Eliciting Stopping Times

Sebastian Ebert Maximilian Voigt[†]

April 25, 2025

ABSTRACT

We propose an experimental method to elicit stopping times. Using an interactive tool, subjects specify complete contingent plans of when to continue or stop taking a given risk. We document five main results: (1) Stopping times differ significantly between subjects. A machine-learning algorithm classifies 39% of the strategies as stop-loss and 29% as buy-and-hold. (2) Trailing stop-loss strategies are 1.5 times more common than threshold stop-loss strategies. Restricting choices to threshold strategies does not affect aggregate stopping times. (3) A structural prospect theory estimation aligns closely with an unsupervised machine-learning algorithm, suggesting a good descriptive fit of prospect theory. (4) Most subjects use path-dependence and randomization if available. (5) 60% of subjects choose their stopping time by forward instead of backward induction (26%). We also compare planned with actual (sequential) risk-taking and document the causal effects of memory, defaults, planning constraints, and planning as such on dynamic consistency.

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

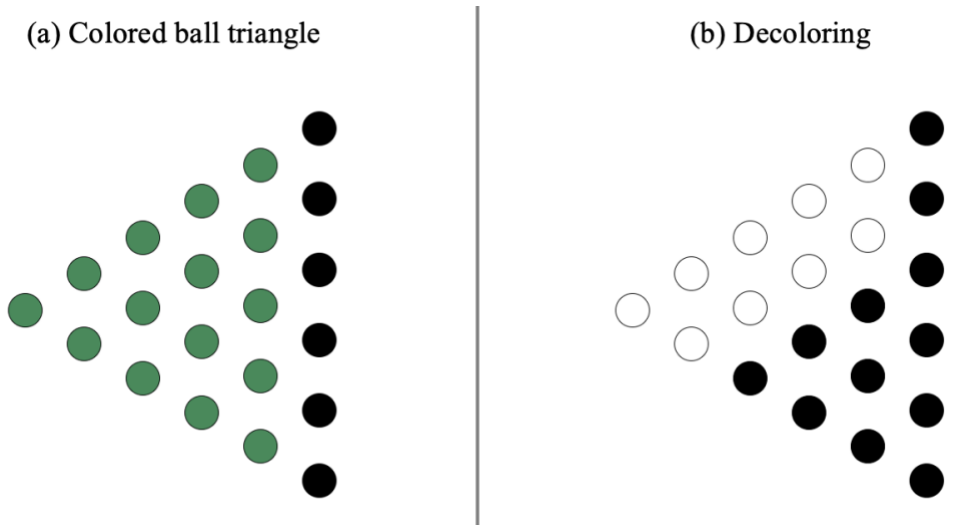
JEL Classification: D01, D81, D91

[†] 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, Maastricht, Oslo (BI), Heidelberg, 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, and the Research in Behavioural Finance Conference 2024 for helpful comments and suggestions. 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. 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. The experiments were pre-registered at the AEA RCT registry (AEARCTR-0009684 and AEARCTR-0013653). We thank Felix Schmücker and Alireza Shadmand for providing excellent research assistance.

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 (see Figure 1). 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 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. 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 shows 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

[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.

binomial tree. The left 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 thus 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 also take the risk sequentially without precommitment (details below).

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, Obłój, 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, thus supporting subjects' perception of the stopping time as a fully contingent risk-taking strategy beyond sequential risk-taking considerations.

Stopping times are complex, high-dimensional data objects. We show how machine learning methods in combination with structural estimation techniques can be used to analyze stopping times, and document five main results 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 stopping times and find that individual risk-taking strategies elicited under precommitment are comprehensively described as stop-loss, take-profit, or buy-and-hold strategies. Neither did the algorithm—by definition—learn these strategy types from training data, nor did the experimental design restrict the risk-taking strategies of subjects to certain types. Yet, we find economically interpretable clusters of unconstrained risk-taking strategies that have been predicted in theoretical work. Most subjects plan on taking the risk often, although the expected value of any induced

outcome distribution S_τ is zero (due to the optional stopping theorem), which is incompatible with risk-averse expected utility preferences. Instead, a relative majority of subjects plans to use stop-loss strategies (39%), followed by buy-and-hold strategies (29%).

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 illustrated in the right panel of Figure 1—are used about 1.5 times as often as threshold stop-loss strategies, which are strategies with a constant stopping barrier. We find no significant change in aggregate measures of stopping times when subjects are constrained to only using threshold strategies. The focus on threshold risk-taking strategies in theoretical, empirical, and experimental work, therefore, seems to be a reasonable approximation to aggregate risk-taking strategies, although it is a binding restriction for individuals.

Third, we extend the non-parametric analysis of risk-taking strategies by developing a procedure 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 an unsupervised machine-learning algorithm, which implies a good descriptive fit of prospect theory. Prospect theory might be a good starting point for modelling stopping behavior.

Fourth, 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.

Fifth, since the order of decoloring the balls of the ball triangle is endogenously chosen by subjects, we gain insight 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 (26%).

In the second part of the experiment, subjects take the risk sequentially. Each subject visually walks through the ball triangle and can freely deviate from her plan at any node along the realized path. Crucially, comparing each subject's unconstrained risk-taking plan to her sequential risk-taking actions allows us to interpret deviations from the plan as dynamic inconsistency. Note that a *mechanical inconsistency* can arise in sequential risk-taking when a subject must choose from a constrained set of strategies that precludes the desired strategy. Deviating from a strategy that one did not want in the first place does not constitute an inconsistency. Our experimental design without constraints allows for testing the overestimation of dynamic inconsistency that results from mechanical inconsistency. We test for mechanical inconsistency by constraining some subjects to threshold strategies and find an insignificant overestimation of dynamic inconsistency.

We furthermore test whether a default, memory, and planning as such causally affect dynamic inconsistency. Overall, subjects follow their risk-taking plan for about 87% of their sequential risk-taking actions. Plan-consistency is significantly lower when subjects do not explicitly plan (as inferred from comparing subjects that did not complete the Decoloring Task with subjects that did; cf. Cubitt, Starmer, and Sugden 1998), and when subjects need to remember their risk-taking strategy. Although consistency is high overall and at most balls, systematic inconsistencies at some balls lead to significant changes in behavior. We find that subjects exhibit a disposition effect in that sequential risk-taking is (more) gain-exit than planned.[2] That is, we find the "standard shift" from loss-exit risk-taking plans to gain-exit risk-taking actions (Barberis 2012; Heimer et al. 2025). The standard shift is weaker if subjects explicitly plan before taking the risk and if subjects have a memory-aid that shows them their plan while taking the risk. We find a small and statistically insignificant overestimation of the standard shift stemming from

[2] The disposition effect describes the tendency to sell winning stocks more frequently than losing stocks (Shefrin and Statman 1985; Odean 1998).

mechanical inconsistency. Offering a plan-consistent default is ineffective for reducing the standard shift.

By eliciting unconstrained stopping times, we hope to inspire experimental or 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 also 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 the repeated risk-taking and stopping literature by proposing and implementing an experimental method to elicit unconstrained stopping times. We also provide the first experimental evidence on the unconstrained stopping times that subjects use.

More directly, we contribute to the experimental stopping literature. A large share of this literature 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.[3] Fischbacher, Hoffmann, and Schudy (2017) show that commitment 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: Subjects plan to take a risk repeatedly by setting an absolutely higher stop-loss than take-profit threshold (a loss-exit plan), while their actual risk-taking behavior deviates from the loss-exit plan by cutting gains and chasing losses, as predicted by Barberis (2012) and as is consistent with the disposition effect. 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 loss-exit plan even if the stochastic process has a negative drift. In contrast, subjects in the experiment by Antler and Arad (2023) often select a gain-exit 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.[4] 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 2024) or risks with time-changing moments (Nielsen 2019), is well possible.

We also contribute to the growing literature 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 further combine machine-learning methods with a structural estimation of cumulative prospect theory preference parameters (Tversky and

[3] 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.

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

Kahneman 1992; Wakker 2010) by extending the method in Bruhin, Fehr-Duda, and Epper (2010) to stopping time data.

We further contribute to the experimental literature on dynamic (in-)consistency 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 investigate the possible overestimation of dynamic consistency that arises from confounding dynamic and mechanical inconsistency. 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.

A second contribution to the dynamic (in)consistency literature is the estimation of the causal effects of memory, defaults, planning constraints, and planning as such on dynamic consistency and the standard shift. Consistent with prior work that highlights the importance of memory for individual-level risk-taking (Gödker, Odean, and Smeets 2024; Jiang et al. 2024; Gödker, Jiao, and Smeets 2025), we find that a memory-aid in the form of a plan-reminder reduces dynamic inconsistency, as does planning as such.

II. Overview of experimental design

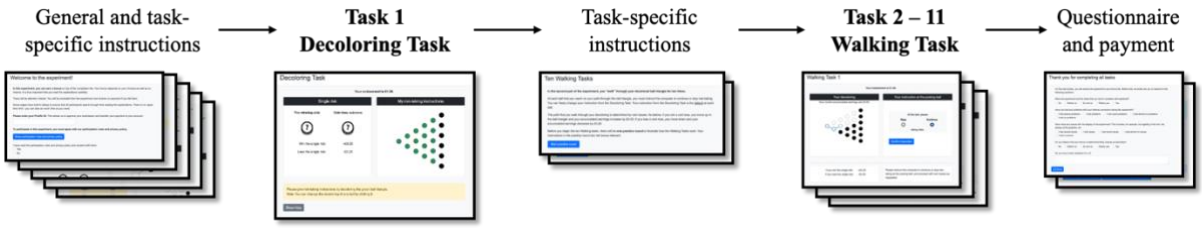
Eliciting unconstrained risk-taking strategies and subsequent stopping behavior requires a controlled experiment to avoid confounding external factors and since complete contingent plans are not observable otherwise. Our preregistered experiment is designed to enable a parsimonious and intuitive elicitation of unconstrained risk-taking strategies, as well as to provide comprehensive comparisons with sequential stopping behavior. We give an overview of the experimental design here and explain the experimental tasks in more detail in Sections III.A and IV.A. An overview video (https://youtu.be/ParT-

P9gX04), videos of all studies, and an interactive version of the experiment are available.[5]

II.A Structure of the experiment

The experiment consists of ten between-subject studies, which contain different 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), stopping times—the complete (time- and outcome-) contingent plan for taking the risk for up to five times—are represented through a recombining binomial tree, visualized and described as a ball triangle to subjects (see Figure 1). Each ball of the triangle represents one possible accumulated profit that can result from taking the risk up to five times.

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.

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

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

decoloring the balls of an initially colored ball triangle.[6] 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 between-subject 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 Supplemental Appendix SA.5).

TABLE 1.—Overview over the ten experimental studies

#	Name	Decoloring Task	Walking Task	Executed on
1	Basic I	Basic	Basic	Nov 16, 2022
2	Path-dependence	Path-dependent	Basic	Nov 16, 2022
3	Path-dependence^{C}	Path-dependent^{C}	Basic	Nov 16, 2022
4	Randomization	Randomized	Randomized	Nov 16, 2022
5	Randomization^{C}	Randomized^{C}	Randomized^{C}	Nov 16, 2022
6	Basic II	Basic	Basic	May 23, 2024
7	No default	Basic	No default	May 23, 2024
8	No reminder	Basic	No reminder	May 23, 2024
9	Threshold	Threshold	Basic	May 23, 2024
10	No planning	-	No reminder	May 23, 2024

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 type 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 was executed.

II.B Studies

Table 1 provides an overview of the ten between-subject studies of the experiment. 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, 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 as well as to continue risk-taking if the first risk is lost and the second risk is won (up-down versus down-up). Note

[6] 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.

that sequential risk-taking in the Walking Tasks is necessarily path-dependent, such that we do not have dedicated path-dependent Walking Tasks.[7] 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.[8] 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 or continue risk-taking for all balls that yield a profit of £0.25. 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.A).

III. Stopping times: Elicitation method and analysis

This section explains the elicitation of stopping times and their statistical analysis. We introduce the Decoloring Task, a novel method for eliciting stopping times. 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 a given risk for up to any fixed number of times, also called risk-taking strategy or stopping time. All choices are incentivized and stopping times are elicited under pre-commitment.

III.A Eliciting stopping times with the Decoloring Task

Figure 3 illustrates the interactive tool that subjects use to indicate, or design, their stopping times during the Decoloring Task. We implement it to elicit complete contingent plans for taking a risk up to five times. The first (second, third, fourth) row is for the *Basic (Path-dependent, Randomized, Threshold)* task. Following Johnson et

[7] 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, the subjects cannot use a costly option during their *Basic* Walking Tasks in study *Path-dependence*^{C}.

[8] 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.

al. (2021), subjects make their decisions by giving instructions to the computer.[9] 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 treatment *Threshold,* for all balls in that ball's row) 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 times according to the decoloring.

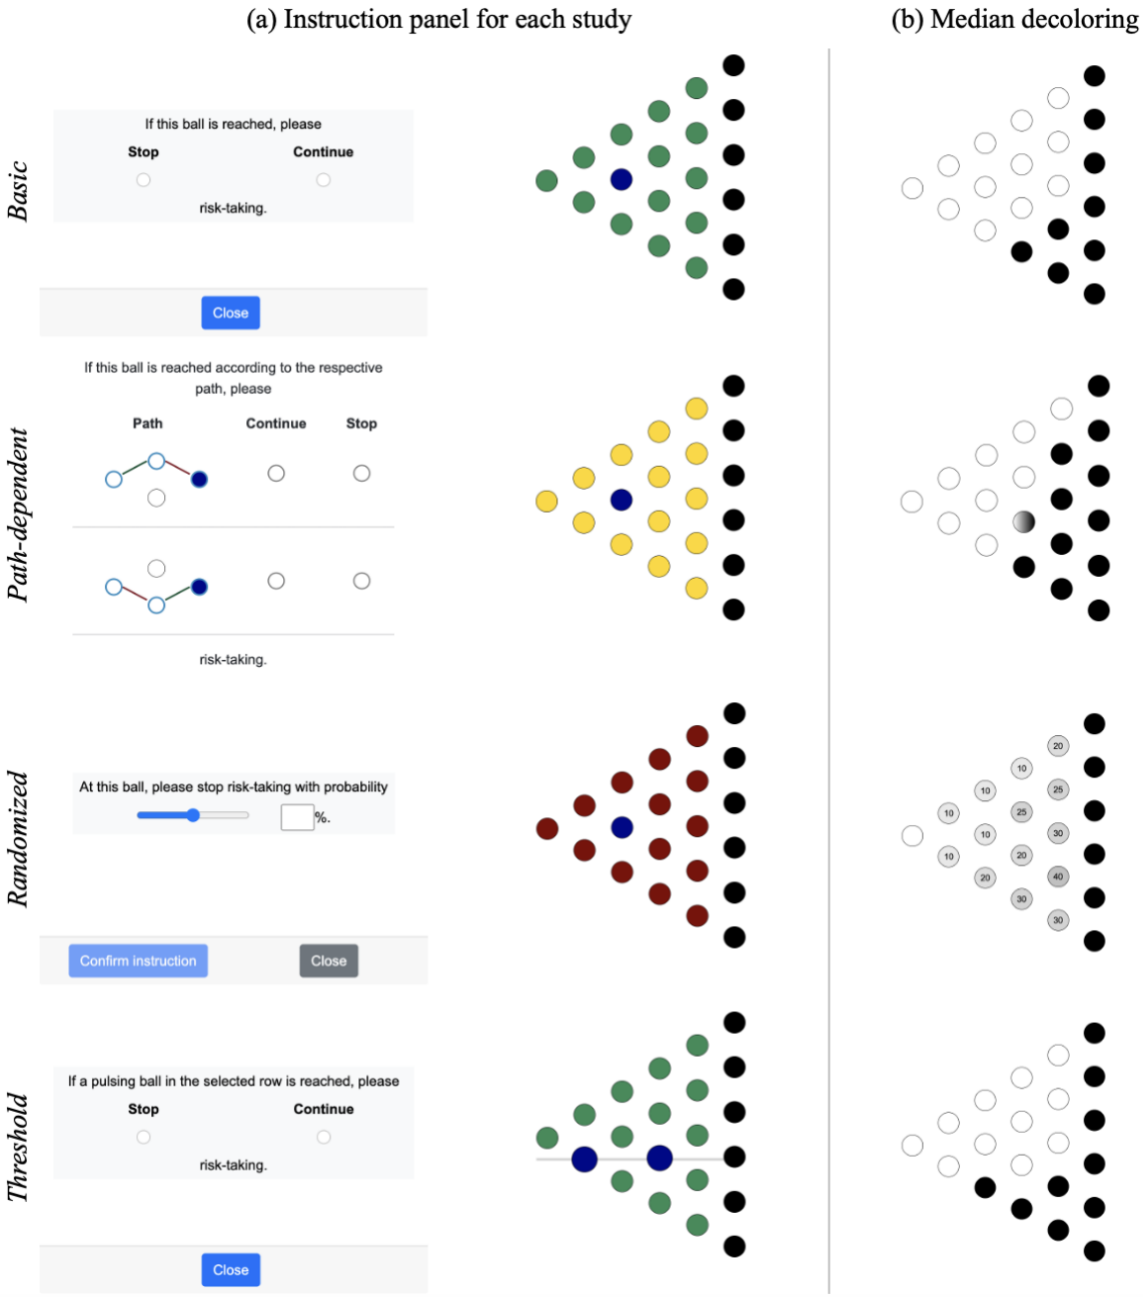
The decoloring scheme is as follows: If a ball is white, the subject takes the 50:50 risk at that ball. If a ball is black, the subject stops risk-taking. Balls with a black-white gradient indicate that stopping at that ball is path-dependent in that the subject continues risk-taking for at least one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

Due to the freedom we give subjects in the various decoloring tasks, the elicited stopping times are not necessarily meaningful. 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). In the analysis, we pool subjects from different studies with the same Decoloring Task.[10]

[9] 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.

[10] 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).

FIGURE 3.—Decision screens and median risk-taking strategies observed in the Decoloring Task



Notes: Subjects in a *Basic (Path-dependent, Randomized, Threshold)* Decoloring Task must decolor all 15 balls of an initially green (yellow, red, green) ball triangle. Subjects click a ball in the right half of the instruction panel (a) and give the risk-taking instruction for that (selection of) ball(s) in the left half of the instruction panel (a). The 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. In the *Threshold* task, colored balls in the same row must be decolored the same (to white or black).

Subjects can go back and change their decoloring. Panel (b) illustrates the median decolorings, whereby the median is taken over the effective decoloring at each ball.

III.B Basic Decoloring Tasks

Panel (b) of Figure 3 shows the aggregate precommitted risk-taking strategies for the *Basic*, *Path-dependent, Randomized*, and *Threshold* Decoloring Tasks. 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.

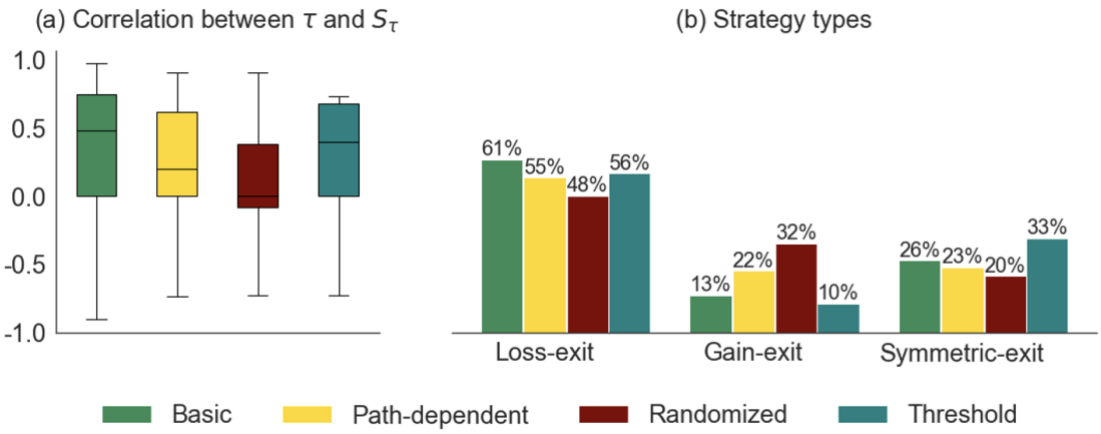
The median decoloring indicates risk-taking (white first ball), continue-when-winning (the upper part of the triangle being white), and stop-when-losing (the lower black balls). 96% of subjects take the risk at least once.[11] Moreover, the median strategy is of the loss-exit type defined by Barberis (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 specifically, the median loss-exit strategy is a level-3 *trailing stop-loss.* The stopping barrier is increasing in the repetition of the risk (trailing), and stopping occurs after losing thrice (level-3), when ignoring that there is a final period. 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. We find that 58% of the subjects decolor the bottom ball in the third column white, which is significantly different from 50% ($t(312) = 2.68$).

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 the induced outcome distribution S_τ for the different Decoloring Tasks. A positive correlation indicates "continuing-when-winning" (or, likewise, stopping when losing; i.e., a stop-loss strategy). For the *Basic* Decoloring Tasks, we find a positive average correlation ($\rho = 0.35, t(312) = 13.71$), which is non-negative for 75% of all risk-taking strategies. The average expected stopping time conditional on stopping with a gain is

[11] The large risk appetite in dynamic settings is in line with that of other experiments. Heimer et al. (2025), for example, report that around 95% of subjects start risk-taking.

4.41 and significantly higher than the expected stopping time conditional on stopping with a loss $(3.54, \mathrm{t}(312) = 12.54)$. Panel (b) of Figure 4 shows that a majority of 58% of the subjects in the *Basic* Decoloring Tasks plan to use a loss-exit strategy. Heimer et al. (2025) also observe a majority of loss-exit plans when subjects choose among threshold strategies. We find that loss-exit strategies even emerge endogenously when allowing for unconstrained risk-taking strategies, and that many of them are of the trailing type (see also Section III.E).

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 strategies that can be classified as loss-exit, gain-exit, or symmetric-exit based on the expected stopping time conditional on stopping with a gain or loss.

Subjects in the *Basic* Decoloring Tasks take the risk often: on average 3.82 out of at most 5 times. By the optimal stopping theorem, the expected earnings of any risk-taking strategy are zero, and thus, risk-averse expected utility preferences are incompatible with the results of our experiment. The induced outcome distribution S_τ is, on average, positively skewed $(\text{skew} = 0.36, t(312) = 12.24)$, in line with Ebert and Köster's (2024) result that relates a positive ρ to 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. We explore this relationship further in Section III.F.

III.C 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 analyze if these more general decoloring options are used by subjects and how they affect the resulting risk-taking strategies.

We find that 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. 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 for achieving randomized stopping decisions (we elaborate on this in Section V.E). Consistently, we find that 81% of subjects in *Randomization* use randomization at least once, 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*. Overall, 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),

[12] The frequencies of path-dependent balls are statistically identical ($\chi^2(6) = 3.898$) in *Path-dependence* and *Path-dependence*^{C}, while the frequencies of randomized balls differ at the 10%-level ($\chi^2(15) = 22.487$) between *Randomization* and *Randomization*^{C}.

with a higher share of gain-exit strategies than in *Basic*. Path-dependent strategies are behaviorally similar to *Basic* risk-taking strategies.

In contrast, risk-taking is more strongly affected by allowing for randomization. The third 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 relatively high popularity of gain-exit strategies in *Randomization*.

Overall, we find that (a) most subjects start taking a fair and symmetric single risk repeatedly, regardless of their available decoloring options, and that (b) aggregate risk-taking strategies resemble stop-loss strategies that induce a positively skewed outcome distribution S_τ. *Basic* and *Path-dependent* risk-taking strategies are behaviorally very similar, whereas *Randomization* results in a weaker stop-loss pattern.

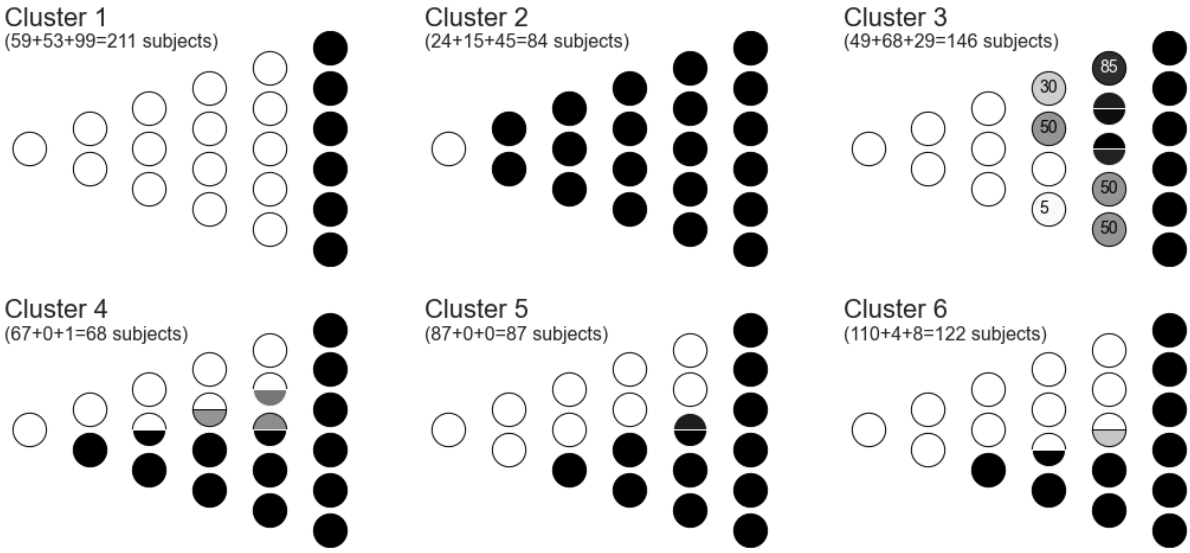
III.D Individual-level analysis of strategy types using unsupervised clustering

We now turn to an analysis of the individual decolorings 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. We find that the clusters suggested by the K-means algorithm have a natural economic interpretation in terms of stop-loss, buy-and-hold, or take-profit strategies. Most individual risk-taking strategies are classified as stop-loss and buy-and-hold strategies. Figure 5 plots the median decoloring over all

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

strategies assigned to a cluster for $K = 6$ clusters.[14] In addition, from left to right, we show the number of loss-exit, gain-exit, and symmetric-exit strategies as well as the total number of strategies in each cluster. Appendix SA.3.1 contains methodological details and robustness results.

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).

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 of cluster 1 is a long-horizon buy-and-hold strategy, that of cluster 2 is taking the risk once. The median strategy of cluster 3 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

[14] Intuitively, a K-means algorithm partitions data into $K \in \mathbb{N}^+$ clusters that have a high within-cluster similarity and are separated from other clusters. We conduct the analysis here with K=6 clusters and use all N=718 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.

contains more gain-exit than loss-exit strategies (68 vs. 49). Thus, we interpret cluster 3 as a take-profit cluster. The median risk-taking strategies of clusters 4, 5, and 6 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.

In summary, we find that a relative majority of all $N = 718$ risk-taking strategies elicited in our experiment is assigned to clusters representing stop-loss strategies ($N = 278$), followed by strategies in buy-and-hold cluster 1 ($N = 211$). The general willingness to take risks in dynamic contexts is evident from the low number of strategies in cluster 2. An unsupervised machine learning algorithm thus 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 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.E Differentiating between threshold and trailing stop-loss strategies

In the previous section, we observed that a large fraction of unconstrained risk-taking strategies is assigned to clusters interpreted as stop-loss strategies. Eliciting unconstrained risk-taking strategies allows us to inquire what *kind* of stop-loss strategies subjects choose and whether threshold or trailing stop-loss strategies are more prevalent. We conduct two analyses. First, we analyze whether a restriction to threshold strategies affects aggregate risk-taking by comparing the unconstrained strategies elicited in the *Basic* Decoloring Tasks to the constrained strategies elicited in *Threshold*. Second, we compare all unconstrained strategies to a comprehensive set of labelled template strategies using matching as well as supervised machine-learning procedures. The first analysis shows that aggregate risk-taking is not strongly affected by restricting subjects' strategy set to threshold strategies, suggesting that threshold strategies might be a good approximation for a wider range of risk-taking strategies in the aggregate. However, as

the median strategies in the three stop-loss clusters in Figure 5 suggest, we observe more trailing than threshold stop-loss strategies at the individual level, with a ratio of trailing to threshold strategies of approximately 1.5:1.

Study *Threshold.* We first compare aggregate risk-taking strategies from the *Threshold* Decoloring Task to aggregate risk-taking strategies from the *Basic* Decoloring Tasks. The bottom row of panel (b) in Figure 3 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 test if the visual difference is statistically significant by comparing the share of subjects stopping at each ball in *Threshold* and *Basic* and find that the share of subjects that stop risk-taking at the bottom ball in the third column is significantly higher in *Threshold* than in *Basic* ($t(389) = 2.22$). The interpretation of the aggregate risk-taking strategy changes from a level-3 trailing stop-loss in *Basic* to a level-2 threshold stop-loss in *Threshold*.

TABLE 2.—Classification of the unconstrained individual risk-taking strategies for each cluster

	Count-based classification							
Cluster	Threshold stop-loss	Trailing stop-loss	Buy-and-hold	Never-start	Threshold take-profit	Trailing take-profit	Not identified	Total
1	18	4	97	-	5	12	43	179
2	-	8	30	25	-	2	11	76
3	4	2	38	-	12	20	67	143
4	26	18	3	-	-	-	8	55
5	3	70	4	-	-	-	10	87
6	44	46	-	-	-	1	9	100
Total	95	148	172	25	17	35	148	640

Notes: Table 2 shows the assignment of strategies to 20 template strategies based on the lowest number of paths for which the elicited risk-taking strategy deviates from each template strategy. We do not assign an elicited risk-taking strategy to a class if the most similar class is not unique, and we exclude strategies elicited in *Threshold*.

Comparing the types of risk-taking strategies in the *Threshold* Decoloring Task to those in the *Basic* Decoloring Tasks (see panel (b) of Figure 4) suggests that fewer subjects

use loss-exit (56% against 61%) and gain-exit (10% against 13%) strategies in *Threshold* than in *Basic*. We observe more symmetric-exit strategies in *Threshold* than in *Basic* (33% against 26%), but 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 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* ($\text{Skew}(\tau) = -0.32$, $t(389) = 3.78$). Overall, aggregate risk-taking measures are not strongly affected by restricting subjects to using threshold risk-taking strategies, although the interpretation of the aggregate risk-taking strategy changes from a trailing to a threshold stop-loss strategy.

Classification. Next, we analyze whether individual risk-taking strategies are more often threshold or trailing stop-loss strategies using different classification procedures. First, we define all *Basic* (neither randomized nor path-dependent) threshold stop-loss, trailing stop-loss, buy-and-hold, never-start, threshold take-profit and trailing take-profit strategies as templates. For example, a level-3 trailing stop-loss strategy (as in the median *Basic* decoloring in Figure 3) is one template, as is a level-2 threshold stop-loss strategy. Figure SA.3.5 in the appendix shows all 20 templates. We then compare all elicited risk-taking strategies except those from *Threshold* to the template strategies using different procedures, including a count-based matching, cosine-similarity, a logistic regression, a decision tree as well as a random forest classifier (for details, see Appendix SA.3.2). In the main text, we focus on the count-based matching procedure and measure similarity by counting the number of paths for which the elicited strategy deviates from each template. A risk-taking strategy is more similar to a template if it deviates from it for fewer paths and is assigned to the class of its most similar template.

Table 2 shows the assignment of the $N = 640$ elicited unconstrained risk-taking strategies (excluding *Threshold*) to the templates. We separately show the assignment for each cluster that the K-means algorithm identified in Section III.D to corroborate our interpretation of the clusters. The bottom line of Table 2 shows that most elicited risk-taking strategies are buy-and-hold strategies (35% of all successfully classified strategies), followed by trailing stop-loss strategies (30%). The share of trailing stop-loss strategies is substantially higher than the share of threshold stop-loss strategies (19%), and we observe 1.56 trailing stop-loss strategies for each threshold stop-loss. In

addition, we find that (trailing) take-profit strategies are far less popular than (trailing) stop-loss strategies, and that the trailing version of take-profit strategies is more prevalent than threshold take-profit strategies. Overall, we observe a large prevalence of trailing stop-loss strategies, both absolutely as well as relative to threshold stop-loss.

III.F Structural estimation of CPT from stopping time data

In the previous sections, we observed that (a) most subjects start risk-taking although the symmetric risk has an expected value of zero, and that (b) most subjects' risk-taking strategies induce positively skewed outcome distributions S_τ. Barberis (2012) shows that CPT (Tversky and Kahneman 1992) can explain such risk-taking behavior. We thus propose and implement an approach to structurally estimate CPT parameters from stopping time data. The estimates may inform theoretical work on the kind of chosen risk-taking strategies that are (not) consistent with CPT. We also identify types of CPT agents from stopping time data. Compared to the non-structural cluster analysis (Section III.D), the analysis speaks to the restrictiveness of CPT's parametric assumptions.

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

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

We assume the probability weighting function as in Tversky and Kahneman (1992)

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

and an exponential value function (Köbberling and Wakker 2005)[15]

[15] As is well known (e.g., Köbberling and Wakker 2005), the exponential value function is econometrically more robust than the commonly used power specification, which confounds the loss aversion and diminishing sensitivity parameters.

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

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. 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 2025). Intuitively, an agent with a low δ overweights the tails of the distribution, and high α implies high marginal utility of large gains. Similarly, 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.

We want to know which CPT parameters rationalize our elicited stopping times and focus on *Basic* stopping times in the estimation.[16] For each of the $j = 1,2,\dots,313$ subjects that completed the *Basic* Decoloring Tasks, we observe one chosen stopping time τ_c^j which induces an outcome distribution (or lottery) $S_{\tau_c}^j$. 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 $\tau_\neg \in \mathrm{T}_{Basic}$, 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 $\tau_\neg$. This 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 $\tau_\neg$ for a given parameter vector $\theta_j = (\alpha_j, \delta_j, \lambda_j, \epsilon_j)$ with error parameter $\epsilon_j \in (0,100]$:

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

$$f(\tau_c|\tau_\neg, \theta_j) = \frac{1}{1 + e^{-\epsilon_j\left(V(S_{\tau_c}|\,\theta_j) - V(S_{\tau_\neg}|\,\theta_j)\right)}}.$$

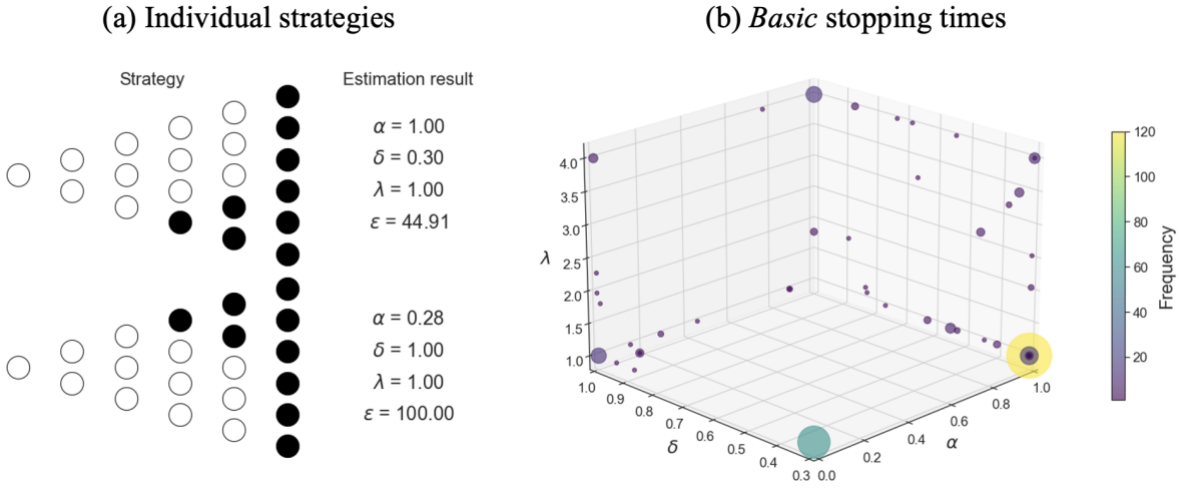
We estimate θ_j for each subject by maximizing the log-likelihood of observing τ_c^j in T_{Basic}, as

$$\theta_j^* = \arg\max_{\theta} \log \prod_{\tau_\neg \in \mathrm{T}_{Basic} \setminus \{\tau_c^j\}} f\left(\tau_c^j | \tau_\neg, \theta\right)$$

$$= \arg\min_{\theta} \sum_{\tau_\neg \in \mathrm{T}_{Basic} \setminus \{\tau_c^j\}} \log\left(1 + e^{-\epsilon_i\left(V\left(S_{\tau_c^j}|\theta\right) - V\left(S_{\tau_\neg}|\theta\right)\right)}\right).$$

Figure 6 shows the results of our maximum-likelihood estimation. To get a first impression, Panel (a) of Figure 6 shows the estimation results for two selected strategies. Our maximum-likelihood estimation finds a high α, low δ, and low λ for loss-exit strategies that induce positively skewed outcome distributions. 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 $(\alpha, \delta, \lambda, \epsilon) = (1.00, 0.30, 1.00, 44.91)$. In contrast, the estimates for a level-3 trailing take-profit strategy as in the bottom row of Panel (a) are $(\alpha, \delta, \lambda, \epsilon) = (0.28, 1.00, 1.00, 100.00)$.

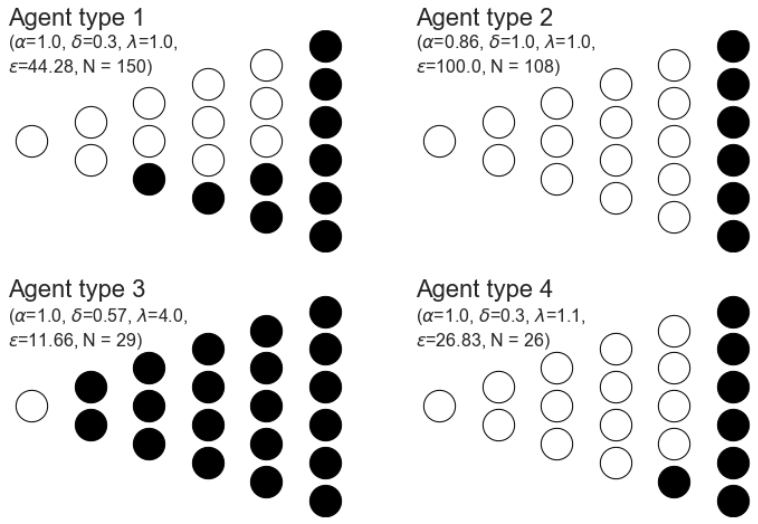
FIGURE 6.—Individual-level CPT parameter estimates



Notes: Panel (a) of Figure 6 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.

Panel (b) of Figure 6 shows the location of all 313 individual-level parameter estimates in the CPT parameter space. 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. Compared with the template strategies, these parameter combinations are obtained for long-duration buy-and-hold strategies, of which we also observe many. Otherwise, the parameter estimates are scattered over the cube and illustrate 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 6. 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.— CPT parameter estimates and median strategies from a finite mixture model



Notes: Figure 7 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.

Finite mixture model. Given the apparent clustering of parameter combinations in Figure 6, we now analyze if a smaller set of distinct 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. Appendix SA.4 discusses methodological details and shows robustness analyses for up to seven agent types. We focus on four agent types here.

Figure 7 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.[17] For agent type 1, we estimate a linear value function ($\alpha = 1.00$), strong probability weighting ($\delta = 0.30$), and no loss aversion ($\lambda = 1$). 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 ($\text{Skew}(S_\tau) = -0.08$, $t(107) = -3.688$) and to $\alpha < 1$.

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.

[17] 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.

Overall, the structural assumption of CPT preferences sorts 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—as a theory of choice under risk—can rationalize most of our elicited risk-taking strategies and is flexible enough to a be a good approximation to the non-parametric clusters identified using machine-learning.

III.G Decoloring process

We now analyze the process by which subjects decolor their ball triangle to gain insights into their approach to finding their desired stopping time. In our setting—with a finite horizon of $T = 5$ single 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—during the decoloring task.[18]

We recorded the decoloring process for $N = 240$ subjects who completed the *Basic* Decoloring Task in our 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, we find that subjects decolor the ball triangle from left to right (forward induction) and from top to bottom, meaning that subjects decolor gain balls before loss balls (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 the tendency to decolor gain balls before loss balls is evidenced by a negative average correlation between the accumulated earnings at a ball and the click order ($\rho_g = -0.09, t(239) = -3.526$). In the aggregate, we thus find evidence of forward induction and a tendency to decolor gain balls before loss balls.

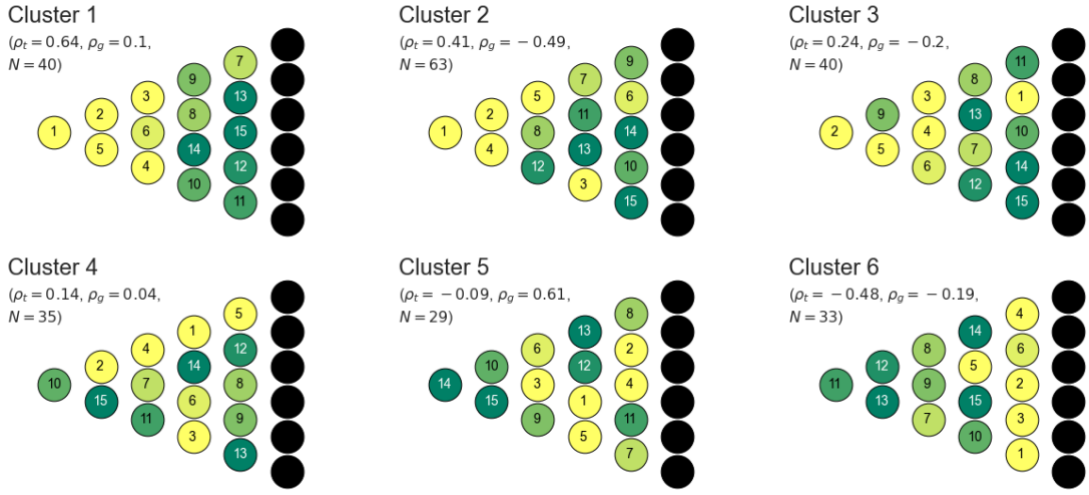
We analyze the heterogeneity in the decoloring process using the same unsupervised machine-learning algorithm as in Section III.D. We conduct the analysis with $K = 6$ clusters, as suggested by model selection criteria.[19] Figure 8 shows, for each cluster, the decoloring process that best describes all decoloring processes assigned to the cluster.

[18] We thank Christian Gollier and Fabio Maccheroni who independently suggested this analysis to us.

[19] The silhouette score has a local maximum, while the AIC and BIC both have a local minimum at $K = 6$.

Intuitively, we search for the click order that minimizes the total squared distance to the click order defining the respective cluster.[20] In addition, Figure 8 shows ρ_t and ρ_g as well as the number of decoloring processes for each cluster.

FIGURE 8.—Types of decoloring orders



Notes: Figure 8 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 the risk repetition (ρ_t) or the accumulated earnings (ρ_g) as well as the number of decoloring processes assigned to each cluster.

Decoloring processes assigned to Cluster 1 (17% of all decoloring processes) are characterized by forward induction, in that balls to the left have a lower order number than balls to the right, and a slight tendency to decolor loss balls before gain balls ($\rho_g > 0$). Cluster 2 (26%)—which is like cluster 3 (17%)—seems to be a forward induction cluster as well, but subjects in these clusters decolor gain balls before loss balls ($\rho_g < 0$). Decoloring processes in Cluster 4 (15%) have a weaker tendency to use forward induction and show almost no correlation between the earnings at a ball and the click order. Clusters 5 (12%) and 6 (14%) both show backward induction ($\rho_t < 0$) which is most pronounced in Cluster 6. While decoloring processes in Cluster 5 decolor loss balls earlier than gain balls ($\rho_g > 0$), the reverse holds for cluster 6 ($\rho_g < 0$).

[20] The Hungarian algorithm solves this problem in polynomial time.

Overall, the decoloring processes of 60% of subjects—those in clusters 1, 2, and 3—show a clear pattern of forward induction ($\rho_t > 0$), while about 26% of subjects (clusters 5 and 6) seem to use backward induction. Our findings thus support prior research which shows that people rarely use backward induction, even if normatively optimal. We present direct individual-level evidence from stopping times that complements previous research which 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).

IV. Dynamic consistency: Elicitation method and analysis

This section analyzes sequential risk-taking with a focus on dynamic consistency—that is, we analyze if subjects follow their risk-taking plan or systematically deviate from it. After having submitted a precommitted risk-taking strategy (in the Decoloring Task), subjects take the risk sequentially during ten Walking Tasks. We interpret subjects' decolorings as their respective plans that they may or may not decide to deviate from in the Walking Tasks.[21] We implement different types of the Walking Task to identify the causes of dynamic (in-)consistency.

IV.A Eliciting risk-taking actions with the Walking Tasks

In the Walking Tasks, subjects make sequential risk-taking decisions while visually "walking" through the ball triangle. Subjects sequentially instruct the computer about their risk-taking action at each ball that is reached, possibly deviating from their previously elicited plan. If a risk is taken, it is directly resolved. Each Walking Task comprises taking up to five risks[22] and each subject completes ten Walking Tasks.

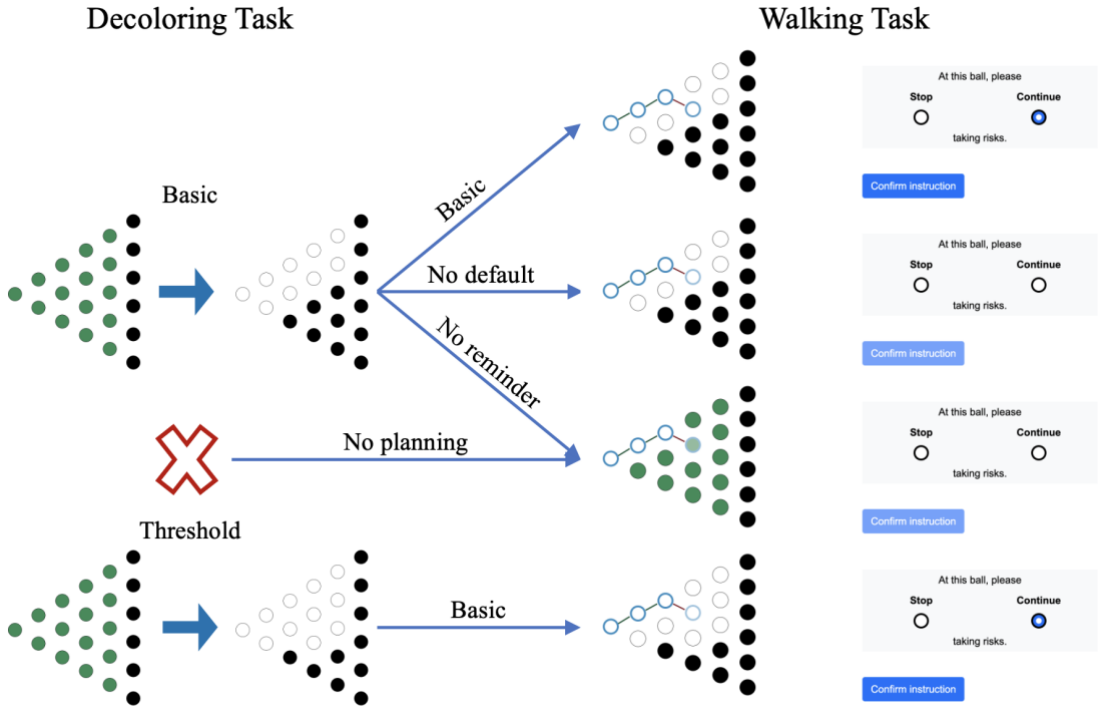
We implement different types of the Walking Task between-subject (see Table 1 and Figure 9). During a *Basic* Walking Task, subjects walk along their previously decolored ball triangle, which serves as a plan reminder, and the plan-consistent risk-taking action

[21] Pooling all subjects, we observe a comprehensive set of 37,452 risk-taking actions. For any node, we observe at least 268 actions, with a maximum of 8,810 actions (= 881 subjects ∗ 10 Walking Tasks) for the very first node in the ball triangle.

[22] Subjects must always resolve all five risks regardless of when risk-taking stopped to prevent finishing a task faster by stopping early. After stopping, subjects no longer observe the risk resolution to prevent regret.

is pre-selected as default. The *No default* Walking Task is like the *Basic* Walking Task, except that no risk-taking action is pre-selected as default. Subjects in a *No reminder* Walking Task see no plan reminder (and have no pre-selected default) when taking the risk sequentially but walk along a colored (green) ball triangle.[23] Moreover, subjects in *No planning* do not complete a *Decoloring Task* before taking the risk sequentially in a *No reminder* Walking Task.

FIGURE 9. — Dynamic consistency treatments



Notes: After submitting the decoloring in a *Basic* Decoloring Task, subjects in a *Basic* and *No Default* Walking Task walk through their decolored ball triangle. Subjects in a *No reminder* Walking Task walk through a green ball triangle. The four left panels of the Walking Task (right part) show exemplarily 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 respective decision screen on the right. In the *Basic* Walking Task, subjects' default is the plan-consistent risk-taking action. Subjects in a *No default* or *No reminder* Walking Task do not have a default risk-taking action and the "Confirm Instruction" button only becomes active after selecting "Stop" or "Continue." If a subject continues risk-taking, the single risk is resolved, the path through the ball triangle is

[23] We also implement a *Randomized* and *Randomized*^{C} Walking Tasks that is like the *Basic* Walking Task, except that subjects can (costly) stop risk-taking with any probability on the grid $0, 10, \ldots, 90, 100$. Results for the *Randomized* Walking Tasks are in the Appendix and unreported details are available from the authors. We do not implement a dedicated *Path-dependent* Walking Task (Footnote 7).

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. Subjects in the *No planning* study do not submit a risk-taking plan (i.e., start the experiment with the Walking Task), and their Walking Task is identical to the *No reminder* Walking Task. Subjects in the *Threshold* study complete *Basic* Walking Tasks.

IV.B Measuring and identifying the causes of dynamic consistency

We designed the Walking Tasks to test for the causes of dynamic (in-)consistency, that is, we analyze if subjects follow their risk-taking plan when taking the risk sequentially. We focus on two different measures of a subject's dynamic consistency. First, the (unconditional) *consistency measure* C_i is the share of plan-consistent risk-taking actions for each subject i:

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

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. C_i is a within-subject measure of dynamic consistency that compares a subject's risk-taking strategy to the same subject's risk-taking actions.[24]

Second, we construct a *disposition measure* D_i, which is inspired by the measures in Odean (1998) and Fischbacher, Hoffmann, and Schudy (2017). The disposition effect describes the tendency of investors to sell winning assets too soon and hold losing assets too long. Our disposition measure D_i consequently compares if subjects stop risk-taking earlier in the gain- or loss-domain. For each subject i, we define

$$\text{Proportion of Losses Stopped } (\text{PLS}_\text{i}) = \frac{\text{Losses stopped}_\text{i}}{\text{Losses reached}_\text{i}},$$

$$\text{Proportion of Gains Stopped } (\text{PGS}_\text{i}) = \frac{\text{Gains stopped}_\text{i}}{\text{Gains reached}_\text{i}},$$

and $D_i = PLS_i - PGS_i$. $D_i > 0$ implies that a subject's proportion of losses stopped is larger than the proportion of gains stopped, which indicates loss-exit risk-taking. Conversely, $D_i < 0$ indicates gain-exit risk-taking. We are interested in whether D_i

[24] 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, computing C_i for each sampled plan, and then taking the average.

measured for the Decoloring and Walking Tasks differs for the various studies we implemented. D_i is directly implementable for risk-taking actions observed during a Walking Task; and we compute the D_i for Decoloring Tasks using the realized paths through the ball triangle from the Walking Tasks (in which subjects must always resolve five risks) but stop risk-taking as if the subject follows her risk-taking strategy. D_i measures behavioral shifts in the type of planned and actual risk-taking.

We investigate how both measures vary across studies to test four factors that influence dynamic consistency. First, the presence of a default risk-taking action may affect consistency. A default makes a plan-deviation salient, and following it is marginally faster in our experiment as it saves a cursor move and one click. Second, the availability of a plan reminder may influence dynamic consistency. Subjects in the *Basic* Walking Task walk along their previously decolored ball triangle, which constitutes a plan reminder. Removing the plan reminder allows us to test the effect of memory on plan-consistent risk-taking. Third, making an explicit plan may positively affect consistency, which is why dynamic consistency is sometimes assessed between-subject. We test the effect of planning on sequential risk-taking by having some subjects complete the Walking Tasks without a previous Decoloring Task. Finally, the flexibility of the strategy space may lead to mechanical dynamic inconsistency, and we can test this because we implement different types of the Decoloring Task that vary in the restrictiveness of permissible stopping times in the following decreasing order: *Randomized*, *Path-Dependent*, *Basic*, and *Threshold* (Hu, Obłój, and Zhou 2022).

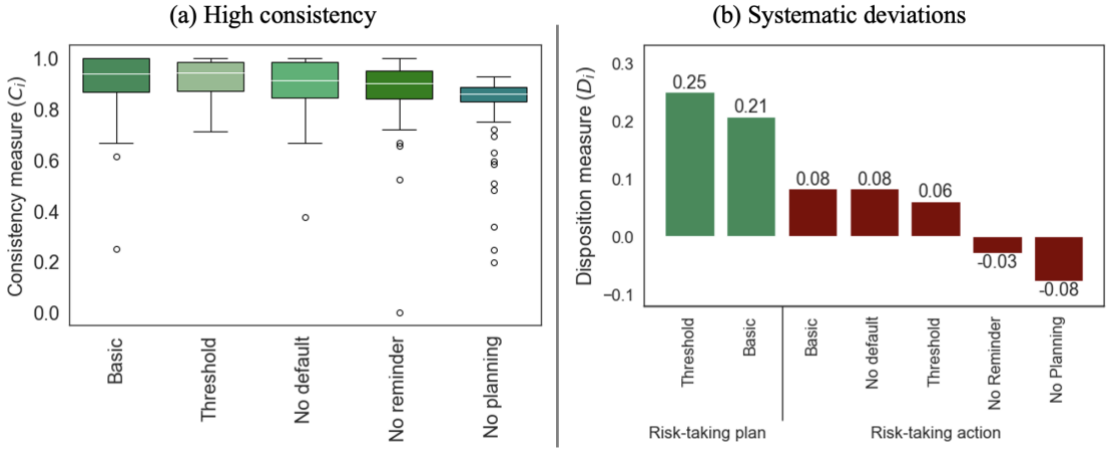
IV.C Causes of dynamic inconsistency

Plan consistency. We find that subjects mostly follow their previously elicited risk-taking plan when taking the risk sequentially. Panel (a) of Figure 10 shows the distribution of the consistency measure C_i for the different experimental studies. On average, 86.9% (median 89.1%) of the risk-taking actions are consistent with the risk-taking plan, and we find a high consistency C_i across experimental studies. The average consistency C_i varies from 83.7% for subjects in *No planning* to 91.6% for subjects in *Basic* (see Table SA.2.2).

The consistency C_i may appear high when having in mind studies like Heimer et al. (2025), who find that a majority of subjects plan stop-loss risk-taking but end up doing

gain-exit. The reason is that we elicit subjects' risk-taking actions at all balls—including the balls at which subjects plan to continue. A subject planning with a low stop-loss level and taking the risk sequentially until reaching a high gain will be consistent for most balls—all balls lying in-between those loss- and gain-levels.[25] Therefore, a substantial behavioral change from a stop-loss plan to gain-exit behavior can be in line with high overall consistency as measured by C_i. The behavioral change results from a few but systematic inconsistencies (at large-loss and large-gain balls).

FIGURE 10.—Average dynamic consistency measures



Notes: Panel (a) of Figure 10 shows the distribution of the consistency measure C_i for the different experimental studies, whereby C_i measures the share of risk-taking actions that follow the risk-taking plan. Panel (b) of Figure 10 plots the average disposition measure D_i according to each subject's risk-taking strategy (green bars) or actions (red bars), whereby D_i measures the difference between the share of stopped losses and gains.

Comparing the consistency C_i across studies, we find statistically significant differences between subjects in a *Basic* study and those in a *No reminder* ($t(226) = 2.51$) or *No Planning* ($t(308) = 6.74$) study.[26] The share of risk-taking actions that follow the risk-taking plan does not vary significantly when removing the default or when restricting

[25] To further assess the level of the consistency measure 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).

[26] If we apply a Bonferroni correction to adjust the significance threshold for multiple testing, the difference between *Basic* and *No reminder* is no longer significant.

the risk-taking strategies to threshold strategies. Subjects instead are significantly less dynamically consistent if not reminded of their plan or if they do not plan as such.

Standard shift. We now analyze if sequential risk-taking is (more) gain-exit than the risk-taking plans. Panel (b) of Figure 10 shows the average disposition measure D_i for the different types of Decoloring and Walking Tasks. First, the positive average D_i for *Basic* and *Threshold* risk-taking plans supports our previous analysis of the Decoloring Task: Subjects plan to use loss-exit strategies and there are no significant differences between D_i in the *Threshold* or *Basic* Decoloring Tasks ($t(372) = 0.95$).[27]

Removing subjects' precommitment—as in the Walking Tasks—leads to less loss-exit risk-taking than planned as shown by the red bars being lower than the green bars in Panel (b) of Figure 10. For subjects in the *Basic* Walking Tasks—who did the *Basic* Decoloring Task—we find an average D_i of 0.08, which is significantly lower than the planned D_i ($t(442) = 4.31$). Such a *standard shift in risk-taking* from loss-exit plans toward gain-exit risk-taking has been identified in previous theoretical and experimental settings (Barberis 2012; Heimer et al. 2025).

Causes of the standard shift. We now analyze drivers of the standard shift's magnitude and causes of dynamic inconsistency in risk-taking more broadly. First, removing the default risk-taking action in the *No Default* treatment does not affect the disposition measure ($t(224) = -0.01$). Second, sequential risk-taking is systematically affected by a strategy reminder. The average D_i for *No Reminder* is -0.03 and significantly smaller than for the *Basic* Walking Tasks ($t(221) = 3.575$). Without the strategy reminder, we obtain a (mildly and non-significantly) negative average D_i, that is, a shift from loss-exit plans to gain-exit risk-taking. Third, comparing D_i between the *No Reminder* and *No Planning* risk-taking actions allows us to assess the additional effect of planning as such on sequential risk-taking. We find a statistically marginally significant difference between the D_i in *No Reminder* and *No Planning* of -0.05 ($t(238) = 1.752$).[28]

[27] We pool all $N = 313$ subjects that complete the *Basic* Decoloring Task because (a) the Decoloring Task is completed before the Walking Task and cannot be affected by differences in the Walking Task, and (b) we do not find significant differences in the measures for subjects in different studies with the same Decoloring Task.

[28] We also analyze changes in the skewness of the outcome distribution (Figure SA.1.4). The results are similar to the disposition measure D_i, but noisier because the skewness of each subject's outcome

Mechanical inconsistency. Next, we analyze if restricting the strategy space leads to a mechanical dynamic inconsistency. The difference in the average planned and actual D_i in *Basic* of 0.13 serves as our baseline dynamic inconsistency. The same comparison finds a lower within-study shift in risk-taking in *Randomization* (difference of 0.04) and *Path-dependence* (0.10), but a larger shift in *Threshold* (0.19). The order of the differences aligns with the order of the flexibility of the strategy space. That is, studies with a more constrained strategy space show greater differences in planned and actual consistency, in line with the hypothesis of mechanical inconsistency. Statistically, however, the effects are insignificant.[29]

Overall, we find that planned risk-taking is more loss-exit than actual risk-taking—that is, we find the standard shift in risk-taking. Our analysis extends previous work by showing that the standard shift emerges even if subjects submit complete contingent plans and have a plan reminder as well as a plan-consistent default risk-taking action. We document that planning as such, and a plan reminder causally reduce the magnitude of the shift in risk-taking, while the presence of a default has no additional effect.

V. Discussion

V.A Experimental design

We introduce and implement the Decoloring Task as a novel method for eliciting unconstrained stopping times in repeated risk-taking. In this first study, to focus on the rich dynamic aspects of behavior, all risk-taking is about the simple 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 by implementing the simplest risk possible. One could adapt the Decoloring Task to study non-fair risks (Antler and Arad 2023), skewed risks (Ebert and Köster 2024), a

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 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).

[29] Table SA.2.3 provides details. Compared to the baseline dynamic inconsistency from *Basic* studies, only the difference to the dynamic inconsistency in the *Randomization* study is statistically significant ($t(221) = 2.78$).

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 node of the binomial tree (Merton and Samuelson 1992) or to consumer search problems (Weitzman 1979).

We incentivize the elicited stopping times τ under precommitment. If the Decoloring Task determines the subject's payment, we draw and visualize a random path through the decolored ball triangle. Risk-taking stops according to the subjects' stopping time τ, such that subjects are exogenously committed to their stopping time τ. 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 the contingencies are shown in a randomized order). Eliciting subjects' planned actions in a pre-specified order, in contrast, 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 decoloring process, specifically regarding backward or forward induction.

V.B Out-of-sample replication of the Basic decoloring task

Study *Basic II* provides an out-of-sample replication of study *Basic I*, because both studies are identical except for taking part at different dates (see Table 1). All major results reported replicate (see also the earlier version of this paper). Specifically, when comparing study *Basic I* conducted Nov 16, 2022 with $N_I = 73$ subjects and study *Basic II* conducted on May 23, 2024 with $N_{II} = 74$ subjects, we find no statistically significant difference between the expected stopping times (3.74 vs. 3.85, $t(145) = 0.598$) 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.D only using the data from *Basic I* and *Basic II*, respectively, we find highly similar clusters with a relative majority of risk-taking strategies assigned to clusters interpreted as (trailing) stop-loss (49% vs. 55%) followed by buy-and-hold (23% vs. 26%). The only behavioral difference between both studies seems to be stronger preference for trailing compared to threshold stop-loss strategies in *Basic I* (2.1:1) than in *Basic II* (1.4:1). The variation from the overall mean of 1.56 ($N = 640$; Section III.E) can be explained by the supervised machine-learning algorithms being sensitive to the smaller sample sizes in studies *Basic I* and *Basic II*, respectively. The differences from the overall mean vary across the supervised machine-learning algorithms used.

Comparing our dynamic consistency measures between subjects that completed study *Basic I* and *Basic II*, we find no statistically significant differences between the consistency measure C_i (0.915 vs. 0.916, $t(145) = 0.063$), the disposition measure calculated on the risk-taking plans (0.19 vs. 0.23, $t(140) = 0.663$), and the disposition measure obtained for sequential risk-taking (0.07 vs. 0.10, $t(141) = 0.710$).

V.C Threshold or trailing stop-loss

We find a relative popularity of trailing over threshold stop-loss strategies of roughly 1.5:1. The preference for trailing stop-loss strategies (in the optimal stopping literature often called Azéma-Yor stopping times) may be important for brokers and policy-makers who wish to promote trading.[30] An important insight from our experiment is that subjects choose trailing stop-loss strategies endogenously—without any restriction on the admissible strategies or without any suggestion given to subjects. At the same time, we find that restricting subjects to choose among threshold strategies (study *Threshold*) does not affect aggregate risk-taking. Focusing on threshold stop-loss strategies in theoretical and experimental work thus seems to be a reasonable approximation to aggregate unconstrained stop-loss strategies. On an individual level, repeating the K-

[30] The brokerage Schwab notes that "[…] standard stop orders can help investors attempt either to limit losses or preserve profits, trailing stop orders offer the opportunity to do both with a single setup;" see https://www.schwab.com/learn/story/trailing-stop-orders-mastering-order-types, accessed on June 3, 2023. Ebert and Köster (2024) show that a trailing stop-loss induces a positively skewed outcome distribution S_τ regardless of the skewness of the single risk if it is taken sufficiently often.

means clustering and matching for *Threshold* strategies, we find more stop-loss strategies (around 60%) and slightly less buy-and-hold strategies (around 20%) than for *Basic* risk-taking strategies (around 50% stop-loss and 25% buy-and-hold). For the analysis of individual-level risk-taking strategies, it thus seems necessary to consider the heterogeneity in stopping strategies by, for example, allowing for trailing strategies.

V.D Buy-and-hold

We find it worth highlighting that, in an experiment that allows for unconstrained stopping strategies, 35% of the subjects choose simple buy-and-hold strategies. Extrapolating this finding to more complex settings might provide evidence on the type of trading strategies that people entertain in the financial market and highlight a popularity of buy-and-hold trading plans for stocks, options, or ETFs.

V.E Path-dependence

Consistent with theoretical predictions by He et al. (2017) and Hu, Obłój, and Zhou (2022), we find that more than 55% of subjects use path-dependence at least once, even if it is costly. Despite the prevalent usage of path-dependence, path-dependent risk-taking strategies are very similar to standard time- and outcome-contingent risk-taking strategies. In contrast to the theoretical motivation for path-dependence (Hu, Obłój, and Zhou 2022), the induced outcome distributions S_τ in the *Path-dependence* studies are not more positively skewed than in the *Basic* studies. Moreover, subjects use path-dependence neither to (a) systematically stop risk-taking at the highest earnings along a path and to continue risk-taking otherwise, as implied by regret preferences (Strack and Viefers 2021), nor to (b) systematically differentiate between downward or upward sloping paths (Rabin 2002; Grosshans and Zeisberger 2018). Relatedly, Magnani et al. (2022) find that path-dependent stochastic control strategies lead to less efficient portfolio choices than standard (path-independent) strategies.

Our evidence on the usage of path-dependence is more consistent with a revealed demand for randomization than with a direct preference for path-dependence. Path-dependence can be used as a randomization device because subjects design their risk-taking strategy before knowing which path will be realized, such that the outcome of the path-dependent risk-taking plan is ex-ante random (see He et al. 2017; Hu, Obłój, and

Zhou 2022). For example, assume that a subject 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 (or vice versa). The ex-ante probability of stopping at the path-dependent middle ball in the third column is then 50%. 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, a substantial fraction of subjects randomizes by varying static choices over 20 repetitions of the same choice situation. Subjects do not know in advance which of the 20 choices will be payoff relevant, such that varying the choice is equivalent to randomizing over the options.

V.F Randomization

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. More than 80% of subjects use randomization at least once, even if it is costly. The high and robust usage of randomized risk-taking strategies is in line with observations in static settings (Dwenger, Kübler, and Weizsäcker 2018; Agranov and Ortoleva 2022),[31] and provides an empirical justification for the usage of randomized stopping times in theoretical work that goes beyond a convexification of the optimization problem (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

[31] 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.

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 (as offered by many stock brokerages), buy-and-hold strategies, path-dependence, as well as randomization. Foremost, 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 also observed in the real world. Prospect theory is flexible enough to accommodate the individual-level heterogeneity.

We hope these results can guide the assumptions imposed in theoretical work or serve as empirical benchmarks for model predictions. Notably, we document that resorting to simple threshold stop-loss strategies approximates aggregate risk-taking behavior well, lending support to the many economic models that emphasize or predict these strategies. Our method can be extended to test stopping behavior in specific economic settings or be adapted for studying fundamental aspects of dynamic risk-taking with and without the possibility of precommitment. We use it to quantify how mechanical inconsistency, explicit planning, defaults, memory-aides, and commitment devices causally affect dynamic consistency—a fundamental property of intertemporal choice.

REFERENCES

Agranov, Marina, Paul J. Healy, and Kirby Nielsen, "Stable randomization," *Economic Journal*, 133 (2023), 2553–2579.

Agranov, Marina, and Pietro Ortoleva, "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.

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*, (2025).

Ebert, Sebastian, and Mats Köster, "On Taking a Skewed Risk More Than Once," *Working paper available at SSRN*, (2024).

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.

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*, (2025).

Gödker, Katrin, Terrance Odean, and Paul Smeets, "Disposed to be overconfident," *Working paper available at SSRN*, (2024).

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, 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," *Working paper available at SSRN*, (2024).

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.

Levitt, Steven D., John A. List, and Sally E. Sadoff, "Checkmate: Exploring Backward Induction among Chess Players," *American Economic Review*, 101 (2011), 975–990.

Machina, Mark J., "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.

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.

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 (Springer).

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.

ELICITING STOPPING TIMES: SUPPLEMENTAL APPENDIX

(FOR ONLINE PUBLICATION)

Sebastian Ebert* Maximilian Voigt†

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.

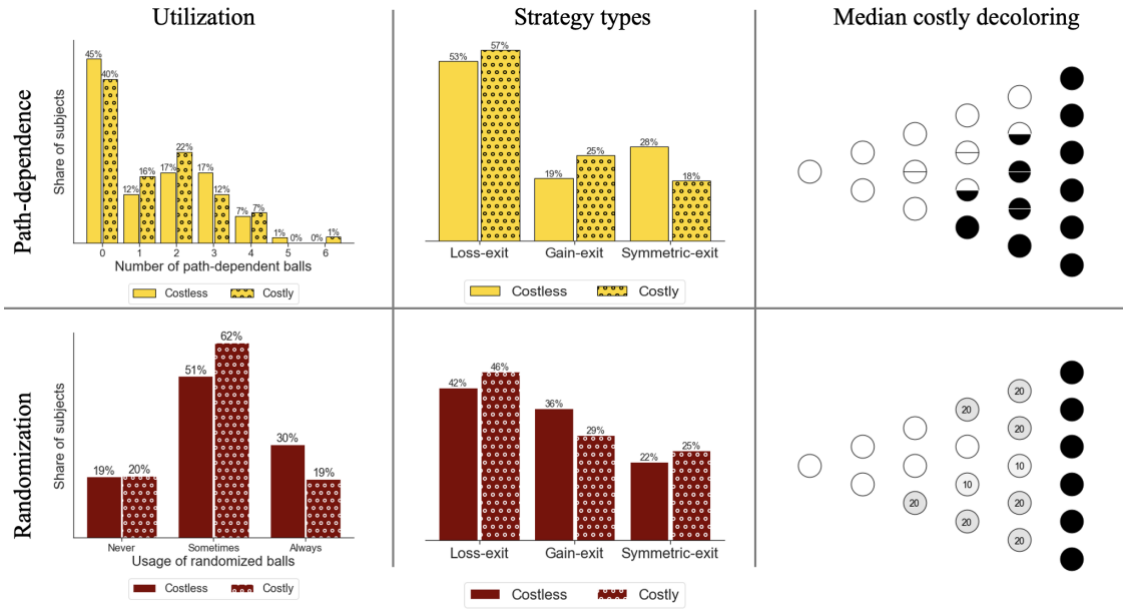
SA.1 Additional Figures

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.1.4 shows the skewness of the outcome distribution induced by both planned and actual risk-taking.

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

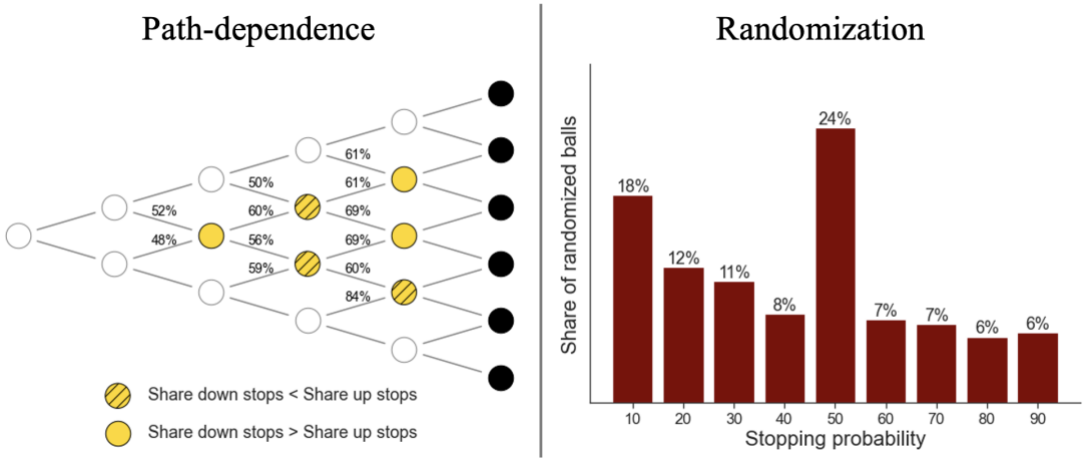
† HEC Montréal; Department of Finance. E-mail: maximilian.voigt@hec.ca.

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



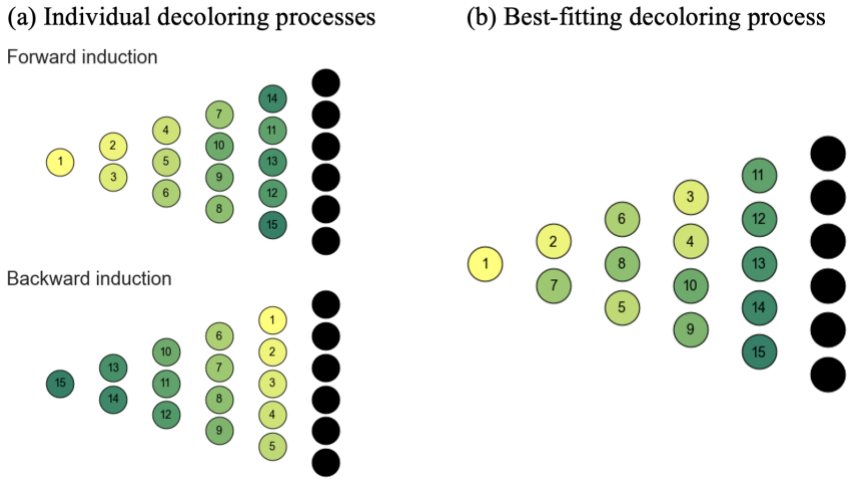
Notes: The upper left panel of Figure SA.1.1 shows the utilization of path-dependent balls when they are costless ($N = 75$) or costly ($N = 67$), while the bottom left panel shows the usage of costless ($N = 88$) or costly ($N = 97$) randomized balls. The middle panel shows the share of loss-exit, gain-exit, 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



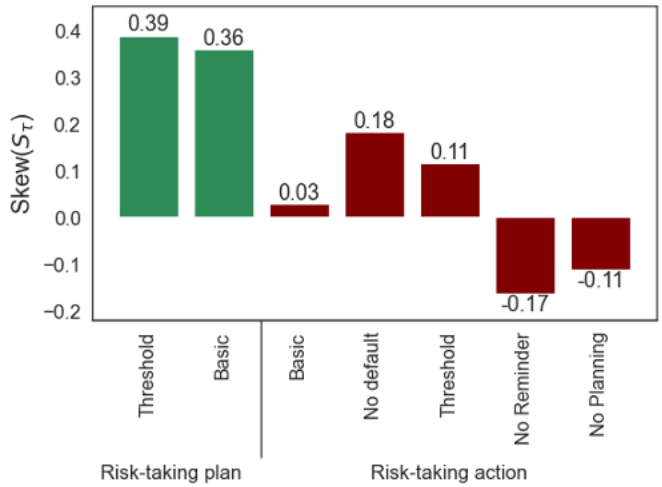
Notes: The left panel of Figure SA.1.2 shows the share of stopped risk-taking when a path is reached from above or below in the *Path-dependent* Decoloring Tasks ($N = 142$), whereas the right panel shows the share of balls with the stopping probabilities among all randomized balls for subjects in the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

FIGURE SA.1.4.—Skewness of the outcome distribution induced by planned and actual risk-taking



Notes: The green bars of Figure SA.1.4 show the average skewness of the outcome distribution induced by the risk-taking strategy, whereas the red bars show the skewness of the outcome distribution realized during the Walking Task.

SA.2 Additional Tables

Table SA.2.1 provides details on the six strategy clusters identified using the K-means algorithm in Section III.D. Table SA.2.2 informs about the within-subject 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	211	4.47	4.66	4.65	-0.01
2	84	1.25	1.52	1.28	0.17
3	146	3.14	3.38	3.72	-0.13
4	68	2.72	4.54	2.01	1.16
5	87	3.72	4.82	2.93	0.69
6	122	4.22	4.81	3.71	0.42

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.D, 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
Basic	0.9155	0.8997	0.9313
No default	0.8995	0.8781	0.9209
No reminder	0.8773	0.8485	0.9061
Threshold	0.9156	0.8979	0.9332
No planning	0.8366	0.8202	0.8530
Path-dependence	0.9032	0.8859	0.9206
Randomization	0.7884	0.7335	0.8432

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						
Threshold	0.2498	0.1581	0.3416	0.3883	0.2537	0.5229
Basic	0.2073	0.1715	0.2430	0.3593	0.3018	0.4168
Path-dependence	0.1285	0.0744	0.1826	0.2288	0.1482	0.3094
Randomization	0.0437	0.0068	0.0805	0.0513	-0.0233	0.1259
Risk-taking action						
Basic	0.0827	0.0399	0.1255	0.0280	-0.1098	0.1659
No Default	0.0831	0.0243	0.1420	0.1824	0.0470	0.3177
Threshold	0.0611	-0.0076	0.1298	0.1146	-0.0317	0.2609
Path-dependence	0.0269	-0.0242	0.0780	0.1293	0.0193	0.2392
Randomization	-0.0047	-0.0397	0.0303	-0.0276	-0.1153	0.0600
No reminder	-0.0289	-0.0719	0.0140	-0.1655	-0.2903	-0.0408
No planning	-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 dynamic consistency measured D_i and $\text{Skew}(S_\tau)$ for different studies and types of the task.

SA.3 Machine-learning methodology and robustness

This Supplemental Appendix complements Sections III.D and III.E. In the first part, we describe the unsupervised machine-learning procedure more closely, 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 III.E.

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

In Section III.D 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 re-performing the analysis of Section III.D with a different

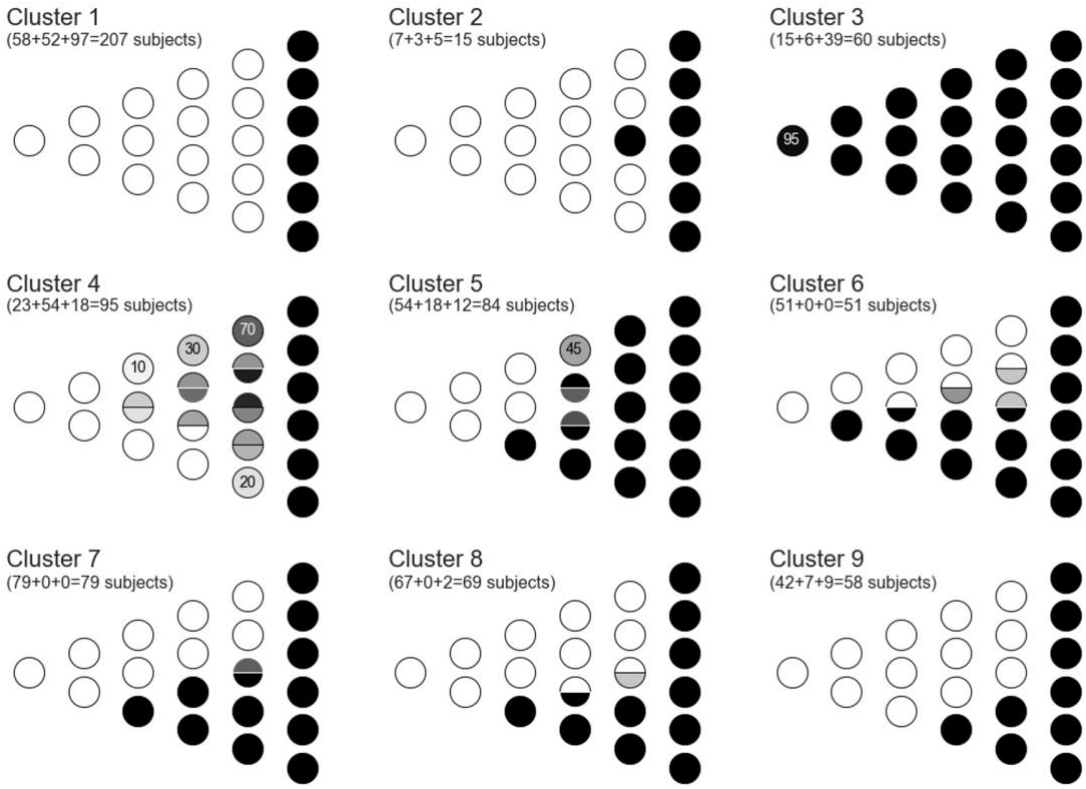
number of clusters[1] and separately for all risk-taking strategies elicited from *Basic*, *Path-dependent,* and *Randomized* Decoloring Tasks, including their costly variations. The findings highlighted in Section III.D—the separation of unconstrained risk-taking strategies into economically meaningful 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 does not learn an assignment of data to specific outcome variables of interest from labeled (training) data. Instead, the K-means clustering algorithm takes a set of unordered data points—718 precommitted risk-taking strategies that are represented as 31-tuples with the planned risk-taking action at each subpath of the ball triangle as entry—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.

[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. 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.

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.

The number of clusters K is a hyperparameter that needs to be chosen before performing the K-means clustering. Using standard methods[3] 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 showed the analysis in Section III.D of the main text for $K = 6$ clusters, because it is more parsimonious and leads to very similar results.

[3] A standard elbow plot of the *inertia* (sum of the squared distances of a datapoint from its cluster center) is inconclusive. While the plot might suggest 5 or 6 clusters, the inertia is declining very smoothly with more clusters. Additionally, we observe local minima of the clustering performance for 9 to 10 clusters in the *AIC* and *BIC* of a *Gaussian mixture* model. Moreover, the AIC and BIC of a Gaussian mixture model has a saddle point at six clusters, such that six clusters seem to provide a reasonable separation of the data. All methods except the *silhouette* method (difference between the average intra-cluster distance and the average distance to the nearest other cluster) suggest that using more clusters yields a better clustering performance.

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 stop-loss 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.

TABLE SA.3.1.—Average moments of the stopping time τ and induced outcome distribution S_τ as well as the Barberis (2012) classification for each cluster

Cluster	N	$\mathbb{E}(\tau)$	$\mathbb{E}(\tau\|\text{Gain})$	$\mathbb{E}(\tau\|\text{Loss})$	$\text{Skew}(S_\tau)$
1	207	4.49	4.67	4.66	0.00
2	15	4.40	4.78	4.69	0.02
3	60	0.76	1.10	0.68	0.29
4	95	2.87	3.15	3.76	-0.24
5	84	2.91	3.32	2.81	0.25
6	51	2.72	4.77	1.89	1.28
7	79	3.75	4.91	2.95	0.69
8	69	4.06	4.84	3.39	0.55
9	58	4.43	4.70	4.19	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 moments of the stopping time τ and induced outcome distribution S_τ for each cluster.

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 strategies (306 compared to 295) and less subjects in clusters that can be interpreted as take-profit or loss-exit strategies. In addition, we now find a never-start cluster, whereas cluster 2 in the main text could be interpreted as a compromise between a buy-and-hold and a never-start cluster. 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.

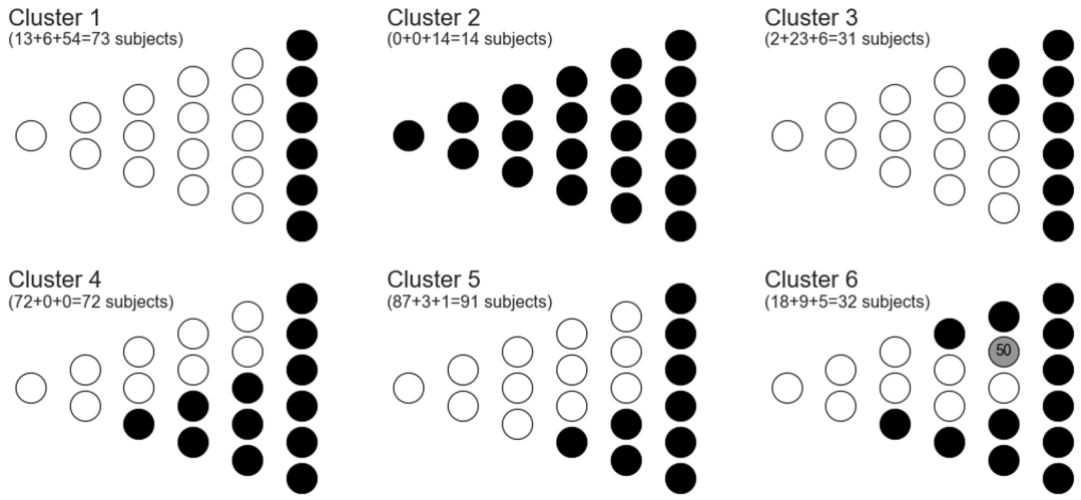
TABLE SA.3.2.—Assignment of clusters to economically meaningful strategy types

Strategy type	Clusters	N	Main text
Buy-and-hold	1, 2, 5	306	295
Never-start	3	60	-
Take-profit	4	95	146
Stop-loss	6, 7, 8, 9	257	277

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.

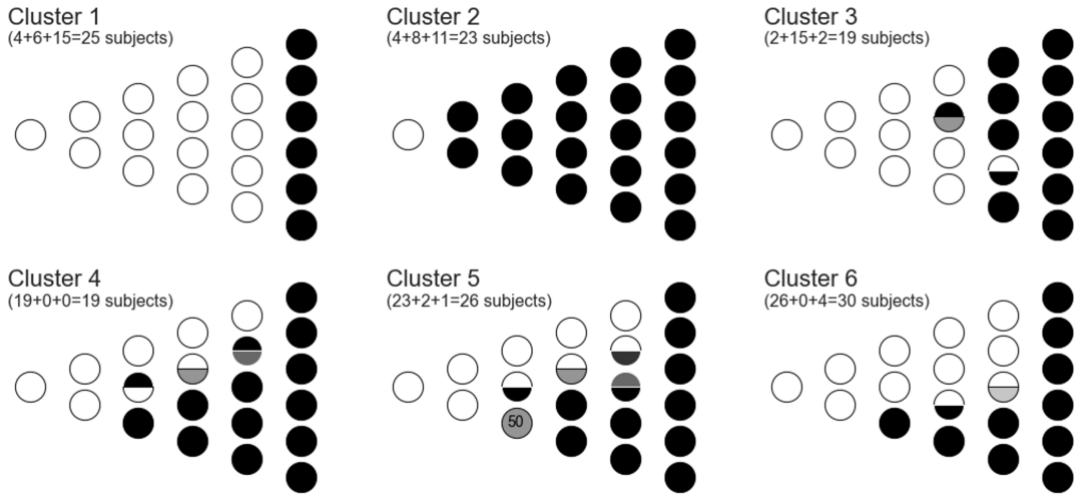
FIGURE SA.3.2.—Median risk-taking strategy per cluster (*Basic*)



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 *Basic* Decoloring Tasks.

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 take-the-risk-at-most-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. A relative majority of the *Basic* risk-taking strategies is assigned to clusters interpretable as (trailing) stop-loss strategies.

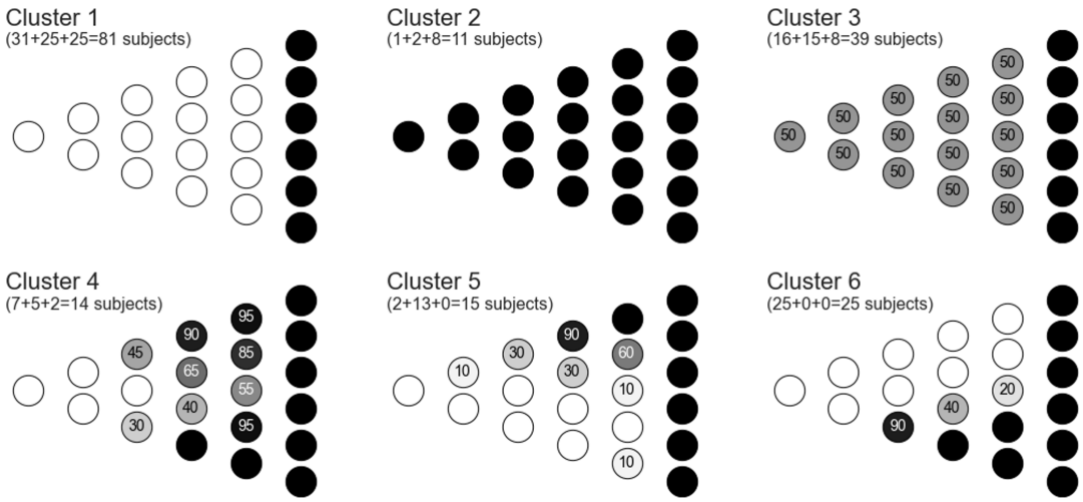
FIGURE SA.3.3.—Median risk-taking strategy per cluster (*Path-dependent*)



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 *Path-dependent* and *Path-dependent*^{C} Decoloring Tasks.

The median risk-taking strategy of cluster 3 among the *Path-dependent* risk-taking strategies 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.3.4.—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 *Randomized* and *Randomized*^{C} Decoloring Tasks.

The median risk-taking strategies for cluster 6 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, while those of the remaining clusters 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.C, 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

Here, we complement the analysis from Section III.E 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 another similarity metric (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:

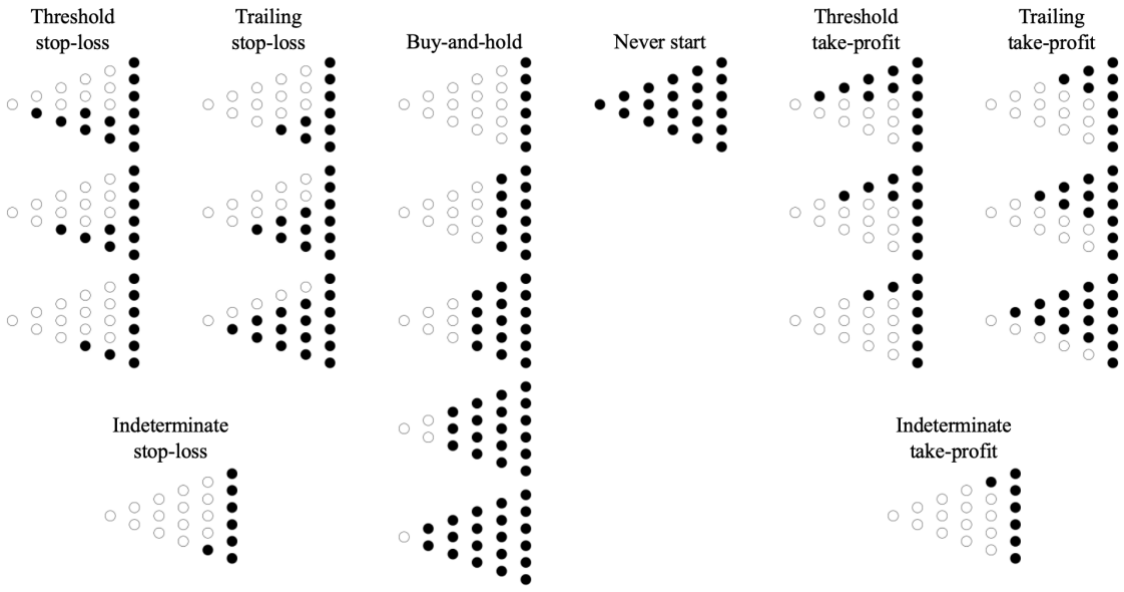
- *Never-start* strategies have a black first ball, such that the decoloring is black for all balls.
- *Buy-and-hold* strategies are time- but not outcome-contingent, such that the balls are white until the stopping repetition is reached, and black to the right of the stopping repetition.
- *Threshold 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. As in the main text, 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.

We also define one indeterminate stop-loss strategy and one indeterminate take-profit 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).[4]

[4] 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.

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

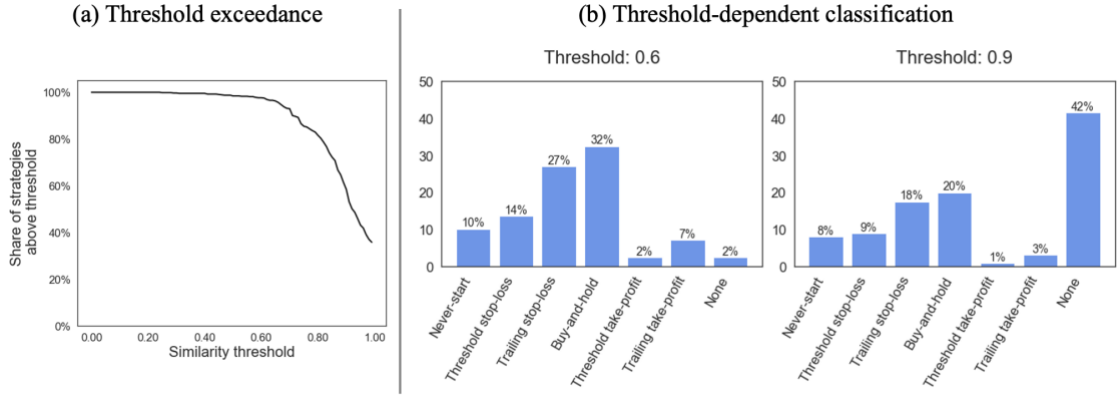
FIGURE SA.3.5.—All Basic risk-taking strategies for each class label



Notes: Figure SA.3.5 shows the comprehensive set of 20 *Basic* template strategies, with 3 threshold stop-loss, 3 trailing-stop-loss, 5 buy-and-hold, 3 threshold take-profit, 3 trailing take-profit and 1 never-start strategy. In addition, one stop-loss as well as one take-profit strategy are neither clearly a threshold nor a trailing strategy due to the finite horizon.

In the main text, we match elicited risk-taking strategies to the most similar labeled template strategy by counting the number of paths for which the elicited strategy deviates from the template. We only classify risk-taking strategies with a unique most-similar strategy class. If, for example, an elicited risk-taking strategy needs to be changed for three paths to become a threshold stop-loss strategy and for three (different) paths to become a buy-and-hold strategy, we do not classify that risk-taking strategy. We classify a risk-taking strategy if multiple templates within one class require, for example, changes for three paths.

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 shows the share of elicited risk-taking strategies whose cosine-similarity to the most similar template strategy exceeds the threshold parameter on the x-axis. We classify a risk-taking strategy as belonging to a class $c \in$ {never-start, buy-and-hold, stop-loss, trailing stop-loss, take-profit, trailing take-profit} if the similarity to the most similar strategy exceeds the threshold. We split indeterminate stop-loss and take-profit strategies evenly between the threshold and trailing type. Panel (b) shows the classification that results from using thresholds 0.6 and 0.9, respectively.

We complement the analysis of the main text by using the more granular cosine-similarity.[5] Each elicited strategy is matched to the template strategy to which it has the highest cosine-similarity. We assign an elicited strategy to a strategy class if the similarity to the matched template strategy is above a threshold and if the class of the most similar template is unique. 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 58% 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 (26% to 13% for a threshold of 0.6,

[5] 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 $||\cdot||$.

and 17% to 8% for a threshold of 0.9). Additionally, Table SA.3.3 shows the classification separately for each type of Decoloring Task, whereby we impose a similarity threshold of 0.6. We also 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 from 1.75:1 (*Basic*) to 3.10:1 (*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.3.3.—Cosine-similarity based classification of elicited risk-taking strategies for each type of Decoloring Task

Study	SL (Threshold stop-loss)	TSL (Trailing stop-loss)	Buy-and-hold	Never-start	Threshold take-profit	Trailing take-profit	Not identified	TSL / SL
Basic	60	105	94	11	5	18	20	1.75
Path-dependent	18	36	58	7	3	10	10	2.00
Randomized	10	31	56	47	8	18	15	3.10
Total	88	172	208	65	16	46	45	1.95

Notes: Table SA.3.3 shows the assignment of strategies to the classes of the 20 template strategies based on the cosine-similarity. We do not assign an elicited risk-taking strategy to a class if two template strategies from different classes are most-similar to the elicited risk-taking strategy or if the highest cosine-similarity to any template strategy is below 0.6. 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.[6]

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.3.4.—Class assignment for different supervised machine-learning algorithms

Class	Logistic regression	Decision tree	Random forest
Never-start	51	28	26
Buy and hold	86	284	156
Threshold stop-loss	112	103	60
Trailing stop-loss	145	160	98
Threshold take-profit	18	30	8
Trailing take-profit	49	35	9
Trailing stop-loss / Threshold stop-loss	1.29	1.55	1.63

Notes: Table SA.3.4 shows the assignment of elicited risk-taking strategies to the six strategy classes using three different supervised machine learning procedures. The bottom rows show the proportion of trailing to threshold stop-loss strategies.

Table SA.3.4 below 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 1.29 (logistic regression), 1.55 (decision tree), 1.56 (count-based

[6] For a motivating example see https://victorzhou.com/blog/gini-impurity/.

similarity), 1.63 (random forest), or 1.95 (cosine-similarity). Across all methods, we find a prevalence of trailing stop-loss over threshold stop-loss strategies of around 1.5.

SA.4 Finite mixture model

We now describe the finite mixture model that we use to estimate the different types of CPT agents in Section III.F 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

$$P\left(\tau_c^i \middle| \mathrm{T}_{Basic}, \theta_k\right) = \prod_{\tau_\neg \in \mathrm{T}_{Basic} \setminus \{\tau_c^j\}} f\left(\tau_c^j \middle| \tau_\neg, \theta_k\right) = \prod_{\tau_\neg \in \mathrm{T}_{Basic} \setminus \{\tau_c^j\}} \left(1 + e^{-\epsilon_i \, \Delta(\tau_c^i, \tau_\neg, \theta_k)}\right)^{-1},$$

where $\Delta\left(\tau_c^i, \tau_\neg, \theta_k\right) = V\left(S_{\tau_c^i} \middle| \theta_k\right) - V\left(S_{\tau_\neg} \middle| \theta_k\right)$. 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 or the probability of group membership. The log likelihood function of the finite mixture model is then given as

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

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-membership 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

$$\tilde{l}(\Theta_i | \tau_c^i, \mathrm{T}_{Basic}) = \prod_{k=1}^{K} \left[q_k P\left(\tau_c^i \middle| \mathrm{T}_{Basic}, \theta_k\right) \right]^{g_k^i}.$$

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

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

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

$$\widetilde{g_k^i}(\tau_c^i, \Theta^{(\mathrm{j})}, \mathrm{T}_{Basic}) = \frac{q_k^{(j)} \, P\left(\tau_c^i \middle| \mathrm{T}_{Basic}, \theta_k^{(j)}\right)}{\sum_{k'=1}^{K} q_{k'}^{(j)} \, P\left(\tau_c^i \middle| \mathrm{T}_{Basic}, \theta_{k'}^{(j)}\right)}.$$

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

$$q_k^{(j+1)} = \frac{1}{313} \sum_{i=1}^{313} \widetilde{g_k^i}\left(\tau_c^i, \Theta_{\mathrm{i}}^{(\mathrm{j})}, \mathrm{T}_{Basic}\right)$$

and

$$\left(\theta_1^{(j+1)}, \ldots, \theta_K^{(j+1)}\right) = \underset{\theta_1, \ldots, \theta_K}{\arg\max} \sum_{i=1}^{313} \sum_{k=1}^{K} \widetilde{g_k^i}(\tau_c^i, \Theta^{(\mathrm{j})}, \mathrm{T}_{Basic}) \, \ln P(\tau_c^i | \mathrm{T}_{Basic}, \theta_k),$$

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.[7] 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 $k+1$ leads 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.

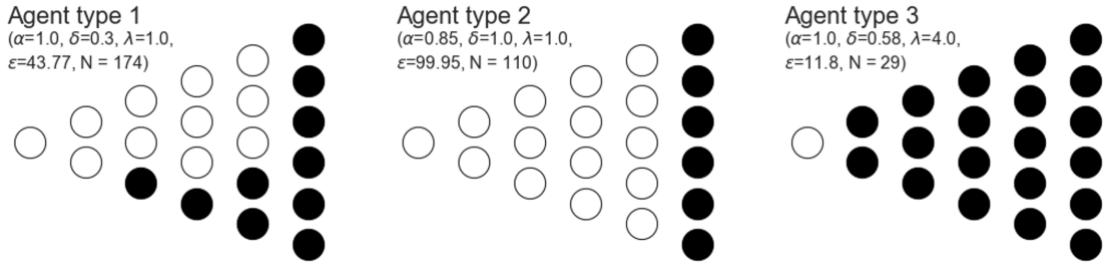
TABLE SA.4.1.— Model selection criteria for finite mixture models with up to seven types

Number of types	Log-Likelihood	AIC	BIC
1	-139,676.95	279,361.90	279,364.74
2	-105,550.47	211,118.95	211,125.32
3	-96,139.69	192,307.38	192,317.29
4	-91,143.47	182,324.95	182,338.40
5	-90,976.29	182,000.59	182,017.58
6	-86,383.44	172,824.89	172,845.42
7	**-85,498.45**	**171,064.90**	**171,088.97**

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.

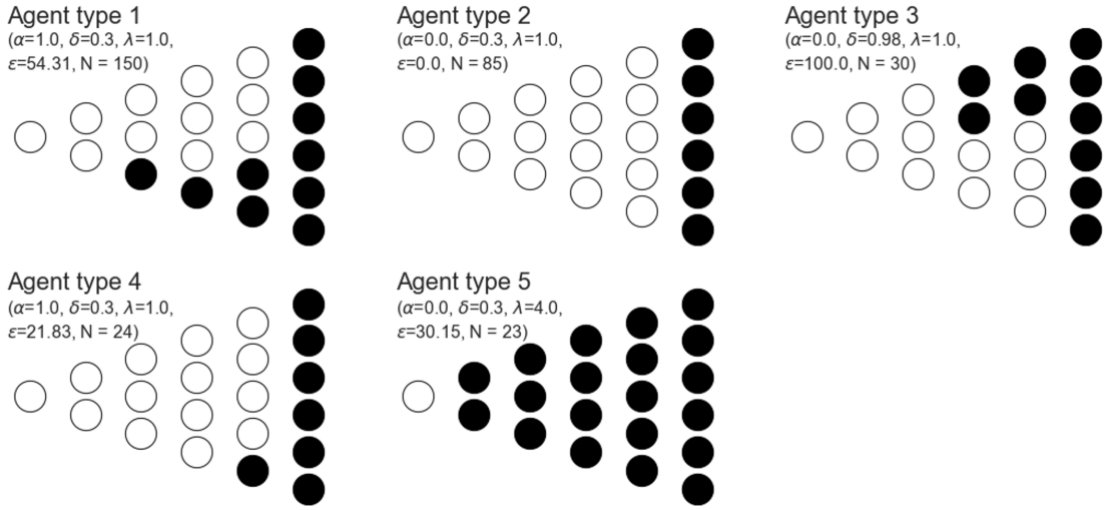
[7] In contrast to the median parameter estimates that we obtained from the individual estimation, the one-cluster specification takes interdependencies between the parameters into account. Therefore, these specifications are not necessarily identical. We estimate $(\alpha, \delta, \lambda) = (1.00,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

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.

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.

Comparing a setting with three agent types (Figure SA.4.1) and four agent types (Figure 7), 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 7—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 7. Behaviorally, the change implies a separation of the stop-loss types into two stop-loss types with different stop-loss levels.

Proceeding in the same way for the comparison between four types (Figure 7) 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 7 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 7 into agent types 2 and 3 in Figure SA.4.2, which singles-out a dedicated take-profit CPT-agent type.

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/tISHWp1Dl4A
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 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. Moreover, one can go through interactive versions of all studies at https://rtsexpdemo.onrender.com/demo.[8] To get a

[8] 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.

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. Below, we discuss the sample obtained from our experiment and payment details.

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.

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).

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.

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	0.0	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.

SA.5.2 Payment

Subjects are exogenously assigned to a study in order of joining the experiment.[9] Before receiving instructions about the experiment, subjects clicked one of 11 (10 for *No*

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

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 overall payment for the experiment consists of a completion fee of £3.00 ($3.57 at the time), an endowment of £1.50 ($1.79 at the time), and the task-dependent bonus.

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.