part, but less strongly so than in *Basic* or *Path-dependence*. An interesting observation is that subjects in *Randomization* often use low stopping probabilities (see Figure SA.1.2). Panel (b) of Figure 4 shows a comparably high popularity of gain-exit strategies in *Randomization*.

III.C Individual-level analysis of strategy types using unsupervised clustering

We now examine the individual-level heterogeneity in unconstrained stopping times by using an unsupervised machine-learning algorithm, specifically a K -means clustering.¹³ The approach in this section is data-driven, since we do not formulate any strategy types or properties of risk-taking strategies as input for the algorithm. Instead, the algorithm partitions the elicited risk-taking strategies into clusters which may or may not represent familiar or economically interpretable strategies. Appendix SA.3.1 contains methodological details and robustness results.


Interpretable heterogeneity in stopping strategies. The clusters suggested by the K -means clustering have a natural economic interpretation in terms of stop-loss, take-profit, buy-and-hold, take-the-risk-once, or never-start strategies. Figure 5 shows the median decoloring of all strategies assigned to a cluster as well as the number of loss-exit, gain-exit, and symmetric-exit strategies it contains for $K = 6$ clusters.¹⁴ Visually, the median strategies of clusters 1 and 2 are time-dependent (buy-and-hold) strategies, and a relative majority of risk-taking strategies in these clusters are symmetric-exit strategies. Whereas the median strategy

¹³ Unsupervised machine-learning algorithms find patterns in unlabelled data. In contrast, supervised machine-learning methods learn a mapping between data $\mathbf{X}$ and output vector $\mathbf{y}$ that is then applied to new, previously unseen data.

¹⁴ Intuitively, a K -means algorithm partitions data into $K \in \mathbb{N}^+$ clusters. We conduct the analysis here with $K = 6$ clusters and use all $N = 640$ elicited risk-taking strategies for simplicity. Appendix SA.3.1 shows the robustness of our results towards using more clusters and towards conducting the analysis separately for *Basic*, *Path-dependent*, and *Randomized* risk-taking strategies. In unreported analyses, we find that the clusters exhibit greater within-cluster similarity and between-cluster separation than matched placebo datasets, suggesting that they capture genuine strategy types (see also Table SA.2.1).

of cluster 1 (27.7% of subjects) is a long-horizon buy-and-hold strategy, that of cluster 2 (12.5% of subjects) is noisily taking the risk once.¹⁵ The median strategy of cluster 3 (19.1% of subjects) appears to be a mixture of a buy-and-hold and a take-profit strategy, as the balls in the upper part of the ball triangle are darker than in the lower part. Cluster 3 also contains more gain-exit than loss-exit strategies (63 vs. 39). Thus, we interpret cluster 3 as a take-profit cluster. The median risk-taking strategies of clusters 4, 5, and 6 (10.8%, 14.7%, and 15.3% of subjects, respectively) are trailing stop-loss strategies, each with a different trailing stop-loss level and a varying blur of the stopping barrier. Almost all strategies in these three clusters are loss-exit strategies. Our visual interpretation of the clusters is supported by an analysis of the cluster constituents, as shown in Table SA.2.1.

FIGURE 5.—Median precommitted risk-taking strategy for each cluster of a K –means clustering



Notes: Figure 5 plots the median decoloring per cluster resulting from a K -means clustering with $K = 6$ clusters. We also show the number of loss-exit, gain-exit and symmetric exit strategies as well as the total number of risk-taking strategies in each cluster below the cluster number. The format is (loss-exit + gain-exit + symmetric-exit = N subjects).

An unsupervised machine-learning algorithm partitions unconstrained risk-taking strategies into economically reasonable clusters. Neither did the algorithm—by definition—learn strategy types from training data, nor did the experimental design restrict the risk-taking strategies of subjects to certain types. Nevertheless, we find a clustering of risk-taking

¹⁵ The analysis in Appendix SA.3.1 shows that when allowing for more than $K = 6$ clusters, Cluster 2 splits into a never-start cluster and short-horizon buy-and-hold strategies. We therefore do not classify this cluster as buy-and-hold when reporting that 28% of all strategies are buy-and-hold.

strategies into types that have previously been identified, analyzed, and discussed in the economics, finance and probability theory literature. We take this as support for theoretical and empirical work that focuses on a set of simple strategies.

III.D Structural estimation of CPT from stopping time data

Our findings that (a) most subjects start risk-taking although the symmetric risk has an expected value of zero and (b) most subjects' risk-taking strategies induce positively skewed outcome distributions S_τ , as predicted by the CPT (Tversky and Kahneman 1992) model of Barberis (2012). We also find (c) significant heterogeneity in stopping times that can be represented by a low number of strategy types (Figure 5). We show below that CPT predicts (c) as well. We then estimate CPT parameters using stopping time data to inform the increasing number of models that apply CPT in dynamic settings and help with model calibration. We identify types of CPT agents that map well to the non-structural cluster analysis (Section III.C). CPT, as a theory of choice under risk, is flexible enough to rationalize most of the elicited stopping times, and the estimated types might be a good starting point for describing and modelling stopping behavior.

Theoretical benchmark. Before describing our methodology, we briefly review CPT. A decision maker considers a lottery $L = (x_{-m}, p_{-m}; \dots; x_{-1}, p_{-1}; x_0, p_0; x_1, p_1; \dots; x_n, p_n)$ that yields outcome x_i with probability p_i . In our setting, $x_{-m} = -1.25$, $x_n = 1.25$, $x_{i+1} - x_i = 0.25$ for all $i = -m, \dots, n-1$, and we assume $x_0 = 0$. CPT posits that the lottery's value is given by:

$$V(L) = \sum_{i=-m}^n \pi_i v(x_i - x_0),$$

with

$$\pi_i = \begin{cases} w(p_i + \dots + p_n) - w(p_{i+1} + \dots + p_n), & \text{for } 0 \leq i \\ w(p_{-m} + \dots + p_i) - w(p_{-m} + \dots + p_{i-1}), & \text{for } 0 > i \end{cases}$$

We assume the probability weighting function as in Tversky and Kahneman (1992) with $w(p) = p^\delta / (p^\delta + (1-p)^\delta)^{1/\delta}$ and an exponential value function (Köbberling and Wakker 2005)¹⁶


¹⁶ The exponential value function is econometrically more robust than the commonly used power specification, which confounds the loss aversion and diminishing sensitivity parameters (Köbberling and Wakker 2005).

$$v(x) = \begin{cases} \frac{1 - e^{-(1-\alpha)x}}{1 - \alpha}, & \text{for } x \geq 0 \\ \lambda \left(\frac{e^{(1-\alpha)x} - 1}{1 - \alpha} \right), & \text{for } x < 0 \end{cases},$$

which is piecewise-linear for $\alpha = 1$. We focus on $\alpha \in [0,1]$, $\delta \in [0.30,1]$ and $\lambda \in [1,4]$.

Barberis (2012) shows that an agent with CPT preferences starts taking a fair and symmetric risk for up to five periods mainly for two wide parametric regions: low δ -high α and high δ -low α . We replicate the analysis from Barberis (2012) with the exponential value function defined above and obtain the optimal *Basic* risk-taking strategy for each of the 8,000 parameter triples in the set $\Delta = \{(\alpha, \delta, \lambda) : \alpha \in \{0, 0.053, \dots, 0.947, 1\}, \delta \in \{0.3, 0.337, \dots, 0.963, 1\}, \lambda \in \{1, 1.16, \dots, 3.84, 4\}\}$.

FIGURE 6.—Optimal risk-taking strategies under CPT



Notes: Figure 6 shows all optimal *Basic* risk-taking strategies under CPT. 8,000 CPT parameter combinations result in only 12 optimal strategies that start risk-taking. The strategies are ordered decreasingly according to the number N of parameter combinations for which the strategy is optimal, as shown in the header of each strategy. The bottom line of each strategy's header shows the median parameter triple over all parameter combinations that yield the strategy.

CPT predicts a low number of distinct risk-taking strategies. We find that a CPT agent starts risk-taking for 2,332 of the 8,000 parameter triples. Figure 6 shows all 12 unique optimal risk-taking strategies. We order the strategies by the number of parameter triples that yield the respective risk-taking strategy. Thus, among a population of CPT agents with different

parameters, we would expect to observe Strategy 1 more often than Strategy 2, and so on. Strategies 1 to 7 are loss-exit strategies, while strategies 8 to 12 are gain-exit strategies, which highlights that loss-exit risk-taking is more common under CPT than gain-exit risk-taking.

The median parameter triples for each risk-taking strategy—which are shown in the bottom line of each strategy’s header in Figure 6—demonstrate that we find loss-exit strategies mostly for a low δ -high α . Agents with low δ and high α like positively skewed risks (large gain with low probability), which are induced by loss-exit strategies (see also Ebert and Karehnke 2026). Intuitively, an agent with a low δ overweights the tails of the distribution, and high α implies high marginal utility of large gains. In contrast, the gain-exit strategies in Figure 6 are more prevalent for high δ -low α . An agent with high δ and low α likes negatively skewed risks (large losses with low probability), as induced by a gain-exit strategy, because she barely overweighs the low probability of a large loss, and the low α makes large losses only slightly more frightening than small losses.

In summary, CPT predicts (a) risk-taking for a wide range of parameters; (b) that loss-exit strategies are more common than gain-exit strategies; and (c) a small number of different risk-taking strategies.

Estimation approach. We next estimate the parameters of CPT using the elicited *Basic* risk-taking strategies.¹⁷ For each of the $j = 1, 2, \dots, 313$ subjects that completed a *Basic* Decoloring Task, we observe one chosen stopping time τ_c^j which induces an outcome distribution $S_{\tau_c^j}^j$. Observing stopping time τ_c^j implies that subject j did not choose any of the other available *Basic* stopping times that yield different outcome distributions. We generically denote non-chosen stopping times as $\tau_- \in T_{Basic}$, where T_{Basic} is the set of all 802 *Basic* stopping times, and assume that choosing τ_c^j is as if choosing τ_c^j in all pairwise binary choices between τ_c^j and τ_- . The approach mirrors structural estimations based on multiple price list data (e.g., Holt and Laury 2002).

Individual estimation. To estimate CPT parameters, we follow Stott (2006) and Bruhin, Fehr-Duda, and Epper (2010) and embed CPT into a stochastic (logit) choice function that gives the probability of choosing τ_c over alternative τ_- for a given parameter vector $\theta_j = (\alpha_j, \delta_j, \lambda_j, \epsilon_j)$ with error parameter $\epsilon_j \in (0, 100]$:

¹⁷ With *Path-dependent* and *Randomized* stopping times, the number of alternatives to the chosen stopping time is too large to be computationally feasible.

$$f(\tau_c | \tau_{\neg}, \theta_j) = \frac{1}{1 + e^{-\epsilon_j(v(s_{\tau_c} | \theta_j) - v(s_{\tau_{\neg}} | \theta_j))}}.$$

The error parameter ϵ_j measures the extent to which the subject's choice between two stopping times is determined by differences in the induced CPT values. Larger values of ϵ_j imply that choice is driven even by small differences in CPT values, while a low ϵ_j implies a more random choice. We estimate θ_j for each subject by maximizing the log-likelihood of observing τ_c^j in T_{Basic} , as


$$\theta_j^* = \arg \min_{\theta} \sum_{\tau_{\neg} \in T_{Basic} \setminus \{\tau_c^j\}} \log \left(1 + e^{-\epsilon \left(v(s_{\tau_c^j} | \theta) - v(s_{\tau_{\neg}} | \theta) \right)} \right).$$

Individual-level parameter estimates are clustered. Figure 7 shows the results of our maximum-likelihood estimation. To get a first impression, Panel (a) of Figure 7 shows the estimation results for two selected strategies. For the level-3 trailing stop-loss strategy in the upper row of Panel (a)—which is identical to the median *Basic* decoloring in Figure 3—we estimate $(\alpha, \delta, \lambda, \epsilon) = (1.00, 0.30, 1.00, 44.91)$. Our maximum-likelihood estimation thus finds a high α , low δ , and low λ for loss-exit strategies that induce positively skewed outcome distributions. In contrast, the estimates for a level-3 trailing take-profit strategy as in the bottom row of Panel (a) are $(\alpha, \delta, \lambda, \epsilon) = (0.28, 1.00, 1.00, 100.00)$.

Panel (b) of Figure 7 shows the location of all 313 individual-level parameter estimates. The most pronounced cluster of estimates is in the bottom-right of Panel (b), for a high α , low δ , and low λ , as is consistent with our observation that most *Basic* stopping times induce a positively skewed outcome distribution. We also find a pronounced cluster in the bottom center of the cube, that is, for a low α , low δ , and low λ . These parameter combinations are obtained for long-duration buy-and-hold strategies, of which we also observe many (see Section III.C). Otherwise, the parameter estimates are scattered over the cube, reflecting the heterogeneity in our data. Both the median and mean parameter triples over all individual estimates are in the high α , low δ , low λ region. The median tuple is $(1.00, 0.30, 1.00, 44.91)$ and the mean is $(0.69, 0.45, 1.47, 47.23)$. Note that we obtain the estimates from a single task (though with many choice options), such that our estimates are more extreme than in typical discrete choice estimations of CPT. Visually, our parameter estimates tend to be on the hull of the cube in Panel (b) of Figure 7. This pattern reflects limited trade-offs in the parameter estimation: the CPT value of a stop-loss strategy, for example, decreases in δ and increases in α , while the attractiveness of any risk-taking strategy is decreasing in λ . Adding further tasks (be it stopping

tasks with different basic gambles or horizons, or different risky choice tasks) would result in more nuanced estimates.

FIGURE 7.—Individual-level CPT parameter estimates



Notes: Panel (a) shows the parameter estimates that we obtain from our maximum-likelihood estimation for a level-3 trailing stop-loss (top) and a level-3 trailing take-profit strategy (bottom). Panel (b) shows the parameter triples that we estimate for all *Basic* risk-taking strategies. Both, the color and the size of the dots visualize the frequency of observing the respective parameter triple.


Finite mixture model. Given the theoretical predictions from Figure 6 and the apparent clustering of parameter combinations in Figure 7, we now analyze if a small number of CPT agent types can capture the heterogeneity in our data by estimating a finite mixture model as in Bruhin, Fehr-Duda, and Epper (2010). A finite mixture model assumes that the data is generated by a fixed number of agent types, each characterized by its own parameter vector. For each type, we estimate the parameters and the likelihood that an observed stopping time has been chosen by that type. Details and robustness analyses are in Appendix SA.4.

A small number of different CPT agent types captures the individual heterogeneity. Figure 8 shows the parameter estimates and median risk-taking strategies using four agent types. Similar to Bruhin, Fehr-Duda, and Epper (2010), we can unambiguously assign each risk-taking strategy to one agent type. About half of all risk-taking strategies are assigned to agent type 1, for whom the median risk-taking strategy is a threshold stop-loss.¹⁸ For agent type 1, we estimate a linear value function ($\alpha = 1.00$), strong probability weighting ($\delta =$

¹⁸ For agent type 1, the third-highest ball in the fourth column is black for 46% of the assigned risk-taking strategies, which is insignificantly different from 50% ($t(149) = -0.97, p = 0.33$). More correctly, the median risk-taking strategy for agent type 1 is thus in-between a trailing and threshold stop-loss.

0.30), and no loss aversion ($\lambda = 1.00$). The median risk-taking strategy of agent type 2, to whom about one-third of all risk-taking strategies are assigned, is a buy-and-hold strategy. The parameter estimates reflect little diminishing sensitivity ($\alpha = 0.86$), no probability weighting ($\delta = 1.00$), and no loss aversion ($\lambda = 1.00$). About one-third of the risk-taking strategies assigned to type 2 stop at some gain balls (unreported results), which leads to a slight negative average skewness of the induced outcome distributions ($\text{Skew}(S_t) = -0.08$, $t(107) = -3.688$) and to $\alpha < 1$.

FIGURE 8.—CPT parameter estimates and median strategies from a finite mixture model



Notes: Figure 8 shows the median risk-taking strategies for each agent type. Next to each header, we show the parameter estimates for the agent type as well as the number of strategies. We find an almost perfect separation of the risk-taking strategies.

Substantially fewer risk-taking strategies are assigned to agent types 3 and 4. Agent type 3 seems to be a buy-and-hold type with early stopping and is characterized by a linear value function ($\alpha = 1.00$), medium probability weighting ($\delta = 0.57$), and high loss aversion ($\lambda = 4$). The strategies assigned to agent type 4 are strategies that take the risk often, but stop at late loss balls, such that the parameter estimates are close to those for type 1.

Overall, CPT predicts a small number of different risk-taking strategies and a prevalence of stop-loss strategies over gain-exit strategies. In the previous sections, we noted that a small number of clusters of economically reasonable risk-taking strategies seems to provide a good approximation to the individual-level heterogeneity in our data, and that most risk-taking strategies are loss-exit strategies. In this section, we additionally highlight that the structural assumption of CPT groups risk-taking strategies into types representing trailing or threshold

stop-loss strategies as well as buy-and-hold strategies. Few subjects, instead, are assigned to types that yield no or little risk-taking, as well as take-profit strategies. CPT is flexible enough to be a good approximation to the non-parametric clusters identified using machine-learning.

IV. Applications

In this section, we use the Decoloring Task to obtain new insights on a number of economic applications. Specifically, we inform about subjects' mental models when planning repeated risk-taking (Section IV.A); we test if threshold strategies—which are often used in practice—provide a good approximation to unconstrained risk-taking strategies (Section IV.B); we show that mostly following a plan is compatible with a systematic dynamic inconsistency (Section IV.C); we introduce and quantify the mechanical dynamic inconsistency that emerges from constraining choice sets (Section IV.D); and we comprehensively test causes of dynamic inconsistency (Section IV.E). Some of these applications involve data from additional experimental tasks and studies (cf. Table 1), and we describe them when they are used.

IV.A How do people plan repeated risk-taking?


Allowing subjects to design unconstrained risk-taking strategies enables us to analyze the mental model by which subjects plan risk-taking. Subjects in our experiment can decolor the ball triangle in any order, such that we analyze the decoloring process as a representation of their mental model (Payne, Bettman, and Johnson 1988).¹⁹ In our setting—with a finite horizon of $T = 5$ risks—the normative benchmark is backward induction, where decoloring proceeds from right to left. Instead, we find that most subjects use forward induction—going from left to right.

In the aggregate, subjects decolor from left to right and from top to bottom. We recorded the decoloring process for $N = 240$ subjects who completed the *Basic* Decoloring Task in the second experimental session on May 23, 2024, and use click order numbers 1, 2, ..., 15 to summarize the decoloring process (a ball gets a one if it is decolored first). In the aggregate, subjects decolor the ball triangle from left to right (i.e., they use *forward induction*) and from top to bottom (i.e., *gains-first*), see Figure SA.1.3. Forward induction is evidenced by a positive average correlation between the risk repetition t and the click order ($\rho_t = 0.20, t(239) = 7.408$), whereas gains-first decoloring is reflected in a negative average

¹⁹ We thank Christian Gollier and Fabio Maccheroni for independently suggesting this analysis. Consistent with assuming that the decoloring process represents the planning process, we observe some pondering in that about 40% of subjects change the decoloring of at least one ball that they previously decolored.

correlation between the accumulated earnings at a ball and the click order ($\rho_g = -0.09, t(239) = -3.526$).

FIGURE 9.—Types of decoloring orders



Notes: Figure 9 shows the best-fitting decoloring process per cluster resulting from a K -means clustering with $K = 6$ clusters. We also show the correlation between the click order and, respectively, the risk repetition (ρ_t) and the accumulated earnings (ρ_g), as well as the number of decoloring processes assigned to each cluster.

Individual-level heterogeneity: few subjects use backward induction. We analyze the heterogeneity in the decoloring processes using a K -means clustering with $K = 6$ clusters.²⁰ Figure 9 shows the decoloring process that minimizes the total squared distance to the click orders of the cluster constituents (computed using the Hungarian algorithm), the number of subjects as well as ρ_t and ρ_g for each cluster.

Decoloring processes assigned to Cluster 1 (16% of subjects) are characterized by forward induction ($\rho_t > 0$ and large) and a gains-first tendency ($\rho_g > 0$). Clusters 2 and 3 contain 27% and 17% of subjects, respectively, and feature forward induction with losses-first. Cluster 4 (14% of subjects) features forward induction with no gains-first or loss-first tendency. Clusters 5 and 6 (12% and 15% of subjects, respectively) feature backward induction ($\rho_t < 0$). While Cluster 5 is gains-first ($\rho_g > 0$ and large), Cluster 6 is losses-first ($\rho_g < 0$).

²⁰ Multiple model selection criteria, including the silhouette score, AIC, and BIC point to $K = 6$. In unreported results, we confirm that the findings replicate economically when using $K = 5$ and $K = 7$ clusters.

Overall, the decoloring processes of 60% of subjects—those in clusters 1, 2, and 3—show a clear pattern of forward induction ($\rho_t > 0$ and large), while about 27% of subjects (clusters 5 and 6) seem to use backward induction. Our direct evidence of pervasive forward induction in designing stopping times complements research that infers dynamic decision-making processes indirectly from outcomes (see Johnson et al. 2002; Levitt, List, and Sadoff 2011; Moinas and Pouget 2013; Klein Teeselink et al. 2024).²¹

IV.B Are strategies threshold or trailing risk-taking strategies?

Section III.C established that many unconstrained risk-taking strategies are stop-loss strategies. Eliciting unconstrained risk-taking strategies allows us to further inquire what *kind* of stop-loss strategies subjects choose and whether threshold or trailing stop-loss strategies are more common. Threshold stop-loss orders appear to be offered by most brokers (see also Linnainmaa 2010), while trailing stop-loss strategies (sometimes called Azéma–Yor stopping times) appear to be less commonly observed despite their theoretical appeal.²²

We investigate the possible preference for trailing over threshold strategies in two steps. First, we compare unconstrained risk-taking strategies to the set of all threshold and trailing strategies, respectively, using matching and supervised machine-learning methods. Trailing strategies are more popular, with a ratio of trailing to threshold stop-loss strategies of approximately 1.5:1. Second, we introduce a *Threshold Decoloring Task* in which subjects must use threshold strategies to test whether this constraint systematically distorts stopping times. Despite the greater prevalence of trailing strategies among unconstrained strategies, restricting subjects to threshold strategies does not significantly affect aggregate properties of stopping times.


²¹ Much of the existing literature infers decision-making processes from observed outcomes. For example, Moinas and Pouget (2013) design an experiment in which asset price bubbles are ruled out by backward induction. The emergence of bubbles suggests that subjects do not apply backward induction. Our design complements the literature by directly observing dynamic decision-making processes.

²² For example, Vanguard seems not to offer trailing stop-loss orders (<https://investor.vanguard.com/investor-resources-education/online-trading/stock-order-types>; accessed 01/30/2026), Fidelity describes trailing stop-loss orders as an advanced order type (<https://www.fidelity.com/trading/faqs-order-types>; accessed 01/30/2026), while Charles Schwab (<https://www.schwab.com/learn/story/trailing-stop-orders-mastering-order-types>; accessed 01/30/2026) and Robinhood (<https://robinhood.com/us/en/support/articles/order-types/>; accessed 01/30/2026) seem to offer both order types. The broker used in Heimer et al. (2025) obliges traders to use threshold strategies when opening a position.

IV.B.1 Matching procedure

We first define the templates used for the matching procedure (see Figure SA.3.5). The threshold templates are all *Basic* (neither randomized nor path-dependent) threshold stop-loss, threshold take-profit and two-threshold strategies. Similarly, the trailing templates are all *Basic* trailing stop-loss, trailing take-profit and two-trailing strategies. We match each elicited, unconstrained risk-taking strategy to the closest threshold or trailing template strategy, respectively, using the lowest Hamming distance. In our setting, the Hamming distance is the number of ball changes (from black to white or vice versa), referred to as *redcolorings*, required to transform an unconstrained strategy into a given template strategy.

FIGURE 10.—Distribution of the distance to threshold and trailing risk-taking strategies



Notes: Figure 10 shows the distribution of the Hamming distance between the elicited risk-taking strategies and the closest threshold or trailing risk-taking strategy. We exclude risk-taking strategies that never take the risk or continue risk-taking at all balls. The left panel includes $N = 259$ *Basic* risk-taking strategies, and the right panel includes all risk-taking strategies ($N = 505$).


More risk-taking strategies resemble trailing than threshold strategies. Figure 10 shows the distribution of redecolorings. We exclude subjects that either always or never take the risk from the analysis, but all results are robust to this choice. On average 1.69 of 15 decolored balls must be redecolored to transform all *Basic* risk-taking strategies into threshold strategies, which is higher than the average number of redecolorings required to transform them into trailing strategies (1.44). We reject identity of the distributions of redecolorings using a Wilcoxon signed-rank test ($W = 4,245, n = 152$). The right panel of Figure 10 shows the same analysis when using all elicited risk-taking strategies: on average, more redecolorings are

required to transform the elicited risk-taking strategies into threshold strategies than into trailing strategies ($W = 13,550, n = 264$).

The ratio of trailing to threshold stop-loss strategies is about 1.5:1. Among the *Basic* risk-taking strategies in the left panel of Figure 10, we identify 61 pure trailing strategies and 44 pure threshold strategies, corresponding to a ratio of 1.39:1. To assess the robustness of this classification, we apply additional matching and supervised machine-learning methods—including cosine-similarity, logistic regression, a decision tree, and a random forest classifier. As reported in Appendix SA.3.2, these approaches yield similar results, with the ratio of trailing to threshold strategies ranging from 1.30:1 to 2.11:1 and centering around 1.5:1.

When free to choose, a majority of subjects prefers trailing to threshold strategies (see also Figure 5). The preference for trailing over threshold strategies emerges despite the relatively short risk-taking horizon with $T = 5$ risks, which limits the effective difference between trailing and threshold strategies.

FIGURE 11.—Elicitation of *Threshold* strategies



Notes: Subjects in a *Threshold* Decoloring Task must decolor all 15 balls of an initially green ball triangle. Subjects click a ball in the right half of the instruction panel (a) and give the risk-taking instruction for all balls in the selected ball's row in the left half of the instruction panel (a). All initially green balls in that row are decolored white (black) if the subject continues (stops) risk-taking for that row. Each subject proceeds like this in an order of her choice. Subjects can go back and change their decoloring. Panel (b) illustrates the median decoloring, whereby the median is taken over the effective decoloring at each ball.

IV.B.2 Threshold study

We next test whether risk-taking strategies differ substantially when subjects must choose among threshold strategies using a dedicated experimental study. Panel (a) of Figure 11 shows the decision screen of the new *Threshold* Decoloring Task ($N = 78$), in which subjects must decolor rows of balls and thus choose the same decoloring for all balls in each row. Risk-taking

strategies in *Threshold* can depend on the profit- and/or loss-level, but not on the risk-repetition.²³

Threshold strategies approximate unconstrained risk-taking well. Panel (b) of Figure 11 shows the median risk-taking strategy of subjects in the *Threshold* Decoloring Task. Compared to the median risk-taking strategy in the *Basic* Decoloring Tasks (top row of panel (b) in Figure 3), the only difference is that the bottom ball in the third column is black for *Threshold* but not for *Basic*. We indeed find that the share of subjects that stop risk-taking at the bottom ball in the third column is significantly larger in *Threshold* than in *Basic* ($t(389) = 2.22$). The median risk-taking strategy thus changes from a level-3 trailing stop-loss in *Basic* to a level-2 threshold stop-loss in *Threshold*.

Subjects in the *Threshold* Decoloring Task and in the unconstrained *Basic* Decoloring Task tend to choose similar types of risk-taking strategies. In *Threshold*, compared to *Basic*, fewer subjects use loss-exit (56% vs. 61%, see panel (b) of Figure 4) and gain-exit strategies (10% vs. 13%), while more subjects choose symmetric-exit strategies (33% vs. 26%). None of these differences is statistically significant. Similarly, we do not find statistically significant differences for summary statistics of the aggregate stopping times τ or the induced outcome distributions S_τ , except for a significantly more negative skewness of the stopping time τ in *Threshold* ($\text{Skew}(\tau) = -0.80$) than in *Basic* ($\text{Skew}(\tau) = -0.32$, $t(389) = 3.78$).

In sum, a majority of subjects prefers trailing over threshold strategies. Threshold strategies, however, approximate unconstrained strategies well, at least in our five-period setting. The focus on threshold strategies may thus constitute a reasonable simplification for many (experimental and theoretical) studies of stopping under precommitment with short horizons.

IV.C Measuring dynamic consistency for unconstrained stopping times


After submitting a precommitted risk-taking strategy in the Decoloring Task, subjects take the risk sequentially during ten Walking Tasks. Our dataset allows for testing dynamic consistency at each possible ball, because we have ball-by-ball data on risk-taking plans (from the Decoloring Task) as well as on risk-taking actions (from the Walking Tasks).

Walking Tasks. Figure 12 shows the decision screen for a *Basic* Walking Task. Subjects make sequential risk-taking decisions while visually “walking” through the ball triangle (left panel of Figure 12). Subjects instruct the computer about their risk-taking action at each ball

²³ On average, subjects need 136 seconds to complete the *Basic* Decoloring Task and 110 seconds to complete the *Threshold* Decoloring Task ($t(389) = 2.463$), suggesting that threshold strategies are simpler to implement.

that is reached, possibly deviating from their previously elicited plan (right panel of Figure 12). If a risk is taken, it is directly resolved. Each Walking Task comprises taking up to five risks and each subject completes ten Walking Tasks. In the *Randomized* and *Randomized^C* Walking Task, subjects can (costly) stop risk-taking with any probability on the grid 0, 10, ..., 90, 100.²⁴

FIGURE 12.—Basic Walking Task



Notes: After submitting the decoloring in a *Basic Decoloring Task*, subjects in a *Basic Walking Task* walk through their decolored ball triangle. The left panel of Figure 12 shows the situation where the risk has been won the first two times and lost the third time. The subject must decide whether to take the risk a fourth time in the panel on the right. Subjects' default is the plan-consistent risk-taking action. If a subject continues risk-taking, the single risk is resolved, the path through the ball triangle is updated, and the subject sees the decision screen again. If a subject stops at a ball, this ball is colored yellow to indicate that this ball determines the subject's payoff if this Walking Task determines payment.

Comparing risk-taking actions from the Walking Tasks with risk-taking plans from the Decoloring Task allows us to test for dynamic inconsistency within-subject at each ball. Subjects are dynamically (in-)consistent at a ball if they (do not) follow their risk-taking plan at that ball. We next report two results on dynamic inconsistency for subjects in the *Basic* studies ($N = 147$) as a benchmark.

Systematic dynamic inconsistency: novel evidence on the standard shift. We first document that the *standard shift in risk-taking*—going from a loss-exit plan to gain-exit risk-taking—is present in our data. The standard shift was theoretically predicted by Barberis (2012) and empirically first confirmed by Heimer et al. (2025). We show that the standard shift prevails with unconstrained planning. To test for the standard shift in unconstrained stopping

²⁴ Subjects must always resolve five risks regardless of when they stopped risk-taking to prevent finishing a task faster by stopping early. After stopping, subjects no longer observe the risk resolution to prevent regret. We do not implement a dedicated path-dependent Walking Task since sequential risk-taking is necessarily path-dependent (see Footnote 6). Pooling all subjects, we observe a comprehensive set of 37,452 risk-taking actions. For any ball, we observe at least 268 actions, with a maximum of 8,810 actions ($= 881 \text{ subjects} * 10 \text{ Walking Tasks}$) for the very first ball in the ball triangle.

time data, we propose to measure the *change in stop-loss intensity* ΔD_i , as follows. With $\circ = A, P$ for actual and planned risk-taking, respectively, we define

$$\text{Proportion of Losses Stopped } (PLS_i^\circ) = \frac{\text{Losses stopped}_i^\circ}{\text{Losses reached}_i^\circ},$$

$$\text{Proportion of Gains Stopped } (PGS_i^\circ) = \frac{\text{Gains stopped}_i^\circ}{\text{Gains reached}_i^\circ},$$

and

$$D_i^\circ = PLS_i^\circ - PGS_i^\circ.$$

$D_i^\circ < 0$ indicates gain-exit risk-taking, while $D_i^\circ > 0$ indicates loss-exit risk-taking. $\Delta D_i = D_i^A - D_i^P$ quantifies the change in stop-loss intensity. The standard shift posits that $\Delta D_i < 0$, that is, actual risk-taking is more gain-exit/less loss-exit than planned. Note that D_i^A coincides with Odean's (1998) disposition measure.²⁵ We compute D_i^P using the realized paths the subject faced in the Walking Tasks (in which subjects must always resolve five risks) but stop risk-taking following the subjects' risk-taking strategy.

We find that both planned ($D_i^P = 0.207, t(140) = 7.366$) and actual risk-taking ($D_i^A = 0.081, t(140) = 3.658$) are of a loss-exit type. The positive average D_i^P resonates with our previous analysis of the Decoloring Task: Subjects plan to use loss-exit strategies. Removing subjects' precommitment—as in the Walking Tasks—leads to the standard shift in that $\Delta D_i = -0.125$ ($t(140) = -4.542$), meaning less loss-exit risk-taking than planned. Section IV.E shows why D_i^A is not negative.

Pronounced overall dynamic consistency. Before we proceed, we analyze another, simpler measure of dynamic (in-)consistency that has not yet received attention because it requires data on ball-by-ball risk-taking plans and actions. The (non-directed) *consistency measure* C_i is the share of plan-consistent risk-taking actions for each subject i :

$$C_i = \frac{\text{Plan executions}_i}{\text{Deviation opportunities}_i} = 1 - \frac{\text{Deviations}_i}{\text{Deviation opportunities}_i}.$$

C_i equals 1 if subject i always follows her plan, and 0 if she never follows her plan. Rescaling the number of plan executions by the number of deviation opportunities controls for between-subject differences, such as general risk appetite or the specific paths observed.

²⁵ As in Odean (1998), $D_i^A < 0$ signifies a disposition effect (or gain-exit risk-taking).

On average, subjects in a *Basic* study follow their risk-taking plan 89.2% of the time (see Table SA.2.2).²⁶ We find this high overall consistency worth reporting since the literature mostly emphasizes inconsistencies in stopping, like the standard shift discussed above. The reason why a high overall consistency is not at odds with the likewise strong evidence for a systematic inconsistency is that C_i considers risk-taking actions at all balls—including those at which subjects plan to continue. A subject planning with a low stop-loss level but sequentially taking the risk until reaching a high gain will be consistent at most balls—all balls lying between the loss and gain levels.²⁷ A behavioral shift from a loss-exit plan to gain-exit behavior is therefore compatible with a high C_i . The evidence of high C_i also implies that the standard shift is *the* dynamic inconsistency observed in stopping. Shifting from buy-and-hold to never starting, for example, is inconsistent with a high C_i . In sum, our subjects exhibit the standard shift but are otherwise quite consistent in their stopping behavior.

IV.D Mechanical inconsistency

Subjects—and economic agents in general—forced to choose a plan from a constrained set may deviate from their plan simply because the desired plan was not available. Such *mechanical inconsistency* is not a preference reversal and a source of bias that has not been acknowledged in the literature thus far. Ex-ante, mechanical inconsistency could either increase or decrease the standard shift. In this section, we measure the mechanical inconsistency induced by a restriction to threshold strategies and estimate its effect on the standard shift.

We first find the set of counterfactual threshold strategies (CTSs) that best approximate the elicited unconstrained risk-taking strategy for each subject in *Basic* using the same method as in Section IV.B.1. Panel (a) of Figure 13 illustrates the CTSs for an exemplary trailing strategy, for which two threshold strategies have a Hamming distance of one. Redecoloring TIB1 (Threshold-Identifying Ball 1) black yields a level-2 threshold strategy (CTS1) whereas redcoloring TIB2 white yields a level-3 threshold strategy (CTS2).²⁸ We assume that, if


²⁶ We find no systematic pattern in the development of C_i over the ten Walking Tasks. The within-subject correlation between the Walking Task number and C_i is virtually zero ($\rho = -0.008$, $W = 2,624$) and a Friedman test does not reject equality of the task-specific C_i ($\chi^2(9, N = 147) = 6.114$).

²⁷ To further assess the level of C_i , note that a randomly choosing subject would exhibit plan-consistent risk-taking in half of her decisions, yielding a C_i of 0.50. Planning to use a level-3 trailing stop-loss strategy but taking the risk following a level-3 (level-2) trailing take-profit strategy would even yield a C_i of 0.84 (0.65).

²⁸ If two redcolorings are required to make the elicited strategy a threshold strategy, we proceed analogously. Then, pairs of TIBs define the CTSs. We define TIBs and CTSs analogously when three, four, five, six

restricted to threshold strategies, a subject would have selected one of the CTSs. We then study the (in)consistency of subjects' risk-taking actions by treating these CTSs as their plans.

FIGURE 13.—Analysis of mechanical dynamic inconsistency



Notes: Panel (a) of Figure 13 illustrates the construction of counterfactual threshold strategies (CTSs) for an exemplary trailing strategy. TIB1 and TIB2 are the Threshold-Identifying Balls that must be redecolored to convert the elicited strategy into a threshold strategy; each yielding a different CTS. Panel (b) shows the distribution of the consistency measure C_i for the elicited *Basic* decolorings, the CTSs, and at the TIBs for subjects in the *Basic* studies ($N = 147$). Panel (c) shows the change in the stop-loss intensity $|\Delta D_i|$ under the CTSs and under the elicited *Basic* decolorings.

Panel (b) of Figure 13 shows that subjects' risk-taking actions are more consistent with their elicited strategy (average $C_i = 0.892$) than with the CTSs (average $C_i = 0.847$, $t(120) = 3.999$). Consistency is particularly low at the TIBs themselves (average $C_i = 0.374$), indicating that the TIBs are precisely the balls at which subjects do not want their plan altered and follow their unconstrained risk-taking plan instead.

Panel (c) of Figure 13 examines the bias in the estimated standard shift induced by mechanical inconsistency. In Section IV.C, we found $D_i^P = 0.21$ and $D_i^A = 0.08$ for subjects in the *Basic* studies. A counterfactual restriction of the risk-taking plan to threshold strategies significantly increases D_i^P to 0.279 ($t(119) = 2.156$). The estimated change in the stop-loss intensity under a restriction to threshold strategies is thus $\Delta D_i = -0.197$ instead of $\Delta D_i = -0.125$ ($t(118) = 2.028$) under unconstrained planning. Put differently, the mechanical dynamic inconsistency biases the magnitude of the standard shift upwards by 58%.²⁹

redecorings are required. We compute the consistency measures as the average over the subjects' CTSs if there is more than one best-approximating CTS for a subject's strategy—as in Panel (a) of Figure 13.

²⁹ The estimate based on the CTSs closely matches the evidence from the *Threshold* study, in which subjects are exogenously restricted to threshold strategies. In *Threshold*, we find $D_i^P = 0.250$ and $D_i^A = 0.063$ (see Table

Overall, we find substantial evidence of mechanical inconsistency. First, subjects' risk-taking actions are, on average, more consistent with their unconstrained plans than with counterfactual thresholds. Second, subjects' risk-taking actions are especially inconsistent with the counterfactual threshold at the balls that need to be redecolored to make their unconstrained strategy a threshold strategy. Thus, although threshold strategies approximate the aggregate distribution of planned risk-taking well (Section IV.B.2), they differ from unconstrained plans exactly at the balls that are consequential for the standard shift.

Imposing threshold strategies overstates the standard shift by about half the effect size, showing that mechanical inconsistency can be severe. To our knowledge, we are the first to identify and quantify mechanical inconsistency separately from preferences. The distinction is likely to matter more broadly whenever choices are compared across environments that differ in their choice set or decision structure. The concern may be particularly important under context-dependent preferences (Bell 1982; Loomes and Sugden 1982; Bordalo, Gennaioli, and Shleifer 2012). For the standard shift, preferences account for about half of the effect, indicating that preference-driven dynamic inconsistency is itself substantial and economically meaningful.


IV.E The causes of dynamic inconsistency: a decomposition of the standard shift

We now analyze causes of dynamic inconsistency and of the standard shift in particular by implementing different types of the Walking Task between-subject (see Figure 14 and Table 1). During a *Basic* Walking Task, subjects walk along their previously decolored ball triangle, which serves as a plan reminder, and the plan-consistent risk-taking action is pre-selected as default. First, the presence of a default risk-taking action may affect consistency. With a default, plan-deviations require active undoing of one's earlier decision, making inconsistencies salient. We implement a *No default* Walking Task that is like the *Basic* Walking Task, except that no risk-taking action is pre-selected as default. Second, the availability of a plan reminder may influence dynamic consistency. Subjects in a *No reminder* Walking Task see no plan reminder (and have no default) when taking the risk sequentially but walk along a colored (green) ball triangle. Third, planning as such may affect consistency, which is why dynamic consistency is sometimes assessed between-subject. We test the effect of planning on sequential risk-taking by having subjects in *No planning* complete Walking Tasks—which are

SA.2.3). The change in the stop-loss intensity of $\Delta D_t = -0.187$ is marginally significantly larger in magnitude than for subjects in *Basic* ($t(212) = 1.301$).

identical to the *No reminder* Walking Tasks—without a previous Decoloring Task. Comparing our measures of dynamic inconsistency across the different treatments causally identifies the effect of a default, a memory-aid, and planning as such on dynamic consistency.

FIGURE 14.—Dynamic consistency treatments




Notes: After submitting the decoloring in a *Basic* Decoloring Task (left panel), subjects in a *Basic* and *No Default* Walking Task (right panel) walk through their decolored ball triangle. Subjects in a *No reminder* Walking Task walk through a green ball triangle. We plot exemplarily the situation where the risk has been won the first two times and lost the third time. Subjects must decide whether to take the risk a fourth time. In the *Basic* Walking Task, subjects’ default is the plan-consistent risk-taking action. Subjects in a *No default* or *No reminder* Walking Task have no default risk-taking action and the “Confirm Instruction” button only becomes active after selecting “Stop” or “Continue.” Subjects in the *No planning* study do not submit a risk-taking plan (i.e., they start with the Walking Task), and complete the *No reminder* Walking Task.

Reminding and planning increases consistency. Across the different Walking Tasks, we find that subjects mostly follow their previously elicited risk-taking plan when taking the risk sequentially. Panel (a) of Figure 15 shows the distribution of the consistency measure C_i for the different experimental studies.³⁰ On average, 85.5% (median 88.4%) of the risk-taking actions are consistent with the risk-taking plan, varying from 82.6% (87.5%) for subjects in *No reminder* to 89.2% (92.3%) for subjects in *Basic* (see Table SA.2.2).

³⁰ We construct C_i for each subject in the *No planning* study by drawing a random sample of 50 *Basic* risk-taking plans from other subjects. We then compute C_i for each sampled plan and take the average.

FIGURE 15.—Average dynamic consistency measures




Notes: Panel (a) of Figure 15 shows the distribution of the consistency measure C_i for the different studies. Panel (b) of Figure 15 plots the average D_i^P (light bars) and D_i^A (dark bars).

Comparing the consistency C_i across studies, we find statistically significant differences between the *Basic* study and *No reminder* ($t(226) = 3.101$) and *No Planning* ($t(308) = 4.308$), respectively. Subjects are thus significantly less consistent if not reminded of their plan or if they do not explicitly plan to begin with. Plan consistency does not vary significantly when removing the default or when restricting the risk-taking strategies to threshold strategies.

Causes of the standard shift. Panel (b) of Figure 15 shows the average stop-loss intensity D_i° for the different types of Decoloring and Walking Tasks. We focus on D_i^A . First, removing the default risk-taking action in the *No Default* treatment does not affect the stop-loss intensity D_i^A of risk-taking actions compared to *Basic* ($t(224) = 0.011$). This result also suggests that subjects make their decisions intentionally, since the heuristic to “always follow the default” would save clicks and cursor moves. Second, sequential risk-taking is systematically affected by a reminder. The average D_i^A for *No Reminder* is -0.03 and significantly smaller than for *Basic* ($t(221) = 3.294$). Without a reminder, we obtain a (mildly and insignificantly) negative average D_i^A , that is, a shift from loss-exit plans to gain-exit risk-taking. Third, comparing D_i^A between *No Reminder* and *No Planning* assesses the additional effect of planning as such on sequential risk-taking. We find a statistically marginally significant difference of 0.05 between D_i^A in *No Reminder* and *No Planning* ($t(238) = 1.721$).³¹

³¹ We also analyze changes in the skewness of the outcome distribution (Table SA.2.3). The results are similar to the disposition measure D_i° , but noisier because the skewness of each subject’s outcome distribution in the Walking Tasks is constructed from only ten outcomes. Nonetheless, the correlation between D_i° and $\text{Skew}(S_\tau)$ is 0.92 for the risk-taking plan, and 0.52 for sequential risk-taking. Comparing the average skewness of the *Basic*

FIGURE 16.—Behavioral decomposition of the standard shift



Notes: Figure 16 shows the decomposition of the standard shift from threshold strategies ($D_i^P = 0.25$ in study *Threshold*) to purely sequential risk-taking ($D_i^A = -0.08$ in study *No planning*).

A behavioral decomposition of the standard shift. We observe a reversal from loss-exit plans to gain-exit risk-taking when comparing precommitted risk-taking with *Threshold* plans ($D_i^P = 0.250 > 0$, $t(72) = 5.163$) to purely sequential risk-taking without planning ($D_i^A = -0.077 < 0$, $t(159) = -4.689$). Our experimental studies allow us to decompose the total shift of $|\Delta D_i| = 0.327$ into its behavioral components, as shown in Figure 16. The largest share, 38% (0.125), comes from the difference in the stop-loss intensity of *Basic* risk-taking plans and actions and is thus caused by removing precommitment. The restriction to threshold strategies—a mechanical dynamic inconsistency—accounts for 13% (0.043) of the total effect. Nudging subjects with a plan-consistent default has no measurable impact. Memory constraints, as captured by removing the plan reminder, explain another 34% (0.112) of the difference. Finally, planning as such accounts for the remaining 15% (0.048) of the total shift.

V. Discussion: Eliciting stopping times with the Decoloring Task

The Decoloring Task constitutes the first method for eliciting unconstrained stopping times, and we now discuss features of the experimental design (Section V.A), its robustness (Section V.B) and make further observations (Sections V.C and V.D).

risk-taking plan to the realized skewness in *No planning*—as is also done in Heimer et al. (2025)—we find that risk-taking plans induce a positively skewed outcome distribution, whereas the realized outcome distribution is negatively skewed (see Table SA.2.3).

V.A Experimental design

We deliberately isolate dynamic aspects of risk-taking by using the simplest possible risk of winning or losing a given amount with equal probability. Just like the large literature on static choice under risk abstracts from dynamic aspects, we abstract from static aspects. One could, however, easily adapt the Decoloring Task to study non-fair risks (Antler and Arad 2023), skewed risks (Ebert and Köster 2026), a multiplicative random walk (Strack and Viefers 2021), or non-stationary risks with time-changing moments (Crosetto and Filippin 2013; Nielsen 2019). Alternatively, one could extend our framework to the elicitation of contingent plans for stochastic control by asking for the wealth shares to be invested at each ball of the ball triangle (Merton and Samuelson 1992) or to consumer search problems (Weitzman 1979).

The incentive-compatibility of the Decoloring Task roots in the strategy method (Selten 1967), and our implementation may offer advantages in the context of repeated risk-taking. Asking a subject what she would do in each possible contingency (“what would you do if you win twice and lose once?”) seems more abstract. The order in which contingencies are presented may prime subjects and affect risk-taking plans due to, for example, failures of contingent thinking (Esponda and Vespa 2024), or it may be unnatural (for example, if contingencies are shown in a random order). Eliciting subjects’ plans in a pre-specified order feels closer to sequential risk-taking in that one asks about one contingency after the other. We believe that it is crucial that subjects choose the order themselves when designing their strategy while viewing each action as part of their complete risk-taking strategy, as is the case in the Decoloring Task. This feature also allows for insight into the mental model by which subjects plan risk-taking (see Section IV.A).

V.B Out-of-sample replication of the Basic study

Study *Basic II* provides an out-of-sample replication of *Basic I*, because both studies are identical except for taking part at different dates (see Table 1). All major results reported replicate. Specifically, we find no statistically significant difference between the expected stopping times (3.74 for *Basic I* vs. 3.85 for *Basic II*, $t(145) = 0.598$) or the skewness of the outcome distribution (0.36 vs. 0.37, $t(145) = 0.157$). The median risk-taking strategies are identical, with 58% of subjects using loss-exit strategies in *Basic I* (62% in *Basic II*), 12% gain-exit (9%), and 30% (28%) symmetric-exit strategies. Conducting the K -means clustering from Section III.C using only the data from *Basic I* and *Basic II*, respectively, yields highly similar clusters with a relative majority of risk-taking strategies assigned to clusters interpreted

as (trailing) stop-loss (49% and 55%) followed by buy-and-hold (23% and 26%). The only behavioral difference between the two studies appears to be a stronger preference for trailing relative to threshold stop-loss strategies in *Basic I*. In *Basic I*, we observe 18 pure trailing and 5 pure threshold stop-loss strategies (as defined in Section IV.B.1), compared with 10 and 6, respectively, in *Basic II*. We also find no statistically significant differences between the consistency measure C_i (0.890 vs. 0.894, $t(145) = 0.182$), the loss-exit intensity of risk-taking plans D_i^P (0.19 vs. 0.23, $t(140) = 0.663$), and the loss-exit intensity of risk-taking actions D_i^A (0.07 vs. 0.10, $t(141) = 0.710$) between *Basic I* and *Basic II*.

V.C Buy-and-hold

We find it worth highlighting that, in an experiment that allows for unconstrained stopping strategies, 28% of the subjects choose simple buy-and-hold strategies (see Figure 5). The prevalence of buy-and-hold strategies suggests that many individuals prefer simple rules even when more complex strategies are available.

V.D Theory considerations

The elicited stopping times are broadly consistent with predictions from non-EU models, particularly CPT. At the same time, the observed strategies differ from theoretical benchmarks in important ways. Theoretical models often predict level-1 stopping frontiers (see Figure 6), whereas the observed strategies in our experiment accommodate losing the single risk more often. Understanding which models of choice under risk can generate the stopping rules observed in the data remains an open question.

Similarly, CPT predicts that path-dependent strategies are used to induce more positively skewed outcome distributions S_τ than time- and outcome-contingent strategies (He et al. 2017; Hu, Obłój, and Zhou 2022). We find, however, no significant difference between the skewness of S_τ in the *Basic* and *Path-dependent* Decoloring Task. Moreover, subjects use path-dependence neither to (a) systematically stop at earnings peaks along a path as predicted by regret preferences (Strack and Viefers 2021), nor to (b) systematically differentiate between upward- and downward-sloping paths (Rabin 2002; Grosshans and Zeisberger 2018). Relatedly, Magnani et al. (2022) find that path-dependent stochastic control leads to less efficient portfolio choices than standard path-independent control.

Our evidence on the usage of path-dependence is more consistent with a revealed demand for randomization. Path-dependence can be used as a randomization device because subjects design their strategy before the path is realized, making the stopping outcome at a path-

dependent ball ex-ante random (see He et al. 2017; Hu, Obłój, and Zhou 2022). For example, a subject who stops risk-taking if the path to the middle ball in the third column of the ball triangle was down-up and continues if it was up-down induces an ex-ante stopping probability of 50% at that ball. This motivation for path-dependent stopping times is similar to the motivation for randomizing static decisions identified by Agranov, Healy, and Nielsen (2023). In their experiment, subjects vary static choices over multiple repetitions of the same choice situation. When doing so, subjects do not know which of the choices will be payoff relevant, such that varying the static choice is equivalent to randomizing over the options.

Consistent with our interpretation of path-dependence as a way to achieve randomization, we observe a higher demand for direct randomization than for path-dependence. The high usage of randomized risk-taking strategies is in line with observations in static settings (Dwenger, Kübler, and Weizsäcker 2018; Agranov and Ortoleva 2022),³² and provides empirical justification for randomized stopping times in theoretical work beyond their role as a convexification device (Henderson, Hobson, and Tse 2017; Hu, Obłój, and Zhou 2022). Although randomization is often used merely as a technical tool in the analysis of optimal stopping, we find that allowing for randomization systematically affects risk-taking strategies. First, because subjects randomize at early risk repetitions, they place a substantial probability mass on stopping early. This leads to a lower expected stopping time in *Randomization* studies. Second, subjects' stopping probabilities at gain nodes are close to the stopping probabilities at loss nodes, resulting in relatively more gain-exit strategies. Randomization is not systematically used to generate more skewness (Henderson, Hobson, and Tse 2017). Third, subjects typically randomize with stopping probabilities below 50%, which is markedly different from a 'pull to the center' effect. We currently have no explanation for this finding.

CONCLUSION

We propose and implement the first experimental method to elicit unconstrained stopping times, which are important objects in economic models and probability theory more broadly. Subjects choose their desired stopping time freely in some studies and under theoretically motivated constraints in others. We find that many subjects want the trailing version of stop-loss strategies, buy-and-hold strategies, path-dependence, as well as randomization. Foremost,

³² The prevalence of randomization in our *Randomization* studies falls within the range of randomization frequencies observed in other studies. Agranov and Ortoleva (2017) find that 71% of subjects randomize, Feldman and Rehbeck (2022) report that 94% of subjects randomize, and Agranov, Healy, and Nielsen (2023) observe that between 64% and 76% of subjects mix in at least one domain.

we document a significant and robust individual-level heterogeneity in stopping times, as well as in how subjects obtain them (through forward and backward induction, respectively). The heterogeneity can be approximated by a small number of distinct strategy types that have been considered in earlier theoretical and empirical research, and that are observed in the real world.

As an application, we show that simple threshold stop-loss strategies approximate aggregate risk-taking plans well, lending support to the many economic models that emphasize or predict these strategies. We quantify the extent to which dynamic inconsistency can be mitigated by offering memory-aides, commitment devices, and the opportunity of explicit planning. A conceptual contribution concerns the definition and measurement of *mechanical inconsistency*, a design-driven upward bias in dynamic inconsistency estimates that is difficult to avoid in many empirical studies. We hope that these first results on general/unconstrained stopping times can guide the assumptions imposed in theoretical work and inform future experiments on stopping times and dynamic consistency measurements more generally.

Acknowledgements. We thank Nicholas Barberis, Lucas Böttcher, Thomas Epper, Katrin Gödker, Christian Gollier, Andreas Grunewald, Alex Imas, Peiran Jiao, Pascal Kieren, Jacopo Magnani, George Malikov, Fabio Maccheroni, Sven Rady, Philipp Strack, Simeon Schudy, Frederik Schwerter, Stefan Trautman, Christoph Vanberg, Jiakun Zheng, Xunyu Zhou, and seminar participants at Aalto, Frankfurt School, Ghent, HEC Montréal, Heidelberg, Maastricht, Oslo (BI), St. Gallen, Trondheim, and Yale, and conference participants at the Experimental Finance Conference 2023, the European Decision Science Day 2023, the Annual Meeting of the European Group of Risk and Insurance Economists 2023, the Canadian Junior Finance conference 2024, Cirano Workshop on Experimental and Behavioural Research 2024, the 40th meeting of the European Economic Association, and the Research in Behavioural Finance Conference 2024 for helpful comments and suggestions. We thank Felix Schmücker and Alireza Shadmand for providing excellent research assistance.

Funding. This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 803380) and from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – project number 462020252.

Disclosure of Potential Conflict of Interests. Sebastian Ebert and Maximilian Voigt have no conflicts of interest to declare. All views expressed in the paper are those of the authors, who assume complete responsibility for its contents.

REFERENCES

- Agranov, Marina, Paul J. Healy, and Kirby Nielsen, “Stable randomization,” *Economic Journal*, 133 (2023), 2553–2579.
- Agranov, Marina, and Pietro Ortoleva, “Stochastic choice and preferences for randomization,” *Journal of Political Economy*, 125 (2017), 40–68.
- , “Revealed Preferences for Randomization: An Overview,” *American Economic Review (Papers & Proceedings)*, 112 (2022), 426–430.
- Andrade, Eduardo B., and Ganesh Iyer, “Planned versus actual betting in sequential gambles,” *Journal of Marketing Research*, 46 (2009), 372–383.
- Antler, Yair, and Ayala Arad, “An Experimental Analysis of the Prize-Probability Tradeoff in Stopping Problems,” *Review of Economics and Statistics*, (2023), 1–45.
- Barberis, Nicholas, “A model of casino gambling,” *Management Science*, 58 (2012), 35–51.
- Barkan, Rachel, and Jerome R. Busemeyer, “Changing plans: dynamic inconsistency and the effect of experience on the reference point,” *Psychonomic Bulletin & Review*, 6 (1999), 547–554.
- Bell, David E., “Regret in decision making under uncertainty,” *Operations Research*, 30 (1982), 961–981.
- Björk, Tomas, Mariana Khapko, and Agatha Murgoci, “Time-Inconsistent Control Theory with Finance Applications,” (Springer, 2021).
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, “Salience Theory of Choice Under Risk,” *Quarterly Journal of Economics*, 127 (2012), 1243–1285.
- Bruhin, Adrian, Helga Fehr-Duda, and Thomas Epper, “Risk and rationality: Uncovering heterogeneity in probability distortion,” *Econometrica*, 78 (2010), 1375–1412.
- Chen, Daniel L., Martin Schonger, and Chris Wickens, “oTree—An open-source platform for laboratory, online, and field experiments,” *Journal of Behavioral and Experimental Finance*, 9 (2016), 88–97.
- Crosetto, Paolo, and Antonio Filippin, “The ‘bomb’ risk elicitation task,” *Journal of Risk and Uncertainty*, 47 (2013), 31–65.
- Cubitt, Robin, Chris Starmer, and Robert Sugden, “Dynamic choice and the common ratio effect: An experimental investigation,” *Economic Journal*, 108 (1998), 1362–1380.
- Dertwinkel-Kalt, Markus, and Jonas Frey, “Optimal Stopping in a Dynamic Salience Model,” *International Economic Review*, 65 (2024), 885–913.
- Dixit, Avinash K., “The art of smooth pasting,” (Taylor & Francis, 1993).
- Dixit, Avinash K., and Robert S. Pindyck, “Investment under uncertainty,” (Princeton University Press, 1994).
- Dwenger, Nadja, Dorothea Kübler, and Georg Weizsäcker, “Flipping a coin: Evidence from university applications,” *Journal of Public Economics*, 167 (2018), 240–250.
- Ebert, Sebastian, and Paul Karehnke, “Skewness Preferences in Choice under Risk,” *Working paper available at SSRN*, (2026).
- Ebert, Sebastian, and Mats Köster, “On Taking a Skewed Risk More Than Once,” *American Economic Journal: Microeconomics*, 18 (2026), 395–425.
- Ebert, Sebastian, and Philipp Strack, “Until the bitter end: on prospect theory in a dynamic context,” *American Economic Review*, 105 (2015), 1618–33.

- Esponda, Ignacio, and Emanuel Vespa, “Contingent Thinking and the Sure-Thing Principle: Revisiting Classic Anomalies in the Laboratory,” *Review of Economic Studies*, 91 (2024), 2806–2831.
- Feldman, Paul, and John Rehbeck, “Revealing a preference for mixtures: An experimental study of risk,” *Quantitative Economics*, 13 (2022), 761–786.
- Fischbacher, Urs, Gerson Hoffmann, and Simeon Schudy, “The causal effect of stop-loss and take-gain orders on the disposition effect,” *Review of Financial Studies*, 30 (2017), 2110–2129.
- Fudenberg, Drew, Jon Kleinberg, Annie Liang, and Sendhil Mullainathan, “Measuring the Completeness of Economic Models,” *Journal of Political Economy*, 130 (2022), 956–990.
- Fudenberg, Drew, and Annie Liang, “Machine learning for evaluating and improving theories,” *SIGecom Exch.*, 18 (2020), 4–11.
- Gneezy, Uri, and Jan Potters, “An experiment on risk taking and evaluation periods,” *Quarterly Journal of Economics*, 112 (1997), 631–645.
- Gödker, Katrin, Peiran Jiao, and Paul Smeets, “Investor memory,” *Review of Financial Studies*, 38 (2025), 1595–1640.
- Gödker, Katrin, Terrance Odean, and Paul Smeets, “Disposed to be overconfident,” *Working paper available at SSRN*, (2026).
- Grosshans, Daniel, and Stefan Zeisberger, “All’s well that ends well? On the importance of how returns are achieved,” *Journal of Banking & Finance*, 87 (2018), 397–410.
- He, Xindi, and Yucheng Liang, “Risk Revisited,” *Working paper available at SSRN*, (2024).
- He, Xue Dong, Sang Hu, Jan Obłój, and Xun Yu Zhou, “Path-dependent and randomized strategies in Barberis’ casino gambling model,” *Operations Research*, 65 (2017), 97–103.
- He, Xue Dong, and Xun Yu Zhou, “Who are I: Time inconsistency and intrapersonal conflict and reconciliation,” in *Stochastic Analysis, Filtering, and Stochastic Optimization* (Springer, 2022).
- Heidhues, Paul, and Philipp Strack, “Identifying Present Bias from the Timing of Choices,” *American Economic Review*, 111 (2021), 2594–2622.
- Heimer, Rawley, Zwetelina Iliewa, Alex Imas, and Martin Weber, “Dynamic Inconsistency in Risky Choice: Evidence from the Lab and Field,” *American Economic Review*, 115 (2025), 330–363.
- Henderson, Vicky, David Hobson, and Alex Tse, “Randomized strategies and prospect theory in a dynamic context,” *Journal of Economic Theory*, 168 (2017), 287–300.
- , “Probability weighting, stop-loss and the disposition effect,” *Journal of Economic Theory*, 178 (2018), 360–397.
- Holt, Charles A., and Susan K. Laury, “Risk aversion and incentive effects,” *American Economic Review*, 92 (2002), 1644–1655.
- Hu, Sang, Jan Obłój, and Xun Yu Zhou, “A Casino Gambling Model Under Cumulative Prospect Theory: Analysis and Algorithm,” *Management Science*, 69 (2022), 2474–2496.
- Imas, Alex, “The realization effect: Risk-taking after realized versus paper losses,” *American Economic Review*, 106 (2016), 2086–2109.

- Jiang, Zhengyang, Hongqi Liu, Cameron Peng, and Hongjun Yan, "Investor Memory and Biased Beliefs: Evidence from the Field," *Quarterly Journal of Economics*, 140 (2025), 2749–2804.
- Johnson, Cathleen, Aurélien Baillon, Han Bleichrodt, Zhihua Li, Dennie Van Dolder, and Peter Wakker, "Prince: An improved method for measuring incentivized preferences," *Journal of Risk and Uncertainty*, 62 (2021), 1–28.
- Johnson, Eric J., Colin Camerer, Sankar Sen, and Talia Rymon, "Detecting Failures of Backward Induction: Monitoring Information Search in Sequential Bargaining," *Journal of Economic Theory*, 104 (2002), 16–47.
- Kahneman, Daniel, and Amos Tversky, "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 47 (1979), 263–291.
- Karni, Edi, and Zvi Safra, "Behaviorally consistent optimal stopping rules," *Journal of Economic Theory*, 51 (1990), 391–402.
- Klein Teeselink, Bouke, Dennie van Dolder, Martijn J. van den Assem, and Jason D. Dana, "High-stakes failures of backward induction," *Games and Economic Behavior*, (2024).
- Köbberling, Veronika, and Peter P. Wakker, "An index of loss aversion," *Journal of Economic Theory*, 122 (2005), 119–131.
- Langer, Thomas, and Martin Weber, "Does commitment or feedback influence myopic loss aversion?: An experimental analysis," *Journal of Economic Behavior & Organization*, 67 (2008), 810–819.
- Lauzier, Jean-Gabriel, Liyuan Lin, Peter Wakker, and Ruodu Wang, "Optimal risk sharing, equilibria, and welfare with empirically realistic risk attitudes," *Working paper available at arXiv*, (2024).
- Levitt, Steven D., "Heads or tails: The impact of a coin toss on major life decisions and subsequent happiness," *Review of Economic Studies*, 88 (2021), 378–405.
- Levitt, Steven D., John A. List, and Sally E. Sadoff, "Checkmate: Exploring Backward Induction among Chess Players," *American Economic Review*, 101 (2011), 975–990.
- Linnainmaa, Juhani T., "Do limit orders alter inferences about investor performance and behavior?," *Journal of Finance*, 65 (2010), 1473–1506.
- Loomes, Graham, and Robert Sugden, "Regret theory: An alternative theory of rational choice under uncertainty," *Economic Journal*, 92 (1982), 805–824.
- Machina, Mark J., "Stochastic choice functions generated from deterministic preferences over lotteries," *The Economic Journal*, 95 (1985), 575–594.
- , "Dynamic consistency and non-expected utility models of choice under uncertainty," *Journal of Economic Literature*, 27 (1989), 1622–1668.
- Magnani, Jacopo, "Testing for the disposition effect on optimal stopping decisions," *American Economic Review (Papers & Proceedings)*, 105 (2015), 371–75.
- Magnani, Jacopo, Jean Paul Rabanal, Olga A. Rud, and Yabin Wang, "Efficiency of dynamic portfolio choices: An experiment," *Review of Financial Studies*, 35 (2022), 1279–1309.
- Merton, Robert C., and Paul A. Samuelson, "Continuous-time finance," (Blackwell, 1992).
- Moinas, Sophie, and Sebastien Pouget, "The Bubble Game: An Experimental Study of Speculation," *Econometrica*, 81 (2013), 1507–1539.
- Mosteller, Frederick, and Philip Nogee, "An experimental measurement of utility," *Journal of Political Economy*, 59 (1951), 371–404.
- Nielsen, Kirby, "Dynamic risk preferences under realized and paper outcomes," *Journal of Economic Behavior & Organization*, 161 (2019), 68–78.

- Nielsen, Kirby, and John Rehbeck, "When choices are mistakes," *American Economic Review*, 112 (2022), 2237–68.
- Odean, Terrance, "Are investors reluctant to realize their losses?," *Journal of Finance*, 53 (1998), 1775–1798.
- Payne, John W., James R. Bettman, and Eric J. Johnson, "Adaptive strategy selection in decision making.," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14 (1988), 534.
- Quiggin, John, "A theory of anticipated utility," *Journal of Economic Behavior & Organization*, 3 (1982), 323–343.
- Rabin, Matthew, "Inference by believers in the law of small numbers," *Quarterly Journal of Economics*, 117 (2002), 775–816.
- Redelmeier, Donald A., and Amos Tversky, "On the framing of multiple prospects," *Psychological Science*, 3 (1992), 191–193.
- Selten, Reinhard, "Die Strategiemethode zur Erforschung des eingeschränkt rationalen Verhaltens im Rahmen eines Oligopolexperiments," *Beiträge zur experimentellen Wirtschaftsforschung*, (1967), 136–168.
- Shefrin, Hersh, and Meir Statman, "The disposition to sell winners too early and ride losers too long: Theory and evidence," *Journal of Finance*, 40 (1985), 777–790.
- Shiryaev, Albert N., "Optimal stopping rules," (Springer, 2007).
- Stott, Henry P., "Cumulative prospect theory's functional menagerie," *Journal of Risk and Uncertainty*, 32 (2006), 101–130.
- Strack, Philipp, and Paul Viefers, "Too proud to stop: Regret in dynamic decisions," *Journal of the European Economic Association*, 19 (2021), 165–199.
- Thaler, Richard H., and Eric J. Johnson, "Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice," *Management Science*, 36 (1990), 643–660.
- Tversky, Amos, and Daniel Kahneman, "Advances in prospect theory: Cumulative representation of uncertainty," *Journal of Risk and Uncertainty*, 5 (1992), 297–323.
- Tversky, Amos, and Eldar Shafir, "The disjunction effect in choice under uncertainty," *Psychological Science*, 3 (1992), 305–310.
- Wakker, Peter P., "Prospect theory: For risk and ambiguity," (Cambridge University Press, 2010).
- Weitzman, Martin L., "Optimal Search for the Best Alternative," *Econometrica*, 47 (1979), 641–654.
- Xu, Zuo Quan, and Xun Yu Zhou, "Optimal Stopping Under Probability Distortion," *Annals of Applied Probability*, 23 (2013), 251–282.
- Yaari, Menahem E., "The dual theory of choice under risk," *Econometrica*, (1987), 95–115.
- Zhang, Xing, and Songfa Zhong, "Putting preference for randomization to work," *Working paper*, (2020).

Eliciting Stopping Times: Supplemental Appendix
(For Online Publication)

Sebastian Ebert*

Maximilian Voigt†

* Heidelberg University; Department of Economics. E-mail: sebastian.ebert@awi.uni-heidelberg.de.

† HEC Montréal; Department of Finance. E-mail: maximilian.voigt@hec.ca.


ABSTRACT

Sections SA.1 and SA.2 of this supplemental appendix (SA) contain figures and tables that report on additional data analysis. Section SA.3 elaborates on the machine-learning methods and results. Section SA.4 provides details on the estimation of the finite mixture models, and Section SA.5 provides details on the experiment and its design.

SA.1 Additional Figures


Figure SA.1.1 below provides additional information on the use of costly versus costless path-dependence and randomization. Figure SA.1.2 presents evidence on the types of paths stopped and the randomization probabilities used. Figure SA.1.3 plots the aggregate best-fitting decoloring process.

FIGURE SA.1.1.—Utilization, strategy types and decoloring for costly and costless path-dependence and randomization




Notes: The upper left panel of Figure SA.1.1 shows the utilization of path-dependent balls when they are costless ($N = 75$) or costly ($N = 67$), while the bottom left panel shows the usage of costless ($N = 88$) or costly ($N = 97$) randomized balls. The middle panel shows the share of loss-exit, gain-exit, and symmetric-exit strategies for the respective studies. The right panels show the median decoloring when path-dependence (upper row) or randomization (bottom row) are costly.

FIGURE SA.1.2.—Usage of path-dependence and randomization



Notes: The left panel of Figure SA.1.2 shows the share of stopped risk-taking when a path is reached from above or below in the *Path-dependent* Decoloring Tasks ($N = 142$), whereas the right panel shows the share of balls with the stopping probabilities among all randomized balls for subjects in the *Randomized* Decoloring Tasks ($N = 185$).

FIGURE SA.1.3.—Individual as well as aggregate decoloring process



Notes: The left panel of Figure SA.1.3 shows the decoloring process for two individual subjects, whereby the upper decoloring process is from a subject that used forward induction, and the lower decoloring process is from a subject that used backward induction. Panel (b) shows the best-fitting decoloring process over all $N = 240$ subjects in the *Basic* Decoloring Task during our second experimental session.

SA.2 Additional Tables

Table SA.2.1 provides details on the six strategy clusters identified using the K -means algorithm in Section III.C. Table SA.2.2 informs about the consistency measure C_i in different studies. Table SA.2.3 reports on the relationship between the dynamic consistency measure D_i° and the skewness of the outcome distribution $\text{Skew}(S_\tau)$ in the different tasks and studies.

TABLE SA.2.1.— Analysis of cluster constituents

Cluster	N	$\mathbb{E}(\tau)$	$\mathbb{E}(\tau \text{Gain})$	$\mathbb{E}(\tau \text{Loss})$	$\text{Skew}(S_\tau)$
1	177	4.24	4.62	4.62	-0.01
2	80	1.43	1.69	1.53	0.14
3	122	3.16	3.36	3.71	-0.13
4	69	2.88	4.09	2.66	0.67
5	94	3.60	4.92	2.77	0.81
6	98	4.27	4.82	3.79	0.40

Notes: Table SA.2.1 shows the number of strategies for each of the six clusters identified using the K -means algorithm in Section III.C, as well as the average moments of the stopping time τ and induced outcome distribution S_τ for each cluster.

TABLE SA.2.2.—Consistency measure C_i

Study	Mean	CI Lower	CI Upper
<i>Basic</i>	0.8919	0.8712	0.9127
<i>No default</i>	0.8553	0.8165	0.8942
<i>No reminder</i>	0.8265	0.7845	0.8685
<i>Threshold</i>	0.8549	0.8072	0.9026
<i>No planning</i>	0.8348	0.8183	0.8512
<i>Path-dependence</i>	0.8357	0.7857	0.8857
<i>Randomization</i>	0.8141	0.7590	0.8692

Notes: Table SA.2.2 shows the mean as well as the upper and lower 95%-confidence interval for the within-subject consistency measure C_i .

TABLE SA.2.3.—Dynamic consistency measures D_i° and $\text{Skew}(S_\tau)$

	D_i°			$\text{Skew}(S_\tau)$		
	Mean	CI Lower	CI Upper	Mean	CI Lower	CI Upper
Risk-taking plan						
<i>Threshold</i>	0.2498	0.1581	0.3416	0.3883	0.2537	0.5229
<i>Basic</i>	0.2073	0.1715	0.2430	0.3593	0.3018	0.4168
<i>Path-dependence</i>	0.1285	0.0744	0.1826	0.2288	0.1482	0.3094
<i>Randomization</i>	0.0437	0.0068	0.0805	0.0513	-0.0233	0.1259
Risk-taking action						
<i>Basic</i>	0.0827	0.0399	0.1255	0.0280	-0.1098	0.1659
<i>No Default</i>	0.0831	0.0243	0.1420	0.1824	0.0470	0.3177
<i>Threshold</i>	0.0611	-0.0076	0.1298	0.1146	-0.0317	0.2609
<i>Path-dependence</i>	0.0269	-0.0242	0.0780	0.1293	0.0193	0.2392
<i>Randomization</i>	-0.0047	-0.0397	0.0303	-0.0276	-0.1153	0.0600
<i>No reminder</i>	-0.0289	-0.0719	0.0140	-0.1655	-0.2903	-0.0408
<i>No planning</i>	-0.0771	-0.1090	-0.0452	-0.1133	-0.2171	-0.0095

Notes: Table SA.2.3 shows the mean as well as the upper and lower 95%-confidence interval for our stop-loss intensity measure D_i° and $\text{Skew}(S_\tau)$ for different studies and types of the task.

SA.3 Machine-learning methodology and robustness

This section complements Sections III.C and IV.B.1. In the first part, we describe the unsupervised machine-learning procedure in more detail, report the results of a robustness analysis, and separately analyze the clusters that emerge for *Basic*, *Path-dependent*, and *Randomized* risk-taking strategies. Thereafter, we corroborate the findings from Section IV.B.1.

SA.3.1 Unsupervised machine-learning—Individual-level analysis

In Section III.C of the main text, we conduct an individual-level analysis of the elicited risk-taking strategies using a K -means clustering algorithm. Here, we first describe the K -means clustering and our implementation thereof. Afterwards, we highlight the robustness of our results by replicating the analysis of Section III.C with a different number of clusters¹ and

¹ We also performed the K -means clustering after reducing the dimensionality of the space spanned by the risk-taking strategies. We first conducted a Principal Components Analysis (PCA) and projected the risk-taking strategies onto the space spanned by the first four principal components, which retains 72% of the variance.

separately for all risk-taking strategies elicited from the *Basic*, *Path-dependent*, and *Randomized Decoloring Tasks*, including their costly variations. The findings highlighted in Section III.C—the separation of unconstrained risk-taking strategies into economically meaningful clusters that can be described as buy-and-hold, never-start, take-profit, and stop-loss strategies—replicate across all implementations.

The K -means clustering algorithm belongs to the class of unsupervised machine-learning algorithms that do not learn an assignment of data to specific outcome variables from labeled (training) data. Instead, the K -means clustering algorithm takes a set of unordered data points—640 precommitted risk-taking strategies that are represented as 31-tuples with the planned risk-taking action at each subpath of the ball triangle—and partitions the data into $K \in \mathbb{N}^+$ clusters by minimizing the within-cluster variance, while maximizing the between-cluster variance. The algorithm selects K random points from the dataset as initial cluster centers (called *centroids*). Each data point is then assigned to the closest centroid based on the Euclidean distance. In the next step, the mean of each cluster is calculated and becomes the new centroid. The algorithm iteratively performs these two steps until convergence. The K -means algorithm is known to converge to local minima (see Bishop 2006), which makes the resulting clustering sensitive to the initial centroids. Whenever we implement the K -means algorithm, we therefore run 400 iterations with random initial centroids and 400 iterations using the k -means++ algorithm (Arthur and Vassilvitskii 2007).² Afterwards, we use a majority voting over the 800 iterations of the K -means clustering.


The number of clusters K is a hyperparameter that needs to be chosen before performing the K -means clustering. Using standard methods³ to determine the number of clusters, we find that $K = 9$ clusters might lead to a partition of the data that provides a good fit. We focus on $K = 6$ clusters in Section III.C because it is more parsimonious and shows economically very similar results.

Thereafter, we implemented the K -means clustering as described here. The results support the robustness of our findings and are available upon request.

² The k -means++ algorithm does not sample the initial centroids randomly but such that they differ from each other. Therefore, the algorithm samples the initial centroids sequentially, with a probability that is higher the further a datapoint in the sample is away from the already sampled centroids.

³ A standard elbow plot of the *inertia* (sum of the squared distances of a datapoint from its cluster center) is inconclusive, as inertia is declining smoothly with the number of clusters. Additionally, we observe local minima of the clustering performance for 6 and 9 clusters in the *AIC* and *BIC* of a *Gaussian mixture* model. All methods suggest that using more clusters yields a better performance.

FIGURE SA.3.1.— Median risk-taking strategy per cluster



Notes: Figure SA.3.1 shows the median decoloring per cluster resulting from a K -means clustering with $K = 9$ clusters. We use a majority voting over 800 iterations of the K -means clustering with 400 iterations for the random and k-means++ initialization.

Figure SA.3.1 shows the median decoloring for each of the nine clusters and the number of loss-exit, gain-exit, and symmetric-exit strategies, while Table SA.3.1 shows the average moments of the stopping time τ and induced outcome distribution S_τ per cluster. The median strategy of clusters 1 and 2 corresponds to a buy-and-hold strategy, whereas the median strategy of cluster 3 corresponds to a never-start strategy and that of cluster 4 to a take-profit strategy. The median strategy of cluster 5 is more ambiguous and could be interpreted as a take-profit or buy-and-hold strategy. In contrast, the median strategies of clusters 6, 7, 8, and 9 are (trailing) stop-loss strategies. We corroborate this interpretation through the analysis in Table SA.3.1.

Combining the findings from Figure SA.3.1 and Table SA.3.1 leads to the cluster interpretations as in Table SA.3.2. Compared to the results in the main text, we find more subjects in clusters that we interpret as buy-and-hold (209 compared to 177) and stop-loss strategies (295 compared to 261), and less subjects in clusters that can be interpreted as take-profit (88 compared to 122) or never-start respectively taking-the-risk-once strategies (48

compared to 80). In addition, we now find a clear never-start cluster, whereas we interpreted cluster 2 in the main text as noisily taking-the-risk-once. Overall, our findings from the individual-level analysis—especially the interpretability of unconstrained risk-taking strategies in terms of simpler strategy classes—is robust to the number of clusters used.

TABLE SA.3.1.—Average moments of the stopping time τ and induced outcome distribution S_τ for each cluster with $K = 9$

Cluster	N	$\mathbb{E}(\tau)$	$\mathbb{E}(\tau \text{Gain})$	$\mathbb{E}(\tau \text{Loss})$	Skew(S_τ)
1	171	4.42	4.64	4.62	0.00
2	18	4.38	4.70	4.69	0.01
3	48	0.74	0.88	0.80	0.07
4	88	2.77	3.05	3.67	-0.24
5	20	4.19	4.07	4.38	-0.12
6	68	3.08	3.40	2.99	0.21
7	58	2.57	4.37	1.93	1.17
8	79	3.76	4.90	2.96	0.69
9	90	4.27	4.87	3.74	0.20

Notes: Table SA.3.1 shows the assignment of strategies to the nine clusters identified using a K -means algorithm, as well as the average (conditional) moments of the stopping time τ and induced outcome distribution S_τ for each cluster.

TABLE SA.3.2.—Assignment of clusters to economically meaningful strategy types

Strategy type	Clusters	N	Main text
Buy-and-hold	1, 2, 5	209	177
Never-start / Taking-the-risk-once	3	48	80
Take-profit	4	88	122
Stop-loss	6, 7, 8, 9	295	261


Notes: Table SA.3.2 shows the assignment of the clusters to economically interpretable strategy types, based on the results from Figure SA.3.1 and Table SA.3.1. The last column shows the number of subjects per class as in the main text.

As a further robustness analysis, we replicate the analysis from the main text for all risk-taking strategies that we elicited in the *Basic* ($N = 313$), *Path-dependent* ($N = 142$) and

Randomized ($N = 185$) Decoloring Tasks, including their costly counterparts, separately using $K = 6$ clusters. The results are shown in Figure SA.3.2, Figure SA.3.3, and Figure SA.3.4, respectively.

For all three sample splits, the median risk-taking strategy of cluster 1 is a buy-and-hold strategy, while that of cluster 2 is a never-start or taking-the-risk-once strategy. Among the *Basic* risk-taking strategies, the median strategy of cluster 3 is a take-profit strategy, those of clusters 4 and 5 are stop-loss strategies, and the median strategy of cluster 6 is a two-threshold loss-exit strategy in which the gain-threshold is further away from a profit level of 0 than the loss-threshold. Note that equally many subjects in cluster 6 plan to continue or stop risk-taking at the second highest ball in column five, leading to the median risk-taking plan at this ball being to stop risk-taking with 50% probability. A relative majority of the *Basic* risk-taking strategies is assigned to clusters interpretable as (trailing) stop-loss strategies.


FIGURE SA.3.2.—Median risk-taking strategy per cluster (*Basic*)



Notes: Figure SA.3.2 shows the median decoloring per cluster resulting from a K -means clustering with $K = 6$ clusters. We use a majority voting over 800 iterations of the K -means clustering with 400 iterations for the random and k-means++ initialization. We only include risk-taking strategies elicited in the *Basic* Decoloring Tasks.


The median risk-taking strategy of cluster 3 among the *Path-dependent* risk-taking strategies in Figure SA.3.3 is a mixture between a buy-and-hold and a gain-exit strategy, while the median strategies of the remaining clusters are (trailing) stop-loss strategies. Again, a relative majority of risk-taking strategies is assigned to clusters interpretable as (trailing) stop-loss strategies.

FIGURE SA.3.3.—Median risk-taking strategy per cluster (*Path-dependent*)



Notes: Figure SA.3.3 shows the median decoloring per cluster resulting from a K -means clustering with $K = 6$ clusters. We use a majority voting over 800 iterations of the K -means clustering with 400 iterations for the random and k-means++ initialization. We only include risk-taking strategies elicited in the *Path-dependent* and *Path-dependent^C* Decoloring Tasks.

FIGURE SA.3.4.—Median risk-taking strategy per cluster (*Randomized*)



Notes: Figure SA.3.4 shows the median decoloring per cluster resulting from a K -means clustering with $K = 6$ clusters. We use a majority voting over 800 iterations of the K -means clustering with 400 iterations for the random and k-means++ initialization. We only include risk-taking strategies elicited in the *Randomized* and *Randomized^C* Decoloring Tasks.

The median risk-taking strategy for cluster 6 in Figure SA.3.4 of the *Randomized* risk-taking strategies also seems to be a (trailing) stop-loss strategy and that of cluster 5 seems to be a take-profit strategy. Cluster 1's (Cluster 2's) median strategy is buy-and hold (never-start.) The remaining clusters 3 and 4 are less clearly interpretable. Indeed, the median risk-taking strategy

of cluster 4 seems to be a mixture between a buy-and-hold strategy and a two-threshold strategy with dark balls in the very top and very bottom of the ball triangle. We visually inspected all 39 risk-taking strategies assigned to cluster 3 of the *Randomized* risk-taking strategies and most strategies seem to be “random randomizers” without clear patterns in the stopping probabilities and a tendency towards stopping with medium probabilities around 50%. In line with our observation from Section III.B, allowing for the usage of *Randomization* has a larger effect on the risk-taking strategies than allowing for *Path-dependence*.

SA.3.2 Classification—Differentiating trailing and threshold stop-loss

We now complement the analysis from Section IV.B.1 of the main text and analyze what kind of stop-loss strategies subjects use. We include buy-and-hold, never-start, and trailing as well as threshold take-profit strategies in our analysis to reduce potential confounds from these strategy types.


The analysis here builds on supervised classification algorithms. In contrast to unsupervised methods (used before), these algorithms learn the assignment between features of the data and pre-defined classes using labeled data. We complement the analysis in the main text by using (a) a broader set of template strategies and (b) a broader set of methods, including cosine-similarity and three machine-learning algorithms. Regardless of the method, we find that trailing stop-loss strategies are more popular than threshold stop-loss strategies.

Basic template strategies. We construct a comprehensive set of all time- and outcome-contingent strategies for the following strategy types, as shown in Figure SA.3.5:

- *Never-start* strategies have a black first ball, such that the decoloring is black for all balls.
- *Buy-and-hold* strategies are time- but not outcome-contingent, such that the balls are white until the stopping repetition is reached, and black to the right of the stopping repetition.
- *Threshold stop-loss* strategies are outcome- but not time-contingent, such that the balls are white above the stop-loss level, and black below the stop-loss level.
- *Trailing stop-loss* strategies are time- and outcome-contingent with a stop-loss level that is increasing with the risk repetition. Trailing stop-loss strategies can be described by the number of risks lost before stopping, called the trailing stop-loss level.
- *Threshold and trailing take-profit* strategies are defined as the mirror image of the trailing or threshold stop-loss strategies.

Furthermore, we include *two-threshold* and *two-trailing strategies*, which are a combination of the threshold or trailing stop-loss and take-profit strategies, respectively.⁴ We also define one indeterminate stop-loss strategy, one indeterminate take-profit strategy, and one indeterminate two-threshold (or two-sided) strategy. Due to the finite time horizon ($T = 5$), they cannot be unambiguously assigned to the threshold or trailing type (see Figure SA.3.5).⁵

FIGURE SA.3.5.—All Basic risk-taking strategies for each class label



Notes: Figure SA.3.5 shows the comprehensive set of 25 *Basic* template strategies, with 3 threshold stop-loss, 3 trailing-stop-loss, 3 threshold take-profit, 3 trailing take-profit, 2 symmetric two-threshold, 2 symmetric two-trailing, 5 buy-and-hold, and 1 never start strategy. In addition, one stop-loss, one take-profit, and one two-threshold strategy are neither clearly a threshold nor a trailing strategy due to the finite horizon.

We define neither path-dependent nor randomized templates and thus expect a lower similarity of path-dependent or randomized risk-taking strategies to the templates. Because we compare the similarity across different templates only within-subject, between-subject differences in the level of similarity do not affect the results. As an example, if a subject randomizes at the very first ball, this ball must be changed for all templates (which are not

⁴ We only include symmetric two-threshold and two-trailing strategies as their asymmetric counterparts are subsumed by the one-sided strategies in all analyses.

⁵ For simplicity, in the analysis we split these indeterminate cases evenly between threshold and trailing strategies rather than letting them define two separate classes, each containing only a single template. This simplification works against our result that trailing stop-loss strategies are more common than threshold stop-loss strategies, although the effect is very small due to the small number of indeterminate stop-loss strategies.

randomized). The classification of path-dependent or randomized strategies relies only on balls that are not path-dependent or randomized.


TABLE SA.3.3.—Hamming-distance based classification of elicited risk-taking strategies for all elicited risk-taking strategies

Cluster	Classification based on the Hamming distance								Total
	Threshold stop-loss	Trailing stop-loss	Buy-and-hold	Never-start	Threshold take-profit	Trailing take-profit	Sym. two-threshold	Sym. two-trailing	
1	22.33	8.33	105.97	1.07	11.08	19.08	4.58	4.58	177
2	0.84	8.51	20.48	25.70	2.09	5.59	1.84	14.93	80
3	8.90	8.23	35.16	3.06	14.20	26.03	15.88	10.54	122
4	15.36	30.53	15.39	0.34	1.04	3.62	1.07	1.66	69
5	19.87	70.87	3.25	-	-	-	-	-	94
6	40.87	40.71	0.25	-	-	-	3.83	12.33	98
Total	108.17	167.18	180.50	30.17	28.41	54.32	27.19	44.03	640
Share	16.9%	26.1%	28.2%	4.7%	4.4%	8.5%	4.2%	6.9%	

Notes: Table SA.3.3 shows the assignment of strategies to the classes of the 25 template strategies based on the Hamming distance. We allow for partial assignment of an elicited risk-taking strategy to multiple classes if multiple template strategies from different classes are most similar to the elicited risk-taking strategy. We split indeterminate strategies evenly between the threshold and trailing type.

In the main text, we estimate the distance of each elicited risk-taking strategy to either the threshold strategies (threshold stop-loss, threshold take-profit, indeterminate stop-loss, two-threshold, and indeterminate two-threshold in Figure SA.3.5) or to the trailing strategies (trailing stop-loss, trailing take-profit, indeterminate stop-loss, two-trailing, and indeterminate two-threshold in Figure SA.3.5) by counting the number of balls that must be changed to transform the elicited risk-taking strategy into the respective template strategy (Hamming distance). Table SA.3.3 shows the classification when including all template strategies in Figure SA.3.5. We distribute indeterminate strategies evenly between the threshold and trailing type and allow for a partial assignment of the elicited risk-taking strategy to several classes. If an elicited risk-taking strategy, for example, has the same Hamming distance to threshold and trailing stop-loss strategies, then we assign it with a weight of 0.5 to both classes. Overall, Table SA.3.3 shows that the highest number of strategies is assigned to buy-and-hold strategies, reflecting the high propensity of planning to take the risk often. We find almost as many trailing stop-loss strategies as buy-and-hold strategies, and a proportion of trailing stop-loss strategies to threshold stop-loss of 1.55:1.

FIGURE SA.3.6.—Share of strategies exceeding the similarity threshold as a function of the threshold and exemplary classifications



Notes: Panel (a) of Figure SA.3.6 plots the share of elicited risk-taking strategies whose cosine similarity to the closest template exceeds the threshold on the x-axis. Strategies are assigned to one of eight classes (never-start; threshold or trailing stop-loss; buy-and-hold; threshold or trailing take-profit; symmetric two-threshold; symmetric two-trailing) if their maximum similarity exceeds the threshold. Indeterminate stop-loss and take-profit strategies are split evenly between threshold and trailing types. Panel (b) reports the resulting classifications for thresholds 0.6 and 0.9; strategies below the threshold are labeled “None.”

We complement the analysis of the main text by matching each elicited strategy to the template strategy to which it has the highest cosine-similarity.⁶ We assign an elicited strategy to a strategy class if the similarity to the matched template strategy is above a threshold. Otherwise, the elicited risk-taking strategy is not classified. Panel (a) of Figure SA.3.6 shows the share of elicited strategies whose highest similarity to a template strategy exceeds the threshold-level for varying thresholds. We find that almost all elicited risk-taking strategies can be classified if the threshold is at or below 0.6, while only around 61% of the elicited risk-taking strategies can be classified if the threshold is 0.9. Panel (b) of Figure SA.3.6 shows the classifications that arise for thresholds 0.6 and 0.9. For both thresholds, we find that most strategies are classified as buy-and-hold or trailing stop-loss strategies. Trailing stop-loss strategies are more common than threshold stop-loss strategies (27% to 13% for a similarity threshold of 0.6, and 18% to 9% for a similarity threshold of 0.9). Additionally, Table SA.3.4 shows the classification separately for each type of Decoloring Task, whereby we do not impose a similarity threshold. We find that trailing stop-loss strategies are more popular than threshold stop-loss strategies, with a ratio of trailing to threshold stop-loss strategies ranging

⁶ We represent the elicited risk-taking strategies as 31-tuples. Let the risk-taking strategy of subject i be denoted by $\tau_i \in \mathbb{R}^{31}$, while the template strategy j in class c is $v_{c,j} \in \mathbb{R}^{31}$. The cosine similarity is then $S_c(\tau_i, v_{c,j}) = \frac{\tau_i \cdot v_{c,j}}{\|\tau_i\| \|v_{c,j}\|}$, with dot product $\cdot$ and the Euclidean norm $\|\cdot\|$.

from 1.61:1 (*Basic*) to 2.82:1 (*Randomized*). Using a cosine-similarity measure instead of the change-count corroborates our finding that trailing stop-loss strategies are more popular than threshold stop-loss strategies.

TABLE SA.3.4.—Cosine-similarity based classification of elicited risk-taking strategies for each type of Decoloring Task

Study	Threshold stop-loss	Trailing stop-loss	Buy-and-hold	Never-start	Threshold take-profit	Trailing take-profit	Sym. two-threshold	Sym. two-trailing	TSL / SL
<i>Basic</i>	56.00	109.50	82.50	11.00	2.00	24.00	9.50	18.50	1.96
<i>Path-dep.</i>	16.50	35.50	48.50	7.00	4.00	12.50	7.50	10.50	2.15
<i>Rand.</i>	11.00	31.00	49.00	47.00	8.50	20.50	1.50	16.50	2.82
Total	83.50	176.00	180.00	65.00	14.50	57.00	18.50	45.50	2.11
Share	13.05%	27.50%	28.13%	10.16%	2.27%	8.91%	2.89%	7.11%	

Notes: Table SA.3.4 shows the assignment of strategies to the classes of the 25 template strategies based on the cosine-similarity. We partially assign an elicited risk-taking strategy to a class if multiple template strategies from different classes are most-similar to the elicited risk-taking strategy and we split indeterminate stop-loss and take-profit strategies evenly between the threshold and trailing type.

Furthermore, we use a logistic regression, a decision tree, and a random forest classifier to map the elicited risk-taking strategies to the template strategies (for a more formal description, see Bishop 2006):

1. *Logistic regression:* We fit a separate logistic regression for each strategy class, using the representation of decolorings as 31-tuples as for the cosine-similarity. We assign an elicited risk-taking strategy to a strategy class if the projection of the elicited strategy onto the logistic regression weights yields the highest odds-ratio among all strategy classes.
2. *Decision tree:* Decision trees sequentially separate data into different branches and nodes of a flowchart-like tree. The algorithm may, for example, choose a certain entry in the 31-tuple that represents the decoloring (a stopping plan after a subpath in the ball triangle) as the feature for the very first node in the decision tree. If the stopping decision is larger than a threshold value, the algorithm moves “right” in the decision tree and “left” otherwise. This continues until a so-called leaf is reached, which assigns a strategy to a strategy class. We use the Gini impurity to design the decision tree, implying that the feature and threshold value for any node of the decision tree are chosen to minimize the Gini impurity.⁷

⁷ For a motivating example see <https://victorzhou.com/blog/gini-impurity/>.

3. *Random forest*: Random forests combine many decision trees. The algorithm repeatedly collects a random subset of the training data and trains a decision tree on each subset. The classification of a strategy is eventually obtained as the majority vote from many decision trees.

TABLE SA.3.5.—Class assignment for different supervised machine-learning algorithms

Class	Logistic regression	Decision tree	Random forest
Never-start	59	28	27
Buy and hold	188	203	228
Threshold stop-loss	120	115	108
Trailing stop-loss	156	169	164
Threshold take-profit	22	25	22
Trailing take-profit	60	35	49
Symmetric two-threshold	18	27	18
Symmetric two-trailing	17	38	24
Trailing stop-loss / Threshold stop-loss	1.30	1.47	1.52

Notes: Table SA.3.5 shows the assignment of elicited risk-taking strategies to the eight strategy classes (first column) using three different supervised machine-learning procedures (columns 2 to 4). The bottom rows show the proportion of trailing to threshold stop-loss strategies.

Table SA.3.5 shows the classification obtained from each of the three supervised machine-learning algorithms. Across the algorithms, buy-and-hold as well as trailing stop-loss strategies are the most common strategy types. All algorithms classify more risk-taking strategies as trailing stop-loss than as threshold stop-loss (and as trailing take-profit than as threshold take-profit). Depending on the method, we find that the ratio of trailing stop-loss to threshold stop-loss strategies is 1.30:1 (logistic regression), 1.47:1 (decision tree), 1.52:1 (random forest), 1.55:1 (count-based similarity), or 2.11:1 (cosine-similarity).

SA.4 Finite mixture model

We now describe the finite mixture model used to estimate the different types of CPT agents in Section III.D of the main text.

Suppose that there are K different types of CPT agents in the data, whose value and probability weighting functions are as in the main text. Each CPT agent type $k = 1, 2, \dots, K$ can be characterized by a parameter vector $\theta_k = (\alpha_k, \delta_k, \lambda_k, \epsilon_k)$, and we denote the share of the different types in our sample as q_k , with $\sum_k q_k = 1$. Recall that we observe one chosen

stopping time τ_c^i for each subject. The likelihood of observing the chosen stopping time τ_c^i if subject i is of type k is

$$P(\tau_c^i | T_{Basic}, \theta_k) = \prod_{\tau_{-} \in T_{Basic} \setminus \{\tau_c^i\}} f(\tau_c^i | \tau_{-}, \theta_k) = \prod_{\tau_{-} \in T_{Basic} \setminus \{\tau_c^i\}} \left(1 + e^{-\epsilon_k \Delta(\tau_c^i, \tau_{-}, \theta_k)}\right)^{-1},$$

where $\Delta(\tau_c^i, \tau_{-}, \theta_k) = V(S_{\tau_c^i} | \theta_k) - V(S_{\tau_{-}} | \theta_k)$. However, we do not know a priori to which type subject i belongs, such that we need to estimate the proportions of the different types in the sample (the probability of group membership). The log likelihood function of the finite mixture model is then given as

$$\begin{aligned} \ln L(\Theta | \{\tau_c^i\}_{i=1}^{313}, T_{Basic}) &= \ln \prod_{i=1}^{313} \sum_{k=1}^K q_k P(\tau_c^i | T_{Basic}, \theta_k) \\ &= \sum_{i=1}^{313} \ln \left(\sum_{k=1}^K q_k P(\tau_c^i | T_{Basic}, \theta_k) \right), \end{aligned}$$

where $\Theta = (\theta_1, \dots, \theta_K, q_1, \dots, q_K)$ summarizes all parameters of the model.

The summation inside the log prevents a closed-form solution, and the log likelihood function is typically not directly maximized due to numerical problems (see Bruhin, Fehr-Duda, and Epper 2010). However, if individual group-memberships were observable and indicated by $g_k^i \in \{0,1\}$, the contribution of subject i to the total likelihood function would be given by

$$\tilde{l}(\Theta_i | \tau_c^i, T_{Basic}) = \prod_{k=1}^K [q_k P(\tau_c^i | T_{Basic}, \theta_k)]^{g_k^i}.$$

We can use the above formulation under the assumption of observing the latent variable g_k^i for all subjects—summarized in matrix g —to obtain the complete-data log likelihood as

$$\ln \tilde{L}(\Theta | \{\tau_c^i\}_{i=1}^{313}, T_{Basic}, g) = \sum_{i=1}^{313} \sum_{k=1}^K g_k^i [\ln q_k + \ln P(\tau_c^i | T_{Basic}, \theta_k)].$$

The Expectation-Maximization (EM) algorithm proceeds iteratively to maximize the complete-data log likelihood. In the $(j+1)$ -th E-step of the EM algorithm, the expectation of the complete-data log likelihood is computed given the estimates from the last step, $\Theta^{(j)}$. By Bayes rule, this yields the posterior probability of individual group-membership

$$\tilde{g}_k^i(\tau_c^i, \Theta^{(j)}, T_{Basic}) = \frac{q_k^{(j)} P(\tau_c^i | T_{Basic}, \theta_k^{(j)})}{\sum_{k'=1}^K q_{k'}^{(j)} P(\tau_c^i | T_{Basic}, \theta_{k'}^{(j)})}.$$

Given the posterior probabilities of individual group-membership, we maximize the remaining parameters in the M-step of the algorithm. The parameters are updated as

$$q_k^{(j+1)} = \frac{1}{313} \sum_{i=1}^{313} \widetilde{g}_k^i(\tau_c^i, \theta_i^{(j)}, T_{Basic})$$

and

$$(\theta_1^{(j+1)}, \dots, \theta_K^{(j+1)}) = \arg \max_{\theta_1, \dots, \theta_K} \sum_{i=1}^{313} \sum_{k=1}^K \widetilde{g}_k^i(\tau_c^i, \theta^{(j)}, T_{Basic}) \ln P(\tau_c^i | T_{Basic}, \theta_k),$$

which we optimize separately for each agent type k . It is well known that the likelihood never decreases from one iteration to another, such that the EM algorithm converges on (local) maxima (Dempster, Laird, and Rubin 1977). We implement the numerical optimization using the *L-BFGS-B* algorithm from the SciPy-library in Python and run the EM-algorithm with 15 different initializations to increase the chance of finding a global instead of a local maximum. We report the results that yield the highest likelihood for a given number of agent types k over all 15 initializations.


The number of types k is determined exogenously, either guided by economic reasoning or model-comparison criteria. We perform the estimation with one, two, three... up until seven agent types. Table SA.4.1 shows the log-likelihood, AIC, and BIC for the different specifications. Although the specification with one CPT-agent is economically informative because it yields the single parameter vector that is most consistent with the data, a representative agent model provides the worst fit to the data among the models that we estimate.⁸ Table SA.4.1 highlights that specifications with more types are preferred over specifications with fewer types and we find that the best-performing model contains seven CPT agent types. If we closely analyze the estimates and type assignments for the different models, we find that increasing the number of types from k to $k + 1$ tends to lead to a division of one of the k types into two smaller types. Such a decomposition of an agent type into two smaller groups generally does not affect the behavioral interpretation. As we want to achieve a parsimonious characterization of different CPT agent types, we report the results with four types in the main text. Figure SA.4.1 and Figure SA.4.2 show the estimates for three and five agent types, respectively.

⁸ In contrast to the median parameter estimates that we obtained from the individual estimation, the one-cluster specification takes interdependencies between the parameters into account. Therefore, these specifications are not necessarily identical. We estimate $(\alpha, \delta, \lambda) = (1.00, 0.30, 1.00)$ with one agent type, which coincides with the median estimates from the individual estimation.

TABLE SA.4.1.— Model selection criteria for finite mixture models with up to seven types

Number of types	Log-Likelihood	AIC	BIC
1	-139,676.95	279,361.90	279,376.89
2	-105,550.47	211,118.95	211,152.66
3	-96,139.48	192,306.97	192,359.42
4	-91,170.89	182,379.79	182,450.96
5	-90,975.99	181,999.98	182,089.89
6	-88,278.00	176,614.01	176,722.65
7	-76,073.60	152,215.20	152,342.57


Notes: Table SA.4.1 shows model selection criteria for finite mixture models with up to seven agent types. We highlight the best value according to the respective criterion using a bold font.

FIGURE SA.4.1.— CPT parameter estimates and median strategies from a finite mixture model with three types

Notes: Figure SA.4.1 shows the median risk-taking strategies for each agent type. Next to each header, we show the parameter estimates for the agent type as well as the number of strategies. We find an almost perfect separation of the risk-taking strategies.

Comparing a setting with three agent types (Figure SA.4.1) and four agent types (Figure 8), we find that agent type 3 is identical in both settings. Moreover, while agent type 2 in Figure SA.4.1 includes two more risk-taking strategies than agent type 2 in Figure 8—which also marginally affects the parameter estimates—the behavioral interpretation of both agent types is identical. Indeed, the largest change when going from three to four agent types is a separation of agent type 1 in Figure SA.4.1 into the two agent types 1 and 4 in Figure 8. Behaviorally, the change implies a separation of the stop-loss types into two stop-loss types with different stop-loss levels.

FIGURE SA.4.2.— CPT parameter estimates and median strategies from a finite mixture model with five types



Notes: Figure SA.4.2 shows the median risk-taking strategies for each agent type. Next to each header, we show the parameter estimates for the agent type as well as the number of strategies. We find an almost perfect separation of the risk-taking strategies.

Proceeding in the same way for the comparison between four types (Figure 8) and five types (Figure SA.4.2), we find that agent type 1 is identical in both setting, while agent types 3 and 4 in Figure 8 are behaviorally equivalent to agent types 4 and 5 in Figure SA.4.2. The only behaviorally significant change is a separation of agent type 2 from Figure 8 into agent types 2 and 3 in Figure SA.4.2, which singles-out a dedicated take-profit CPT-agent type.

SA.5 Experimental details

The instructions and comprehension questions of the experiment are interactive, including dynamic screens, pop-up windows, and behavior-contingent feedback. As a static textual version of the instructions would be unnatural, excessively long, and inconvenient, we provide videos of all experimental studies using the links in Table SA.5.1. The videos allow for viewing any aspect of the experiment (e.g., instructions, comprehension questions, the payment screen) quickly. To get a comprehensive overview of the experimental design and all ten studies, we recommend viewing the overview video at <https://youtu.be/ParT-P9gX04>. All videos are posted on the channel “Eliciting Stopping Times – Experiment” at https://www.youtube.com/playlist?list=PLr_xX-ToY8MNLtLLa3176gLOlw8rcNLmy.

Moreover, one can go through interactive versions of all studies at

<https://rtsexpdemo.onrender.com/demo>.⁹ Below, we discuss the sample obtained from our experiment and payment details.

TABLE SA.5.1.—Overview of the experimental studies

#	Name	Decoloring Task	Walking Task	Link
1	Basic	Basic	Basic	https://youtu.be/ZRKuLjEtVUw
2	Path-dependence	Path-dependent	Basic	https://youtu.be/tlSHWp1Dl4A
3	Path-dependence ^C	Path-dependent ^C	Basic	https://youtu.be/zX3cpH7Og_k
4	Randomization	Randomized	Randomized	https://youtu.be/VcSZRVzFgVQ
5	Randomization ^C	Randomized ^C	Randomized ^C	https://youtu.be/cuSl8P-dQH8
6	Basic II	Basic	Basic	(same as Basic)
7	No default	Basic	No default	https://youtu.be/-GuQFqQMCgw
8	No reminder	Basic	No reminder	https://youtu.be/hQ4cFsMxToY
9	Threshold	Threshold	Basic	https://youtu.be/dW9xLnrcGUY
10	No planning	-	No reminder	https://youtu.be/O2bX87olnLU

Notes: Table SA.5.1 shows the different between-subject studies and the link to review the experimental instructions.

SA.5.1 Procedural details and sample

We programmed the experiment using oTree (Chen, Schonger, and Wickens 2016) and ran it via Prolific on November 16, 2022, and on May 23, 2024. We obtained $N = 881$ completes (440 female and 441 male). As pre-registered, we excluded subjects who failed an attention check or who did not complete all experimental tasks. Overall, 1,718 subjects started our experiment. 320 (19%) of them left before the first attention check. Out of the remaining 1,398 subjects, 186 (13%) failed the first attention check. Of the remaining 1,212 subjects, 58 (5%) left before the second attention check, and 181 of the remaining 1,154 subjects (16%) failed the second attention check. Of the remaining 973 subjects, 948 finished the experiment. 67 of these submissions had to be removed because the subject joined the experiment previously and failed an attention check. We only consider the data of the remaining 881 completes. Our participant selection filters lead to a final sample that is biased towards sophisticated and attentive subjects, as in standard lab experiments. The subject pool contains a small share of students (13%), and a relative majority of subjects are in full-time employment (49%, see Table SA.5.2).

⁹ Opening the browser window may take up to one minute. Once the window appears, click on “Risk-taking Strategies: Experiment” located under the “Demo” header. Then, within the new window, click the link shown below “Session-wide demo link” to open the study selection window of the experiment.

TABLE SA.5.2.—Descriptive sample statistics

		N	%
Gender	Male	441	50.06%
	Female	440	49.94%
Nationality	United Kingdom	692	78.55%
	United States	103	11.69%
	Ireland	31	3.52%
	Nigeria	18	2.04%
	Australia	13	1.48%
	India	5	0.57%
	Others	19	2.16%
Student	Yes	113	12.83%
	No	704	79.91%
	N/A	64	7.26%
Employment	Full-Time	431	48.92%
	Part-Time	136	15.44%
	Unemployed	59	6.70%
	Not in paid work	111	12.60%
	Starting a job soon	5	0.57%
	Other	39	4.43%
	n/a	100	11.35%

Notes: Table SA.5.2 shows the number and share of subjects by gender, nationality, student status, and employment status.

Table SA.5.3 shows that the mean age of our subjects is 39.6 years, and that they are experienced experimental subjects with an average of 860 completed studies on Prolific. On average, subjects took 28.5 minutes to complete the experiment, whereby the fastest subject took 10.2 minutes. The distribution of total earnings of the subjects is symmetric and centered at the mean and median of £4.50.

SA.5.2 Payment

Subjects are exogenously assigned to a study in order of joining the experiment.¹⁰ Before receiving instructions about the experiment, subjects clicked one of 11 (10 for *No Planning*) buttons to determine their individually payment-relevant task. After completing all experimental tasks and a short questionnaire, subjects learn their payoff-relevant task. The remaining risks are resolved if the Decoloring Task is the payoff-relevant task. If the bonus payment is determined by a Walking Task, the outcome of that task is paid as bonus. The

¹⁰ In the experimental session on November 16, 2022, we gave priority to studies with fewer completes or currently active subjects when assigning subjects to studies. We dropped this additional balancing mechanism for the second run of the experiment.

overall payment for the experiment consists of a completion fee of £3.00 (\$3.57 at the time), an endowment of £1.50 (\$1.85 at the time), and the task-dependent bonus.

TABLE SA.5.3.— Summary statistics of the sample and experimental procedure

	Mean	Median	Standard Deviation	Max	Min
Age (in years)	39.6	37.0	13.0	83.0	18.0
Duration (in min)	28.5	26.1	10.5	68.1	10.2
Total earnings (in GBP)	4.5	4.5	0.5	5.75	3.25
Studies completed on Prolific	859.5	595.5	863.0	4,984.0	0.0

Notes: Table SA.5.3 shows summary statistics of the age of subjects, their experience as measured by the number of studies they completed on Prolific before taking part in our experiment, and summary statistics of the experiment, mainly the duration and the total earnings.

REFERENCES

- Arthur, David, and Sergei Vassilvitskii, “K-means++ the advantages of careful seeding,” in *Proceedings of the eighteenth annual ACM-SIAM symposium on discrete algorithms*, 2007.
- Bishop, Christopher M., “Pattern recognition and machine learning,” Information science and statistics (New York, Springer, 2006).
- Bruhin, Adrian, Helga Fehr-Duda, and Thomas Epper, “Risk and rationality: Uncovering heterogeneity in probability distortion,” *Econometrica*, 78 (2010), 1375–1412.
- Chen, Daniel L., Martin Schonger, and Chris Wickens, “oTree—An open-source platform for laboratory, online, and field experiments,” *Journal of Behavioral and Experimental Finance*, 9 (2016), 88–97.
- Dempster, A. P., N. M. Laird, and D. B. Rubin, “Maximum Likelihood from Incomplete Data Via the EM Algorithm,” *Journal of the Royal Statistical Society*, 39 (1977), 1–22.