Eliciting Stopping Times

Sebastian Ebert Maximilian Voigt[†]

November 7, 2023

Abstract

We propose an experimental method to elicit *stopping times*—each subject's complete contingent plan for taking a risk for up to five times—to study repeated risk-taking under precommitment. In addition to time- and outcome-contingent risk-taking, we allow some subjects to use path-dependent or randomized stopping times. Our experimental design thus allows for hundreds of different risk-taking plans. Using an unsupervised machine-learning algorithm, we find that individuals' risk-taking strategies map well to stop-loss, take-profit, or buy-and-hold strategies. Most strategies are of a continue-when-winning and stop-when-losing type, with a profit-trailing stopping barrier. Path-dependence and randomization are used extensively, even if they are costly. We further analyze dynamic consistency in a sequential risk-taking task and find that subjects largely follow the unconstrained plans that we elicited.

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

JEL Classification: D01, D81, D91

[†] We thank Nicholas Barberis, Lucas Böttcher, Andreas Grunewald, Peiran Jiao, Pascal Kieren, Sven Rady, Simeon Schudy, Frederik Schwerter, Stefan Trautman, Christoph Vanberg, Jiakun Zheng (discussant), Xunyu Zhou, and seminar participants at the 50th Seminar of the European Group of Risk and Insurance Economists (EGRIE), European Decision Sciences Day 2023, Experimental Finance Conference 2023, Heidelberg University, and Yale 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 the Frankfurt School of Finance and Management and at the Department of Economics at Heidelberg University. E-mail: m.voigt@fs.de. 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). The experiment was pre-registered at the AEA RCT registry (AEARCTR-0009684).

I. Introduction

We propose an experimental method to elicit stopping times—the complete (time- and outcome-) contingent plan for taking a risk repeatedly. A large literature studied risk-taking in static contexts and identified various aspects that affect the attractiveness of static risks.[1] 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 risk-taking can systematically affect the risks' properties. Stopping times and their properties have received significant attention in the classical theory literature assuming standard preferences. More recently, stopping under non-standard preferences ("behavioral stopping") also received extensive attention.[2]

In this study, we experimentally elicit unconstrained—that is, complete contingent—stopping times for taking a risk repeteadly for up to five times. We visualize the five-times repeated binary, fair and symmetric risk as a ball triangle that represents a recombining binomial tree. The left 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 illustrates an exemplary time- and outcome-contingent stopping time whereby risk-taking stops at black balls and continues at white balls. If the Decoloring Task determines the subject's payment, we draw and visualize a random path through the decolored ball triangle. Risk-taking is stopped according to the subjects' stopping time τ. The Decoloring Task thus constitutes an incentivized method for the elicitation of unconstrained stopping times. The incentive-compatibility of the Decoloring Task roots in the strategy method (Selten 1967). Our visual and operational implementation is inspired by the behavioral stopping model of Barberis (2012) and subsequent theory articles such as He et al. (2017).

[1] Factors that affect the attractiveness of a (static) risk are summarized in preference functionals as, for example, disappointment aversion (Gul 1991), prospect theory (Kahneman and Tversky 1979; Tversky and Kahneman 1992), rank-dependent utility (Quiggin 1982; Yaari 1987), regret theory (Bell 1982; Loomes and Sugden 1982), or salience theory (Bordalo et al. 2012).

[2] 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.

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 must *decolor* the ball triangle by clicking each ball in an order of their choosing and thereby design their preferred stopping time τ, as exemplary shown in the right panel of Figure 1.

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 (He et al. 2017; Hu et al. 2022). Although our experimental design allows for hundreds of behaviorally different risk-taking strategies, the graphical interface helps subjects to comprehend the repeated risk-taking environment and ensures 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.

We find that individual risk-taking strategies elicited under precommitment are comprehensively described as stop-loss, take-profit, or buy-and-hold strategies. Using an unsupervised machine-learning algorithm—that could identify arbitrary patterns in the data—we find economically interpretable clusters of unconstrained risk-taking strategies. Our findings thus indicate that the restriction to these strategy types in theoretical work consitutes a reasonable approximation to actually desired strategies. Almost all subjects plan on taking the risk often, although the expected value of any induced outcome distribution S_τ is zero, which is incompatible with risk-averse expected utility preferences. Aggregate as well as individual risk-taking strategies are foremost stop-loss strategies that entail continue-when-winning and stopping-when-losing. Eliciting unconstrained risk-taking strategies further 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 3.5 times as often as threshold stop-loss strategies.

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.

In the second part of the experiment, subjects take the risk sequentially. After we elicit the unconstrained risk-taking strategy, each subject visually walks through her decolored ball triangle and can freely deviate from her plan at any node along the realized path. Crucially, comparing each subjects' *unconstrained* risk-taking plan to her sequential risk-taking actions allows us to interpret deviations from the plan as dynamic inconsistency. Note that, if a subject has to choose her plan from a constrained set of strategies, deviations to a different pursued strategy in sequential risk-taking would be fully rational if the pursued strategy was not included in the constrained choice set. In our experiment, we observe a high degree of dynamic consistency. Subjects follow their unconstrained plan for about 86% of their sequential risk-taking actions. If subjects deviate from their risk-taking plan, we find that they stop risk-taking somewhat earlier in the gain domain and continue risk-taking somewhat longer in the loss domain.

Many if not most risks that people face can be taken repeatedly and stopping problems abound. Typical economic applications concern the sale of an asset, job search, real options, or the decision to conclude data collection. By eliciting unconstrained stopping times, we hope to inspire experimental or empirical work on applied research questions. The experimental method introduced in this paper can be easily adapted for such purposes (see Section V). We also hope to inform the fundamental and applied theoretical literature on (optimal) stopping. Given this large literature (summarized, for example, in the textbooks of Shiryaev 2007 or Björk et al. 2021), it may be surprising that such unconstrained stopping times were not elicited yet. Allowing subjects to choose unconstrained risk-taking strategies also delivers insights into what subjects precisely want. We find that subjects use path-dependence, randomization, and the trailing version of stop-loss strategies often. Offering flexible strategies may increase individuals' ability to stick to their plan beyond commitment devices that focus on restricted plans.

Related literature. On a general level, we contribute to the literature on repeated risk-taking and stopping decisions. Classical experiments in this literature are Tversky and Shafir (1992), Redelmeier and Tversky (1992), Thaler and Johnson (1990), 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 et al. 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 complete contingent stopping times. We provide experimental evidence on the strategies that subjects use and document various properties of these strategies.

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 began to elicit stopping times for longer horizons and under constraints suitable for the application under study.[3] Fischbacher, Hoffmann, and Schudy (2017) show that the opportunity to commit to threshold stop-loss and take-profit strategies reduces the Disposition Effect.[4] Heimer et al. (2023) document a dynamic inconsistency, in that subjects start taking a risk repeatedly with a loss-exit plan by setting a (in absolute values) higher stop-loss than gain-exit threshold. Subjects' 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. In order to test implications of salience theory (Bordalo et al. 2012), Dertwinkel-Kalt and Frey (2022) show

[3] Previous research that elicited risk-taking strategies typically focused on threshold strategies in which the agent stops risk-taking as soon as the underlying stochastic process of gains or losses leaves an interval (a, b), $a < b$. Threshold strategies often yield analytically tractable results (Dixit 1993; Xu and Zhou 2013; Ebert and Strack 2015; Magnani 2015; Henderson et al. 2018), and are well-implementable in experiments even if the underlying stochastic process is defined in continuous time.

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

that subjects start taking a risk repeatedly with a loss-exit plan even if the stochastic process has a non-positive drift. In contrast, subjects in the experiment by Antler and Arad (forthcoming) 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 the efficiency of dynamic portfolio choices.[5] We contribute to this literature by proposing the Decoloring Task as a method to elicit complete contingent stopping times. In addition to designing path-dependent risk-taking strategies, we also allow 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 risks with non-zero mean (Dertwinkel-Kalt and Frey 2023), skewed risks (Ebert 2020), or risks with time-changing moments (Nielsen 2019) is well possible, as discussed in Section V.

We also contribute to the experimental literature on dynamic (in-)consistency in stopping decisions as theoretically predicted under non-standard preferences (Machina 1989; Cubitt et al. 1998; Barberis 2012; Ebert and Strack 2015; He and Zhou 2022). Our 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. Subjects in our study are largely dynamically consistent and mostly follow their risk-taking strategies. 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 previously choose the rule or express agreement therewith.

II. Experiment and Data

Eliciting unconstrained risk-taking strategies and identifying if people stick to unconstrained plans requires a controlled experimental setup to avoid confounding external factors. Our experiment is designed to enable a parsimonious and intuitive elicitation of unconstained risk-taking strategies. We describe the structure of the experiment and discuss the experimental design in this section. We also comment on procedural details.

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

II.A Structure of the experiment

The experiment consists of five between-subject studies, each having eleven tasks. In each task, subjects are endowed with £1.50 ($1.79 at the time) 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 subject's complete (time- and outcome-) contingent plan for taking the risk for up to five times—are represented through a recombining binomial tree, described as a ball triangle to subjects. Each ball of the triangle represents one of the possible accumulated profits that can result from taking the risk up to five times (see Figure 1).

Before the first task begins, subjects receive general instructions about the experiment and the risk as well as specific instructions about the first task. Figure 2 illustrates the course of the experiment. In the first ("main") task of the experiment, called *Decoloring Task*, each subject uses an interactive tool to design her precommitted risk-taking strategy by decoloring[6] the balls of an initially colored ball triangle. After the first task, subjects receive another short instruction about the second part of the experiment. The second part of the experiment consists of ten tasks, called *Walking Tasks*. In each Walking Task, subjects sequentially decide whether to take (another) risk depending on how the previous risk(s) turned out while seeing their plan from the first task of the experiment. The risk-taking strategy from the Decoloring Task is the default for each decision, but subjects can freely deviate from it. At the end of the experiment, subjects complete a short questionnaire and receive their payment, which is determined by the subject's behavior in one of the eleven tasks (selected at random). The experimental instructions are interactive and contain practice elements, comprehension questions, and attention checks.

[6] We refer to the task and the subject's interaction with the tool as decoloring, meaning that the ball triangle initially has a color (green, yellow, or red), and the subject's task is to transform the ball triangle into a greyscale ball triangle.

FIGURE 2.—Schematic storyboard of the experiment

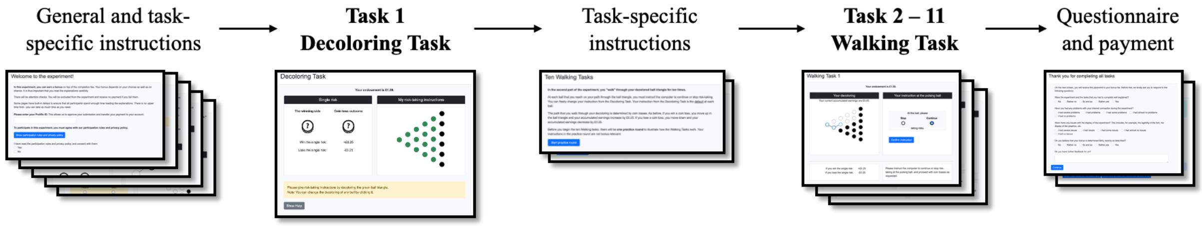

Notes: Subjects first receive general and task-specific instructions, accompanied by comprehension questions. We then elicit the precommitted stopping time in Task 1, the Decoloring Task. Subjects then receive another short instruction for Tasks 2 – 11, the Walking Tasks, and complete these ten tasks.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

II.B Studies

The five between-subject studies vary the available risk-taking options. Each subject in study *Basic* can decide to stop or continue risk-taking at any ball of the ball triangle. Subjects assigned to the *Path-dependence* study can additionally condition their risk-taking strategy 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 and to continue risk-taking if the first risk is lost and the second risk is won (up-down versus down-up).[7] Subjects in the *Randomization* study have access to an independent randomization device that allows them to stop risk-taking with a probability of their choice.[8]

Study *Path-dependence*^{C} *(Randomization*^{C}*)* is analogous to its counterpart described above, except that path-dependence (randomization) costs £0.01 per ball. Because risk-taking in the Walking Tasks is necessarily path-dependent (see Footnote 7), we only charge the cost in study *Path-dependence*^{C} in the Decoloring Task. The cost per ball as well as the total costs are saliently shown on subjects' screens.

[7] The Walking Tasks in studies *Basic, Path-dependence,* and *Path-dependence*^{C} 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 sequential risk-taking behavior during the Walking Tasks is thus necessarily conditional on the path taken to a ball and can no longer differentiate between counterfactual paths that could have occurred to a ball. Consequently, the subjects also cannot use a costly option during the 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.

II.C Decoloring Tasks

Figure 3 illustrates the new interactive tool that subjects use to design risk-taking strategies in the Decoloring Task. The first (second, third) row is for study *Basic (Path-dependence, Randomization).* Following Johnson et 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 part of the instruction panel—which is shown in panel (a) of Figure 3—and gives risk-taking instructions for that ball in the left part of the instruction panel, proceeding like this for each ball in an order of her choice. She can also go back and change the decoloring of any ball. Panel (b) of Figure 3 shows possible outcomes of this process, that is, exemplary decolorings, each specifying a possible stopping time. 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 does not take the 50:50 risk at that ball but stops risk-taking. Balls with a black-white gradient indicate that the stopping decision 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. If a ball is numbered with a greyscale decoloring, the subject uses an independent randomization device that yields a stopping outcome with the probability shown on the ball. Subjects use the decoloring scheme available in the respective study to design their preferred stopping time.

[9] In our setting, subjects give instructions to the computer and not to the experimenter. First, the computer is the eventual recipient of the instructions and implements them, such that we do not deceive subjects. Second, because computers strictly follow instructions given to them, this frame highlights the decisiveness of the instructions.

FIGURE 3.—Examplary decision screens and decolorings from the Decoloring Task

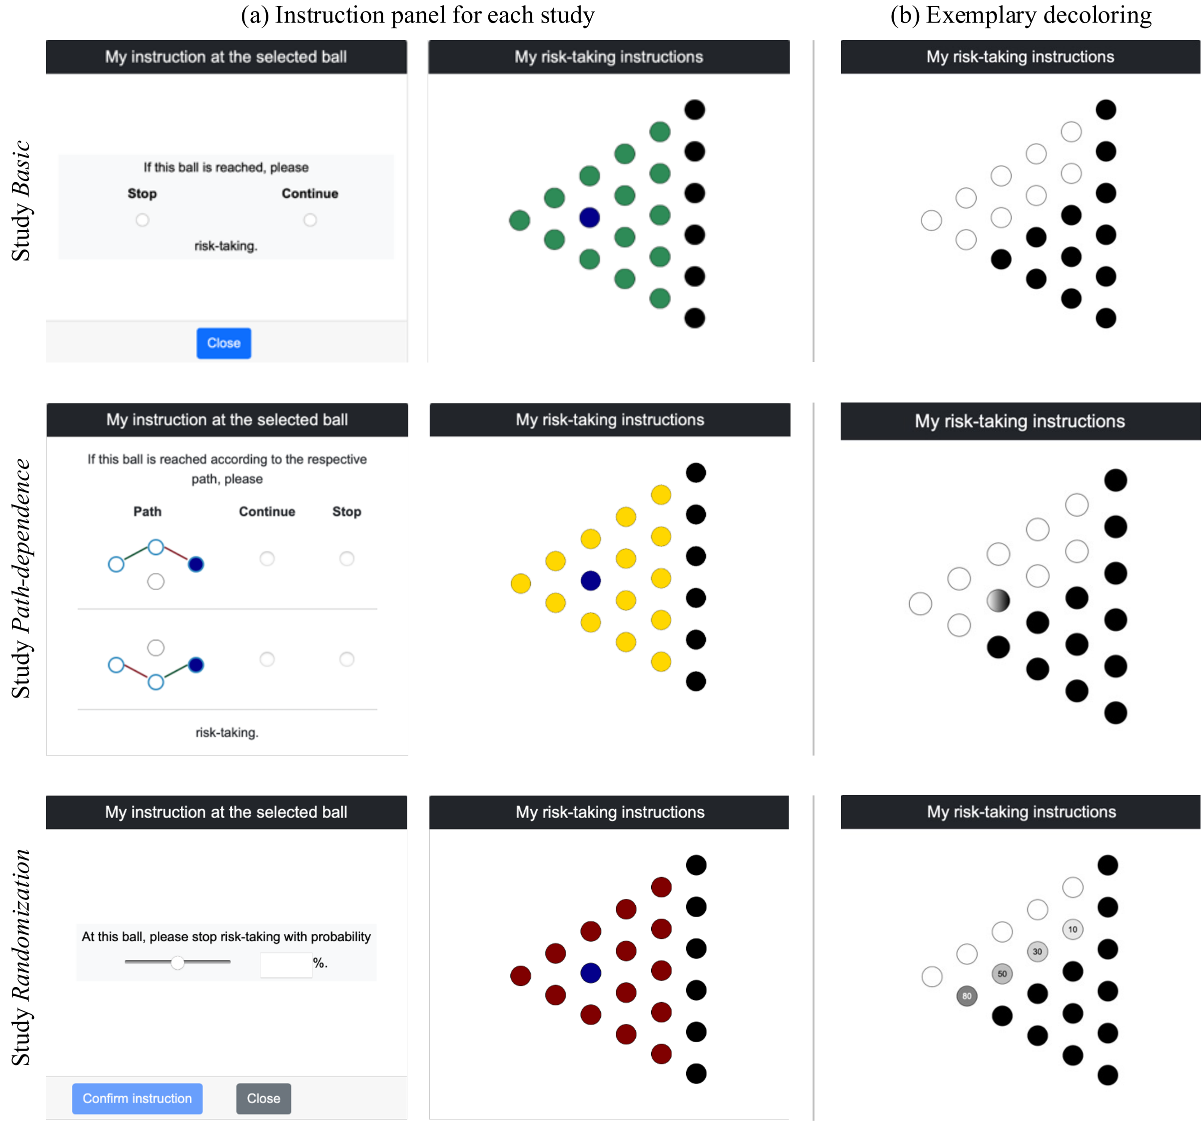

Notes: In the Decoloring Task, a subject in study *Basic (Path-dependence, Randomization)* must decolor all 15 balls of an initially green (yellow, red) ball triangle. Subjects click a ball in the right part of the instruction panel (a) and give the risk-taking instruction for that ball in the left part of the instruction panel (a). In study *Basic,* if the subject continues risk-taking, the ball is decolored white; if the subject stops risk-taking, the ball is decolored black. In studies *Path-dependence* and *Randomization,* subjects can also decolor balls to become black-white gradient (indicating a path-dependence) or numbered (the number indicating the stopping probability). Each subject proceeds like this for each ball in an order of her choice. She can also go back and change the decoloring of a ball. Panel (b) of Figure 3 shows exemplary decolorings, each specifying a possible strategy for repeatedly taking the risk up to five times.

II.D Walking Tasks

In the Walking Tasks, subjects make sequential risk-taking decisions, thereby "walking" through the ball triangle as decolored in the Decoloring Task. Figure 4 shows an example.

FIGURE 4. — Examplary decision screen from a *Basic* Walking Task

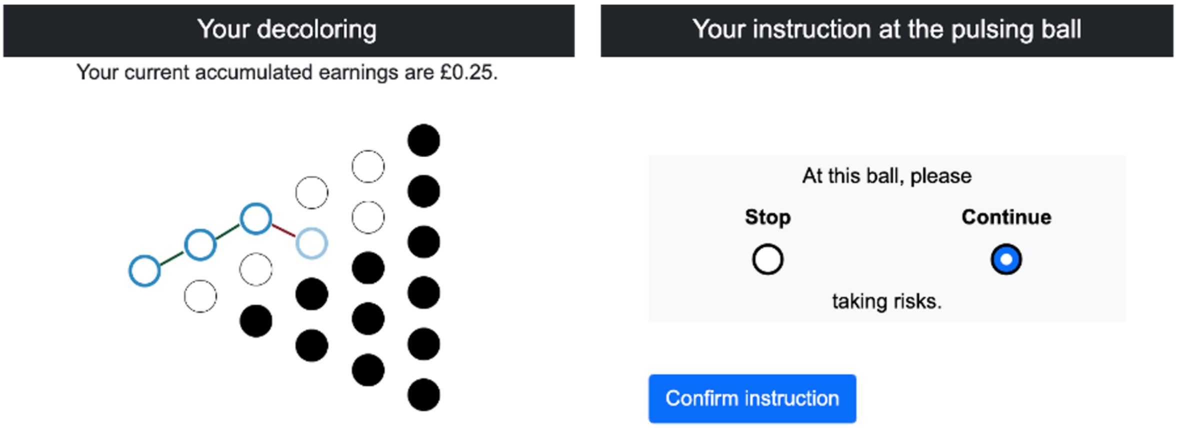

Notes: The subject sees the strategy she submitted in the Decoloring Task and the path she took thus far. She took the risk three times, won the first two risks and then lost the third risk. She must now decide whether to take the risk a fourth time. The default is to continue, in line with the strategy that she submitted in the Decoloring Task, but she may choose to deviate and stop risk-taking. The decision screens in the *Path-dependence* studies look as in the *Basic* study. The decision screens in the *Randomization* studies look the same, except that the instruction panel on the right is as in Figure 3, allowing for randomization.

During each Walking Task, subjects sequentially instruct the computer about their risk-taking action at each ball that is reached, possibly deviating from their previously elicited plan. Each Walking Task comprises taking up to five risks[10] and subjects complete ten Walking Tasks.

Subjects in studies *Basic, Path-dependence,* and *Path-dependence*^{C} can continue or stop risk-taking at any ball along their path (see Figure 4). Recall that sequential risk-taking is necessarily path-dependent (Footnote 7) and thus, there is no difference in the Walking Tasks across these studies. Subjects in the *Randomization* studies can stop risk-taking with any probability on the grid $0, 10, \ldots, 90, 100$. We implement an independent randomization device that subjects must resolve before taking a risk. If the randomization device yields a continuation

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

(stopping) outcome, risk-taking continues (stops) and the remaining procedure is as before. Subjects in study *Randomization*^{C} pay £0.01 for each randomization along their path.

II.E Randomization

Subjects are exogenously assigned to a study in order of joining the experiment while giving priority to studies with fewer finished and/or active subjects. Prior to receiving instructions about the experiment, subjects clicked one of 11 buttons to determine their individually payment-relevant task. Any 50:50 risk of winning or losing £0.25 was determined by a virtual coin toss (subjects had to click "Heads" or "Tails"). Subjects in the *Randomization* studies made virtual card draws (click a card from a set of cards, numbered from 1 to 10) to determine whether risk-taking stops or continues according to the chosen probability.[11] We used saliently different randomization devices to emphasize the different nature of uncertainty in the risk-taking decision and in the payoff determination, actively involving subjects in the risk resolutions (through selecting a side of the coin or a card). All risks were independent.

II.F Procedural details and sample

We programmed the experiment using oTree (Chen et al. 2016) and ran it via Prolific[12] on November 16, 2022. We obtained $N = 400$ completes (200 female and 200 male) and around 80 subjects participated in each of the five studies. The experiment took on average 30 minutes and the mean remuneration was £4.50 ($5.36 at the time). As pre-registered, we exclude subjects who failed an attention check or who did not complete all experimental tasks. Our participant selection filters lead to a final sample that is biased towards sophisticated and attentive subjects as in standard lab-experiments. Appendix A.3 provides further details and statistics on the sample.

[11] Subjects must also draw a card when stopping with probability 0 or 100 percent to ensure that subjects cannot finish the Walking Tasks faster by avoiding randomization.

[12] Online platforms, such as Prolific, are increasingly used in economics and the broader social sciences to recruit subjects for experiments. Studies have shown that laboratory results broadly replicate using subjects from online platforms (Horton et al. 2011; Gupta et al. 2021; Snowberg and Yariv 2021; Haaland et al. 2023). Douglas, Ewell, and Brauer (2023) find that subjects on Prolific and CloudResearch tend to be more attentive, work slower, and recall previously presented information better than subjects on other comparable platforms, including MTurk.

III. Results for the Decoloring Task: Stopping times

This section analyzes the stopping times designed during the Decoloring Task. Throughout the analysis of the Decoloring Task, we focus on 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 the approach taken by Barberis (2012).

III.A Study Basic

The left (right) panel of Figure 5 shows the median (mean) precommitted risk-taking strategy obtained from $N = 73$ subjects who completed study *Basic*. The median and mean are taken ballwise, such that a ball in the median decoloring is black if more than 50% of subjects stop risk-taking at that ball and white otherwise. The number on each ball in the mean decoloring in Panel (b) shows the percentage of subjects who instructed stopping at that ball, with balls being darker the more subjects stop. The overall picture is similar to the median decoloring, but noisier, so we focus on the median decoloring.

FIGURE 5. — Aggregate *Basic* precommitted risk-taking strategies

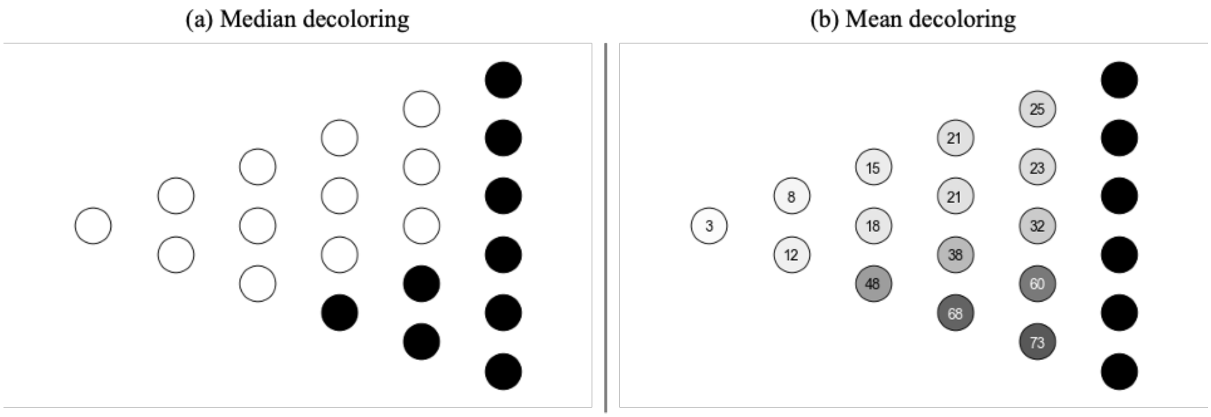

Notes: Panel (a) illustrates the median (panel (b) the mean) decoloring of subjects in study *Basic*, whereby the median (mean) is taken over the effective decoloring at each ball.

The median decoloring indicates risk-taking (white first ball), continue-when-winning (the upper part of the triangle being white), and stopping-when-losing (the lower black balls). 97% of subjects take the risk at least once (Panel (b) of Figure 5).[13] The median strategy is of the loss-exit type as defined by Barberis (2012). A strategy is loss-exit if the expected number of

[13] The large risk appetite in dynamic settings is in line with that of other experiments. Heimer et al. (2023), for example, report that around 95% of subjects start risk-taking.

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 so-called trailing stop-loss, in which the stopping barrier is increasing in the repetition of the risk. The trailing stop-loss level is 3, which, when ignoring that there is a final period, means that stopping occurs after losing thrice. The mean decoloring shows that the ball in the third row of the third column was, in fact, black for 48% of subjects. If it was black, the strategy would be of a threshold (stop-loss) type, that is, stopping once profits decline to -50p. We pursue the question whether trailing or threshold stop-loss is more common in Section III.G.

The nature of most strategies being continue-when-winning is also observed at the individual level. Panel (a) of Figure 6 shows that the stopping time τ and the induced outcome distribution S_τ correlate positively for most subjects ($t(72) = 6.75$). The average expected stopping time conditional on stopping with a gain is 4.32 and significantly higher than the mean stopping time conditional on stopping with a loss ($3.48, t(72) = 5.55$). Moreover, panel (b) shows that a majority of 58% of subjects plan to use a loss-exit strategy. Heimer et al. (2023) also observe a majority of loss-exit plans when subjects choose among threshold strategies. Here, we show that loss-exit strategies emerge endogenously when allowing for unconstrained risk-taking strategies and provide support for the popularity of loss-exit strategies using a different experimental design.

FIGURE 6.— Summary statistics of *Basic* precommitted risk-taking strategies

(a) Correlation between τ and S_τ (b) Strategy types

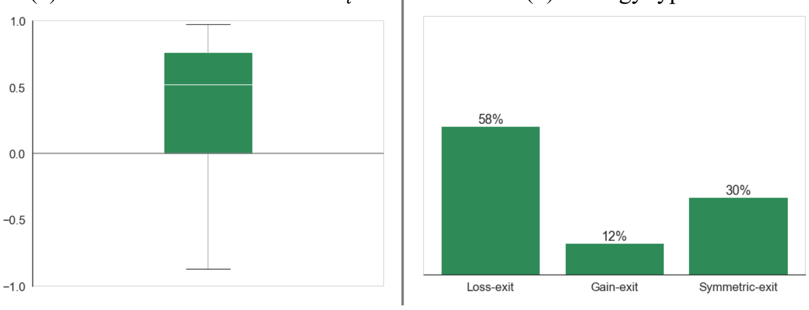

(c) Moments of τ and S_τ

	τ	S_τ
$\mathbb{E}[\cdot]$	3.7432	0.0000
$\mathrm{Std}[\cdot]$	0.8392	0.4722
$\mathrm{Skew}[\cdot]$	−0.2255	0.3553

Notes: Panel (a) of Figure 6 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. The table in panel (c) shows the average of the first three moments of the precommitted stopping time τ and induced outcome distribution S_τ.

The table in panel (c) of Figure 6 contains sample statistics of the stopping time τ and the induced outcome distribution S_τ that show that subjects take the risk often; on average 3.74 out of 5 times. They do so even though the mean of the induced outcome distribution is zero (due to the optimal stopping theorem), and thus risk-averse expected utility preferences are incompatible with the results of our experiment. Finally, note that S_τ is positively skewed. As observed by Barberis (2012), taking risks repeatedly in a way that skews the profit distribution to the right is consistent with a wide parametric range of cumulative prospect theory (CPT) preferences; we explore this relationship further in Section III.E.

III.B Study Path-dependence

In addition to time-and outcome-contingent risk-taking, subjects in study *Path-dependence* ($N = 75$) can costlessly condition their risk-taking strategy on the path taken to the ball. Panel (b) of Figure 7 highlights the six inner balls that can be reached by more than one path. 55% of subjects effectively use path-dependence at least once. Moreover, conditional on using path-dependence, 78% of subjects use it more than once. Panel (a) of Figure 7 shows the exact shares of subjects that use zero, one, …, five, or six path-dependent balls.[14]

We do not observe a simple systematic pattern in the usage of path-dependence. Panel (b) of Figure 7 visually separates path-dependent risk-taking when reaching a ball from above or below. For example, 65% of subjects that decolor the second-highest ball in the fourth column of the ball triangle as a path-dependent ball stop risk-taking if the path comes from above, and 57% of those subjects stop if the path comes from below. In panel (b) of Figure 7, subjects stop risk-taking more often for paths coming from above than from below at the non-hatched yellow balls. In contrast, subjects stop risk-taking more often for paths coming from below than from above at the hatched yellow balls. The position of the hatched and non-hatched balls in Panel (b) of Figure 7 suggest no simple systematic pattern in the usage of path-dependence (see Section V for further discussion).

[14] We reject the null that subjects never use path-dependence against the alternative that subjects use path-dependence ($t(74) = 7.968$). Note that we only consider *effective* path-dependence, whereby we ignore paths that cannot be taken when following the subjects' decoloring. We find more usage of path-depedence in the actual decolorings, where 72% of subjects use path-dependence at least once and 8% use it always.

FIGURE 7.—Prevalence and usage of costless path-dependent risk-taking

(a) Prevalence of path-dependence (b) Usage of path-dependent stopping

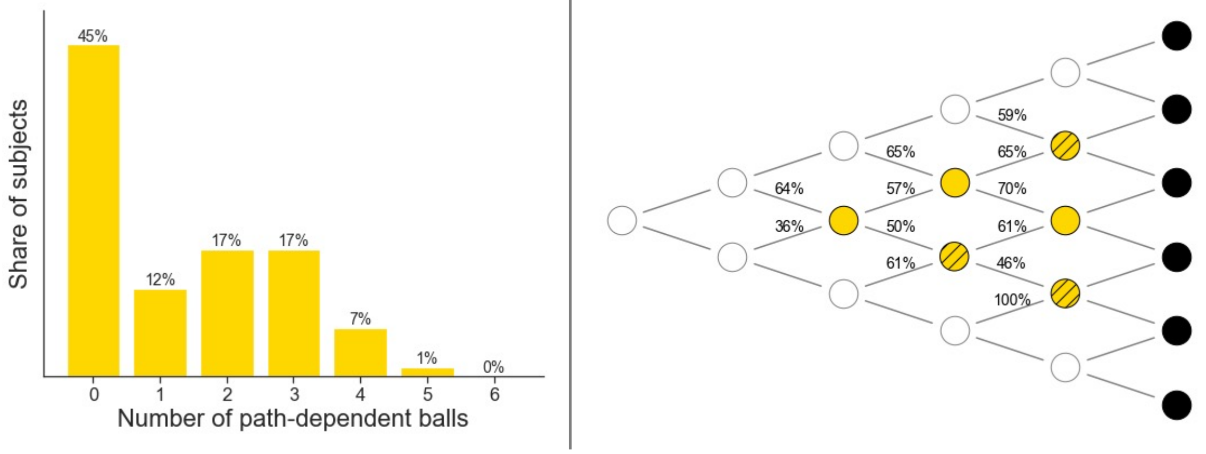

Notes: Panel (a) of Figure 7 shows the share of subjects for each number of potentially path-dependent balls. Subjects could use at most six path-dependent balls. Panel (b) shows the share of stopping decision for paths that reach a path-dependently decolored ball from above or below.

Although subjects use path-dependent risk-taking strategies if given the opportunity, we continue to find a popularity of trailing stop-loss strategies as in study *Basic*. The top left panel of Figure 8 shows the median decoloring for subjects in study *Path-dependence*, whereby we separately display the median risk-taking plan for paths that reach a ball from above or below by using a horizonal line (see, for example, the middle ball in the third column).

The median decoloring for subjects in study *Path-dependence* resembles the *Basic* median decoloring in that it shows risk-taking (first ball being white), and a continue-when-winning strategy. In total, 92% of subjects start risk-taking (see Figure B. 1 in the appendix) and the median strategy is (roughly) a trailing stop-loss strategy, in which subjects stop risk-taking after two to three losses. Since 48% of subjects plan to stop risk-taking at the bottom ball in the third column (see Figure B. 1), this ball could also have been black. Allowing for path-dependence leads to less overall risk-taking, which is consistent with some balls above the level-3 trailing stop-loss frontier that are white in study *Basic* being black in Figure 8 below. Moreover, a majority of subjects stops risk-taking at the second-highest ball in column five, which yields a gain of £0.50.

FIGURE 8.—Precommitted risk-taking strategies with costless *Path-dependence*

(a) Median decoloring

(b) Strategy types

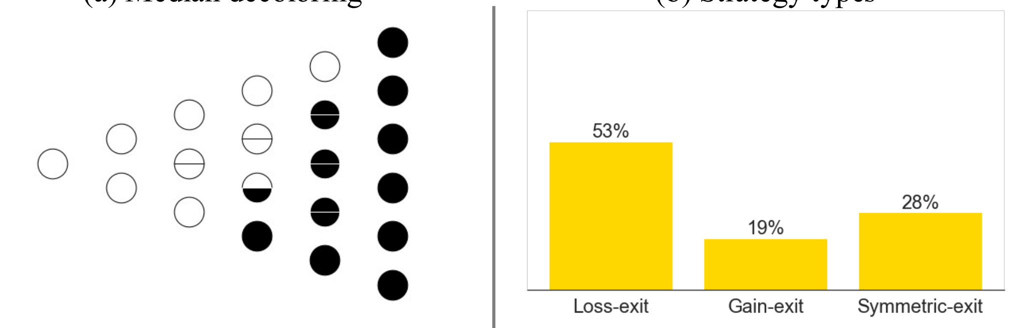

(c) Correlation between τ and S_τ

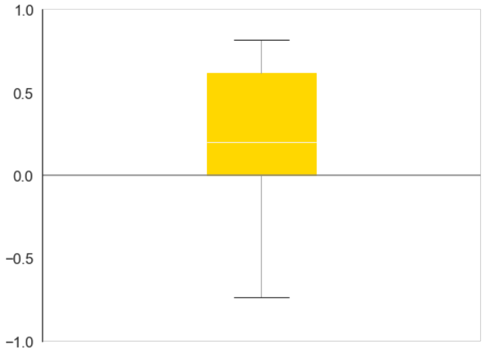

(d) Moments of τ and S_τ

	τ	S_τ
$\mathbb{E}[\cdot]$	3.3567	0.0000
Std$[\cdot]$	0.7998	0.4352
Skew$[\cdot]$	0.0154	0.2561

Notes: Panel (a) of Figure 8 shows the ballwise median decoloring of subjects in study *Path-dependence,* while 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. Panel (c) in the bottom row shows the distribution of the correlation between the stopping time τ and the induced outcome distribution S_τ. The table in panel (d) displays the average first three moments of τ and S_τ.

On an individual level, we observe that most subjects use a loss-exit strategy (panel (b) of Figure 8), with a higher share of gain-exit strategies than in study *Basic* (19% vs. 12%). Since most subjects use a continue-when-winning strategy, the mean expected stopping time conditional on stopping with a gain is larger than conditional on stopping with a loss (3.77 vs. 3.22, $t(74) = 4.350$); consistently, panel (c) shows that the correlation between the stopping time τ and the induced outcome distribution S_τ is positive for most subjects (average correlation 0.23, $t(74) = 5.126$). On average, subjects plan on taking the risk often, with an unconditional mean expected stopping time of 3.36, and use risk-taking strategies that induce a positively skewed outcome distribution (skewness of 0.26, $t(74) = 4.714$), as shown in panel (d) of Figure 8. Despite the prevalent usage of path-dependence, precommitted path-dependent strategies are behaviorally similar to *Basic* risk-taking strategies.

III.C Study Randomization

Next, we analyze if subjects that can costlessly use randomized stopping ($N = 88$, study *Randomization*) actually employ randomization. Panel (a) of Figure 9 shows that a majority of subjects uses randomization (81%), with more subjects randomizing at all 15 balls than at none of the balls (30% vs. 19%, respectively). Panel (b) of Figure 9 plots the distribution of randomization probabilities conditional on randomizing. Subjects display a tendency to use low randomization probabilities as 51% of all randomized balls stop risk-taking with a probabilitiy below 50 percent, while 25% stop with a probability higher than 50 percent.

FIGURE 9.— Prevalence and usage of costless *Randomization*

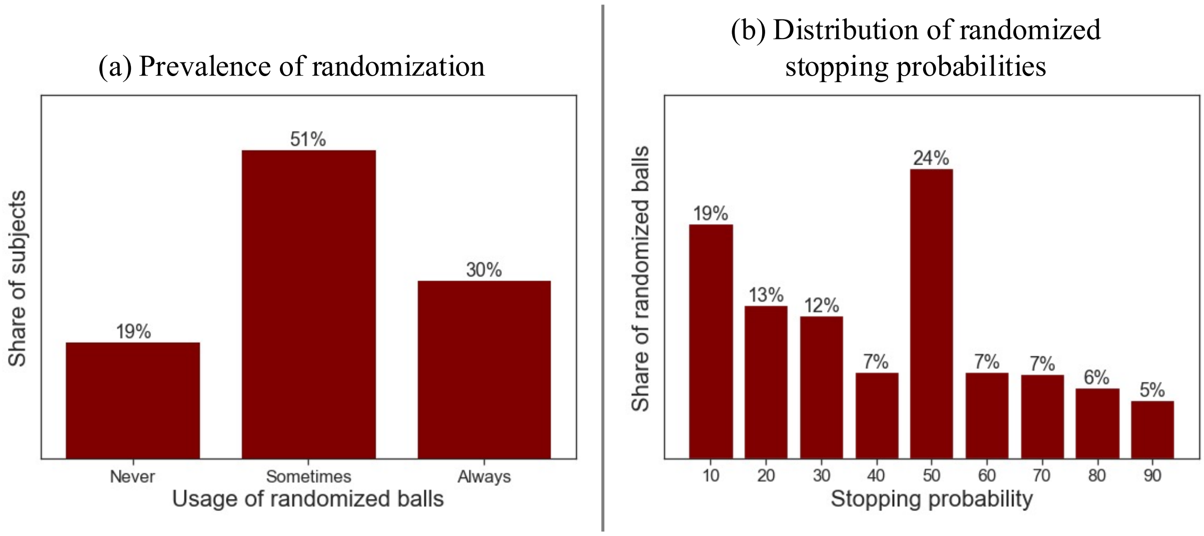

Notes: Panel (a) shows the share of subjects that never, sometimes or always (for all 15 balls) use randomization, while panel (b) shows the share of each stopping probability conditional on using randomization.

The high prevalence of randomization in combination with the usage of low randomization probabilities systematically affects risk-taking strategies. Panel (a) of Figure 10 shows the ball-wise median decoloring, wherby a darker greyscale corresponds to a higher median stopping probability that is shown inside each ball. 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% (see Figure B. 1). The median decoloring is of the continue-when-winning and stop-when-losing type in that the balls in the bottom of the ball triangle tend to be darker than in the upper part, although the pattern is less stark than in studies *Basic* (Figure 5) and *Path-Dependence* (Figure 8). All median stopping probabilities are strictly below 50 percent.

FIGURE 10.— Precommitted risk-taking strategies with costless *Randomization*

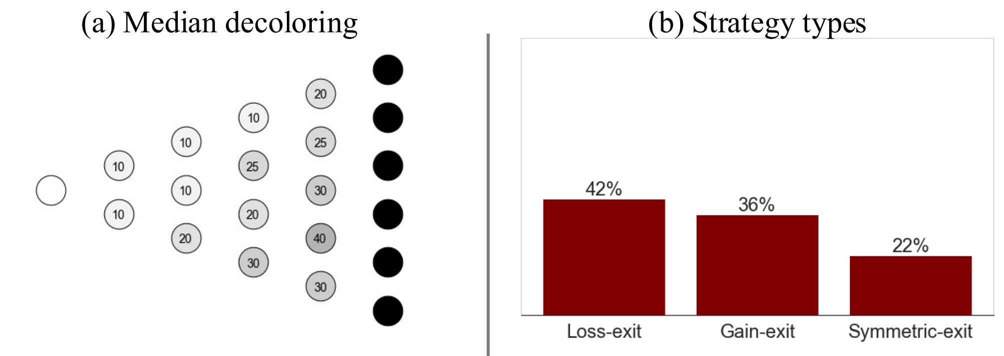

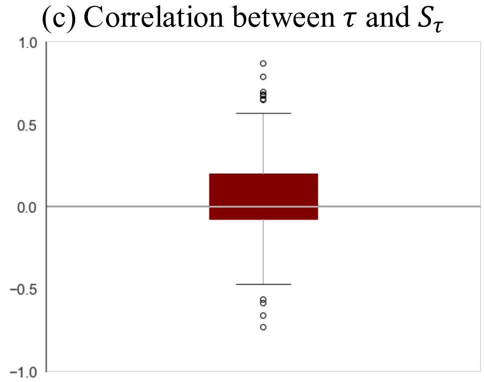

(d) Moments of τ and S_τ

	τ	S_τ
$\mathbb{E}[\cdot]$	2.9172	0.0000
Std$[\cdot]$	1.1280	0.4064
Skew$[\cdot]$	0.1238	0.0044

Notes: Panel (a) of Figure 10 shows the ballwise median decoloring of subjects in study *Randomization,* while 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. Panel (c) in the bottom row displays the distribution of the correlation between the stopping time τ and the induced outcome distribution S_τ. The table in panel (d) contains the average first three moments of τ and S_τ.

Panel (b) of Figure 10 shows that loss-exit strategies are also the most popular strategy type when allowing for randomization, which is in line with our observation that the median decoloring is of a continue-when-winning type. Compared to the previously discussed studies *Basic* and *Path-dependence*, we find an increased popularity of gain-exit strategies, whose share is almost as high as the share of loss-exit strategies. The average expected stopping times conditional on stopping with a gain or loss are statistically identical (3.39 vs. 3.32, $t(87) = 0.590$). Consequently, we do not find a significantly positive correlation between the stopping time τ and the induced outcome distribution S_τ for most subjects (average 0.06, $t(87) = 1.616$). In addition, the average expected stopping time of 2.92 (panel (d) of Figure 10) is lower than in studies *Basic* (3.74) and *Path-dependence* (3.36) since subjects place more weight on stopping at early repetitions. The induced outcome distribution is, on average, symmetric (skewness 0.004, $t(87) = 0.067$).

Overall, allowing for randomization systematically affects precommitted risk-taking strategies, leading to more cautious risk-taking at early repetitions, a higher weight on stopping in the gain domain, and less weight on stopping in the loss domain. Subjects take the risk with positive but low stopping probabilities at almost all balls in the triangle.

III.D Robustness towards costly path-dependence and randomization

In this section, we show that the preferences for using path-dependence and randomization are strict by making both options costly. The majority of subjects has a strictly positive willingness to pay for non-standard risk-taking options as offered in studies *Path-dependence*^{C} and *Randomization*^{C}. Imposing a cost on the usage of path-dependence or randomization does not affect their prevalence, which is thus unlikely to be the result of an experimenter demand effect. Figure B. 2 in Appendix B shows precise numbers.

Figure 11 highlights that imposing a cost on the usage of path-dependence (top row) or randomization (bottom row) does not affect aggregate risk-taking strategies. For ease of comparison, the left column of Figure 11 replicates the median decolorings without a cost, while the middle column shows the median decolorings for the respective costly study. The right column of Figure 11 shows the share of strategy types per study.

FIGURE 11.—Aggregate precommitted risk-taking strategies and individual-level strategy types for costless and costly *Path-dependence* and *Randomization*

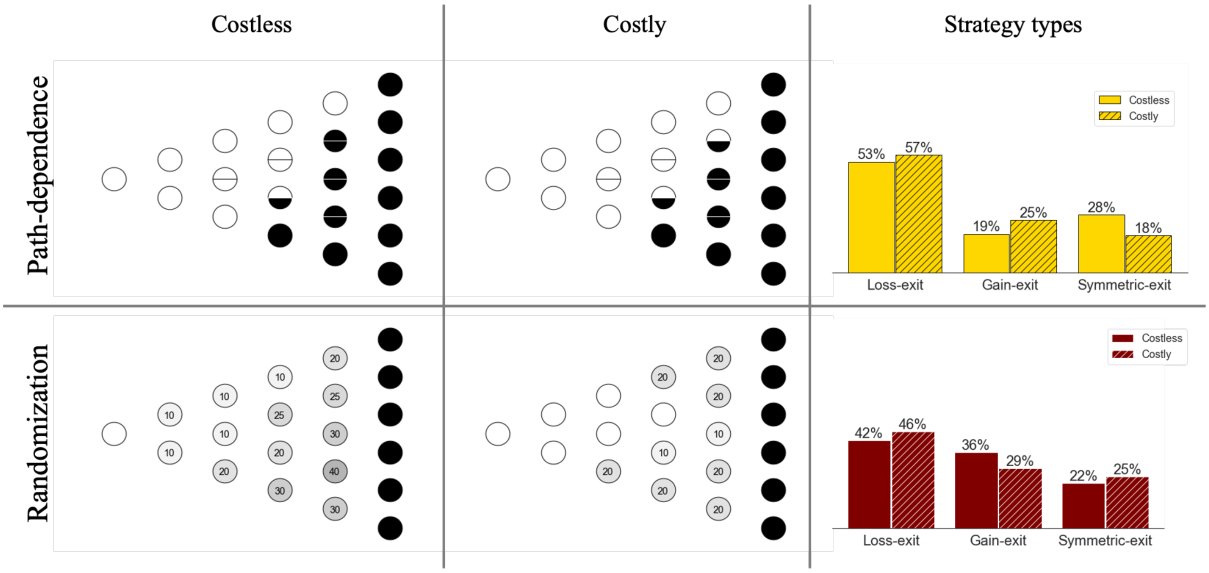

Notes: The left panel in each row of Figure 11 shows the median decoloring for subjects that can costlessly use path-dependence (top row) or randomization (bottom row), while the middle panel shows the median decoloring for subjects that must pay 1p for the usage of path-dependence (top row) or randomization (bottom row). The right panels show the classification of the strategy types in the respective studies.

The median decoloring of subjects in study *Path-dependence*^{C} is almost identical to the median decoloring of subjects in *Path-dependenc*e. The only difference is the white decoloring for paths that reach the second-highest ball in the fifth column from above, which does not affect the behavioral interpretation as a trailing stop-loss strategy. The median decoloring for study *Randomization*^{C} indicates more risk-taking than that for study *Randomization*, as we observe more white balls and lower median stopping probabilities for most balls.[15] In addition, we observe an increased (reduced) share of subjects that use loss-exit (gain-exit) strategies in study *Randomization*^{C} compared to study *Randomization.* Overall, studies *Path-dependence*^{C} and *Randomization*^{C} may largely be regarded as replications of their non-costly counterparts, evidencing that the preference for using path-dependence and randomization, respectively, is strict.

[15] The mean stopping probabilities at all nodes are, however, almost identical in both studies (see Figure B. 1), such that the shift in the median decoloring is due to an increased individual heterogeneity in the stopping probabilities in study *Randomization*^{C} compared to study *Randomization*.

III.E Structural CPT-parameter estimation

In the previous sections, we observed that (a) most subjects start risk-taking although the symmetric risk has an expected value of zero, and (b) most subjects' risk-taking strategies induce positively skewed outcome distributions S_τ in studies *Basic* and *Path-Dependence*. Barberis (2012) shows theoretically that CPT can explain such risk-taking behavior. In this section, we estimate parameters of the classical CPT formulation (Tversky and Kahnemann, 1992) that are consistent with our elicited risk-taking strategies. We find that subjects' risk-taking strategies are rationalizable using CPT, and that they are best explained by strong probability weighting and low diminishing sensitivity.

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. Outcomes are ordered such that $x_{-m} < x_{-m+1} < \cdots < x_0 < \cdots < x_n$, and x_0 corresponds to the reference point, taken to be the status quo. Tversky and Kahnemann (1992) assume that the utility derived from this gamble can be represented by

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

where

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

Following Kahnemann and Tversky (1992), we assume the probability-weighting function

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

and value function

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

where $\alpha \in (0,1)$, $\delta \in (0.28, 1)$ and $\lambda > 1$.

There are four ingredients that differentiate CPT from standard expected utility models. First, the carrier of value is the gain or loss relative to the reference point, since the argument of $v(\cdot)$ is $x_i - x_0$, rather than x_i. Second, $v(\cdot)$ is concave over gains and convex over losses to capture the finding that people are risk-averse over moderately-sized gains and risk-seeking over moderately sized losses. The degree of concavity/convexity is determined by α, with a lower value of α implying more curvature (i.e., more diminishing sensitivity). Third, the value function $v(\cdot)$ has a kink at the reference point x_0, so that the agent is more sensitive to losses than to gains of equal size. A higher value of λ implies a greater sensitivity to losses. Fourth, the agent does not use objective probabilities but transformed decision weights obtained according to the formulation above. The main consequence of probability weighting is that the agent overweights the tails of the distribution, with more overweighting the lower δ is.

Barberis (2012) shows that an agent with CPT preferences starts taking a fair and symmetric risk for up to five periods for a wide range of parameters. For each of the 8,000 parameter triples in the set $\Delta \coloneqq \{(\alpha, \delta, \lambda): \alpha \in \{0, 0.053, \dots, 0.947, 1\}, \delta \in \{0.3, 0.337, \dots, 0.963, 1\}, \lambda \in \{1, 1.16, \dots, 3.84, 4\}\}$, Barberis (2012) numerically searches for the basic risk-taking strategy that yields the highest CPT value and reports that a CPT agent starts risk-taking for 1,813 of the 8,000 parameter triples. Panel (a) of Figure 12 shows the parameter triples in Δ for which the CPT agent optimally starts risk-taking, replicating Figure 3 in Barberis (2012). A CPT agent starts risk-taking in mainly two parameter regions: For low δ and high α (towards the right of the cube), as well as for high δ and low α (towards the left of the cube). Agents with low δ and high α find positively skewed risks—as induced by loss-exit strategies—attractive. Intuitively, an agent with a low δ overweights the tails of the distribution more heavily, while the marginal utility of large gains diminishes less rapidly for an agent with high α, making positively skewed risks—which feature a large gain with low probability—attractive. Contrary, an agent with high δ and low α finds negatively skewed risks—risks featuring a large loss with low probability as induced by a gain-exit strategy—attractive since she does barely overweight the low probability of a large loss, and the low α makes the large loss only slightly more frightening than a small loss. Barberis' (2012) theoretical analysis finds more parameter triples in the low $\delta -$ high α than in the high $\delta -$ low α region that yield risk-taking, but it is ultimately an empirical question which parameter region rationalizes more observed risk-taking strategies.

We therefore take the flip-side of the theoretical approach developed in Barberis (2012) and Hu, Obloj, and Zhou (2023): For any elicited risk-taking strategy, we obtain the induced outcome distribution S_τ and search for each parameter triple in Δ that yields a positive CPT value, thus rationalizing the risk-taking strategy. Our procedure finds regions of parameters for which a CPT agent prefers the precommitted risk-taking strategy over not taking the risk.[16]

[16] It is not possible to obtain a unique estimate of preference parameters from the one risk-taking strategy that we elicit per subject. To see this, note that the CPT value of every risk-taking strategy decreases in the value of loss aversion λ, such that only parameter triples with $\lambda = 1$ maximize the CPT value for any (non-degenerate) risk-taking strategy. Similar arguments hold for parameters α and δ that determine subjects' skewness preference (see Ebert and Karehnke 2022).

We find that 302 of the 400 elicited risk-taking strategies achieve a positive CPT value for at least one of the 8,000 parameter triples.[17] Panel (b) of Figure 12 shows the parameter triples that rationalize elicited risk-taking strategies, with more transparent dots corresponding to parameter triples that rationalize fewer strategies. The mean and median parameters that yield a positive CPT value coincide and are $(\bar{\alpha}, \bar{\delta}, \bar{\lambda}) = (0.65, 0.50, 2.50)$.

FIGURE 12.—Structural CPT parameters that yield positive CPT values

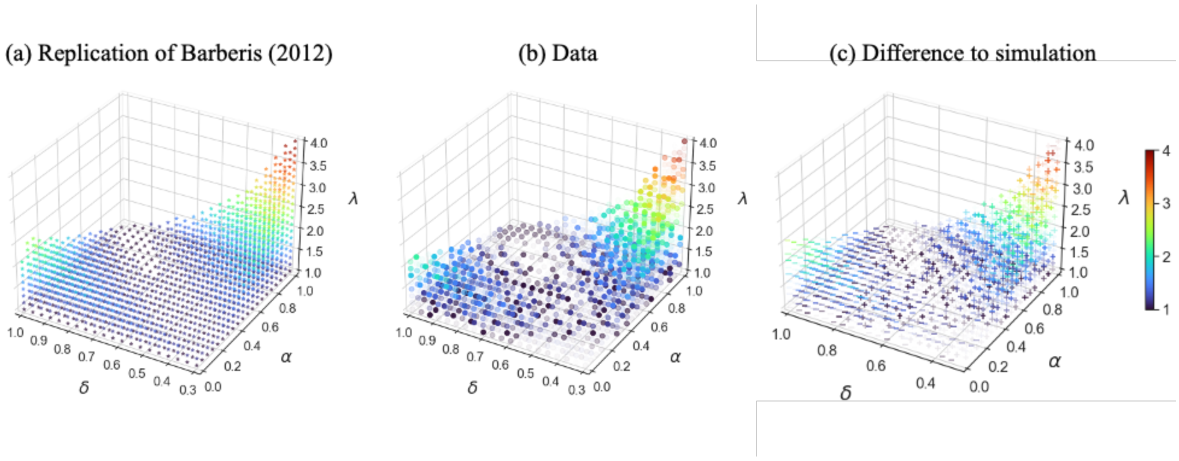

Notes: Panel (a) of Figure 12 shows parameter triples that yield a positive CPT value for some risk-taking strategy and replicates Figure 3 in Barberis (2012), except for omitting his classification as gain-exit and loss-exit strategies. Note that Barberis (2012) only considers time- and outcome-contingent risk-taking strategies. Panel (b) then shows the parameter triples that yield a positive CPT value for the risk-taking strategies elicited during the Decoloring Task. Panel (c) compares the empirical prevalence of each parameter triple in panel (b) with simulation results (i.e., with parameter triples obtained from randomly generated risk-taking strategies). A minus (plus) sign means that we empirically observe a parameter triple less (more) often than in the simulation results. The color of each ball corresponds to the value of λ as shown in the colorbar, and balls are more transparent if they yield positive CPT values for fewer risk-taking strategies.

Panel (b) of Figure 12 shows that we observe more strategies that are rationalizable with a high degree of probability weighting (low $\delta -$ high α region) than with a low degree of probability weighting (high $\delta -$ low α region), as highlighted by the more pronounced dots in the right part of the cube. In addition, less loss aversion, as captured by lower values of λ, better describes the risk-taking strategies of our subjects. The shape of our empirical parameters in Panel (b) is similar to the shape of the theoretically possible parameters shown in Panel (a). 1,681 of the 8,000 parameter triples yield a positive CPT value for at least one elicited risk-

[17] All 98 subjects with a symmetric-exit strategy achieve a maximal CPT value of 0 and are thus indifferent between their risk-taking strategy and not taking the risk.

taking strategy, which is close to the 1,813 parameter combinations for which Barberis (2012) finds risk-taking.

We next investigate if parameter combinations with a high degree of probability weighting (low $\delta -$ high α region) are in fact more pronounced in our sample by comparing our findings to simulation results: We simulate 400 random risk-taking strategies in the same proportion as the elicited risk-taking strategies, i.e., 73 time- and outcome dependent strategies, 142 path-dependent risk-taking strategies, and so on. For each simulated risk-taking strategy, we compute the induced outcome distribution and search for parameter triples in Δ that yield a positive CPT value. We reduce noise in the results by repeating the simulation 100 times and report average values. Panel (c) compares the simulation to the empirical results: A plus (minus) sign shows more (fewer) rationalizable risk-taking strategies in our elicited data than in simulations, with more transparent symbols indicating smaller differences. Panel (c) shows that parameter combinations with high probability weighting (the low $\delta -$ high α region) indeed rationalize more risk-taking strategies than expected under random sampling, and that we observe fewer risk-taking strategies rationalized by the high $\delta -$ low α region in the left part of the cube than under random sampling.

We thus find (i) that many parameter combinations rationalize at least one risk-taking strategy, indicating that subjects' risk-taking strategies are consistent with CPT, and (ii) that our subjects' risk-taking strategies indicate a high degree of probability weighting (low δ), a low diminishing sensitivity (high α), and low loss aversion (low λ). Unreported results replicate the findings on the study level, with less differences between elicited and simulated strategies for *Randomization* studies.

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

We now turn to an individual-level analysis of the precommitted risk-taking strategies by using an unsupervised machine-learning algorithm, specifically a K-means clustering.[18] 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 similar clusters which may either be economically

[18] Unsupervised machine-learning algorithms aim at finding patterns in unlabeled data. In contrast, supervised machine-learning algorithms first learn a mapping between data X and an output vector y that is then applied to new, previously unseen data. We describe and employ different supervised algorithms in Appendix D. Appendix D also contains further methodological details and robustness results.

interpretable or nudge us to consider entirely new strategy types. 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 assigned to clusters that we interpret as stop-loss strategies.

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 in the main text using $K = 6$ clusters for simplicity, but Appendix D shows the robustness of our results towards using more clusters.

Our K-means clustering with $K = 6$ clusters separates the elicited risk-taking strategies into economically interpretable groups. Figure 13 plots the median risk-taking strategy for each cluster. Visually, the median strategy of cluster 1 is a time-dependent (buy-and-hold) strategy, while the median strategy of cluster 2 does not start risk-taking with 50% probability and otherwise takes the risk exactly once. The median strategy of cluster 3 is roughly a (trailing) take-profit strategy, and that of clusters 4, 5, and 6 is a trailing stop-loss, each with a different trailing stop-loss level and a varying blur of the stopping barrier.

FIGURE 13.—Median precommitted risk-taking strategy for each cluster of a K −means clustering

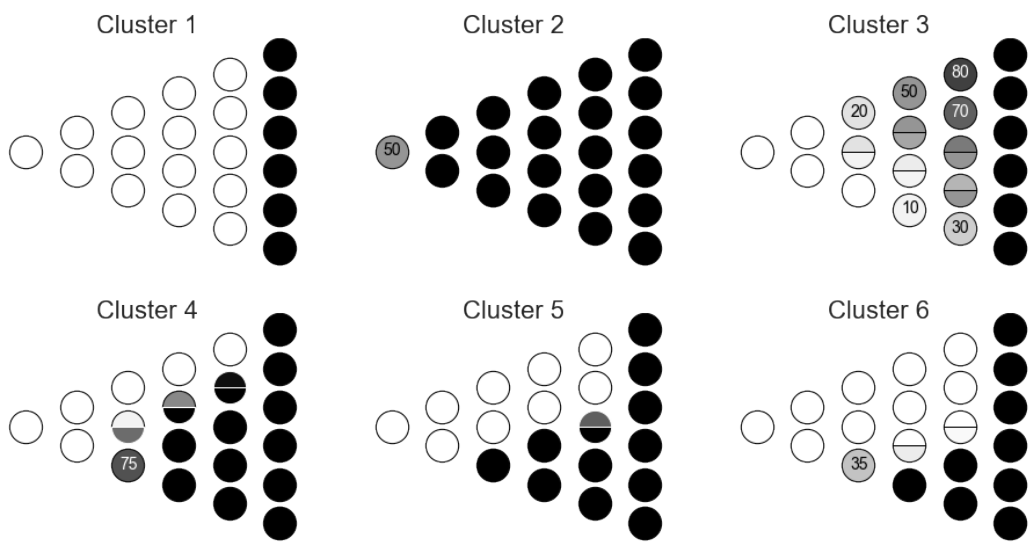

Notes: Figure 13 plots the median decoloring per cluster resulting from a K-means clustering with $K = 6$ clusters. We achieve stability in the cluster assignment by using a majority voting over 200 iterations of the K-means clustering with 100 iterations for the random and k-means++ initialization, respectively.

Our visual interpretation is supported by an analysis of the cluster constituents. Table 1 provides an overview of the number of elicited strategies per cluster, the average moments of the stopping time τ, the skewness of the induced outcome distribution S_τ, as well as the Barberis (2012) strategy types. Strategies in cluster 1, for example, have a high expected stopping time with an equal number of risks taken in the gain and loss domain, inducing symmetric outcome distributions. Most individual strategies in cluster 1 correspond to symmetric-exit strategies as defined by Barberis (2012), such that we interpret cluster 1 as containing buy-and-hold strategies ($N = 121$). Proceeding in the same way for all six clusters, we interpret clusters 4, 5, and 6 as 'stop-loss' strategies ($N = 163$), cluster 2 as 'never-start' strategies ($N = 37$), and cluster 3 as 'take-profit' strategies ($N = 79$). A relative majority of subjects uses strategies that belong to clusters interpreted as stop-loss strategies, while few subjects use take-profit strategies. Moreover, the general willingness to take risks in dynamic contexts is evident from the low number of subjects in the never-start cluster.

TABLE 1.—Analysis of cluster constituents

						Barberis (2012) classification		
Cluster	N	$\mathbb{E}(\tau)$	$\mathbb{E}(\tau\|\text{Gain})$	$\mathbb{E}(\tau\|\text{Loss})$	$\text{Skew}(S_\tau)$	Loss-exit	Gain-exit	Symmetric-exit
1	121	4.20	4.50	4.51	-0.03	34	37	50
2	37	0.92	1.00	1.04	0.01	6	8	23
3	79	2.84	3.03	3.52	-0.21	22	44	13
4	70	2.73	3.63	2.68	0.50	54	8	8
5	41	3.54	4.86	2.73	0.81	41	0	0
6	52	4.11	4.66	3.70	0.37	45	3	4
Total	400	-	-	-	-	202	100	98

Notes: Table 1 shows the assignment of strategies to the six 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. The last three columns show the number of subjects with a loss-exit, gain-exit, and symmetric-exit strategy in each cluster, based on each subject's expected stopping time conditional on stopping with a gain or loss (Barberis 2012).

We find that an unsupervised machine learning algorithm partitions unconstrained risk-taking strategies into economically reasonable clusters. Neither did the algorithm—by definition—learn strategy types from training data, nor did the experimental design restrict the risk-taking strategies of subjects to certain types. Nevertheless, we find a clustering of risk-taking strategies into types that have previously been identified, analyzed, and discussed in the economics and finance literature. Our findings support theoretical and empirical work that restricts attention to these strategy types. Even when subjects are not offered particular strategy types but can design an unconstrained risk-taking strategy, they choose simple strategy types known from the theoretical literature.

III.G Differentiating between threshold and trailing stop-loss strategies

In the previous section, we observed that a relative majority of risk-taking strategies are assigned to clusters interpreted as stop-loss strategies. Eliciting unconstrained risk-taking strategies allows us to further analyze what *kind* of stop-loss strategies subjects choose. We specifically analyze if subjects use threshold stop-loss or trailing stop-loss strategies by first defining a comprehensive set of labelled template strategies, and then assigning each elicited risk-taking strategy to the class of its most similar template strategy. As the median strategies in the three stop-loss clusters in Figure 13 already suggest, we find a large prevalence of trailing stop-loss strategies in absolute terms as well as relative to threshold stop-loss strategies, with a ratio of trailing to threshold stop-loss strategies of approximately 3.5.

As a first step, 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, we define a level-3 trailing stop-loss strategy (as in the median *Basic* decoloring in Figure 5) as one template, but also include a level-2 and level-1 trailing stop-loss as well as corresponding threshold stop-loss strategies in the set of templates. Figure D. 5 in the appendix shows all 20 templates. Defining a comprehensive set of templates allows us to cleanly identify threshold as well as trailing stop-loss strategies among the elicited risk-taking strategies and avoids a confounded classification of, for example, buy-and-hold strategies.

We measure the similarity between an elicited risk-taking strategy and each template by counting the number of paths for which the elicited risk-taking strategy deviates from each template. An elicited risk-taking strategy is closer to a template if it deviates for fewer paths, and is assigned to the class of its most similar template.[19,20] We define neither path-dependent

[19] In Appendix D.2, we also measure the similarity between the elicited risk-taking strategies and the templates using cosine-similarity and supervised machine-learning procedures. Our results are robust to the choice of a specific method.

[20] 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.

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 thus relies only on balls that are not path-dependent or randomized. Table C. 4 in the appendix shows that the majority of risk-taking strategies that cannot be uniquely classified are strategies that use randomization, but our results are robust to using more granular similarity measures (see Appendix D.2).

Table 2 shows the assignment of elicited strategies to the templates. We separately show the assignment for each cluster that the K-means algorithm identified in Section III.F 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 (43%), 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 (10%). Overall, trailing stop-loss strategies are chosen more frequently than threshold stop-loss strategies, with a ratio of 3.19.[21] The analysis also corroborates the finding 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.

[21] On an aggregate level, we find similar proportions when using a cosine-similarity measure (3.60:1, see Table D. 4) or different supervised machine-learning procedures (3.16:1, 3.44:1, and 3.96:1, see Table D. 5). The variation is higher on a study level, ranging from 1.67:1 to 4.75:1 for the count-based measure (see Table C. 4).

TABLE 2.—Deviation-count based classification of elicited risk-taking strategies

	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	6	-	67	-	5	2	41	121
2	-	-	16	16	-	-	5	37
3	-	-	15	-	5	15	44	79
4	-	23	18	-	-	2	27	70
5	1	37	-	-	-	-	3	41
6	19	23	1	-	-	-	9	52
Total	26	83	117	16	10	19	129	400

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

In addition, Table 2 shows that our interpretation of the six clusters identified by the K-means algorithm in the previous section aligns with the assignment of individual strategies. Most strategies in cluster 1 are buy-and-hold strategies, cluster 2 contains all never-start strategies, cluster 3 contains most trailing and threshold take-proft strategies, while clusters 4, 5 and 6 contain mostly trailing stop-loss strategies, with cluster 6 also containing many threshold stop-loss strategies.

Overall, we observe a large prevalence of trailing stop-loss strategies, in absolute terms as well as relative to threshold stop-loss strategies. When subjects can choose unconstrained risk-taking strategies, they mostly do *not* select threshold strategies. While a relative majority of subjects chooses strategies that are closest to a buy-and-hold strategy, the ratio of trailing stop-loss strategies to threshold stop-loss strategies is around 3.5. Moreover, trailing take-profit strategies—although not observed often—are also chosen more often than threshold take-profit strategies.

IV. Results for the Walking Tasks and dynamic consistency

This section analyzes the sequential risk-taking actions and dynamic consistency of subjects. After submitting a precommitted risk-taking strategy (in the Decoloring Task discussed above), subjects take the risk sequentially during ten Walking Tasks and can freely deviate from their risk-taking strategy. Our within-subject design allows us to compare each subject's risk-taking strategy and actions to analyze if subjects are dynamically consistent.[22] We construct a "reverse decoloring" that maps subjects' risk-taking actions onto the ball triangle. Since sequential risk-taking is necessarily path-dependent (see Footnote 7), the reverse decoloring of each subject is also path-dependent. We therefore differentiate between subjects that can or cannot use randomization in the following analyses.

IV.A Reverse decolorings

Panel (a) in Figure 14 shows the reverse decoloring for $N = 215$ subjects that can continue or stop risk-taking at each ball (as in the *Basic* and *Path-dependence* studies), and panel (c) shows the reverse decoloring for $N = 185$ subjects that can use an independent randomization device (as in the *Randomization* studies*)*. We aggregate risk-taking actions for paths that reach a ball from above or below, respectively, and show the average stopping probability inside each (half) ball. The table in each row displays the first three moments of the realized stopping times τ and profits S_τ.[23]

[22] Pooling all subjects, we observe a a comprehensive set of 17,877 risk-taking actions. For any node, we observe at least 115 actions, with a maximum of $4{,}000$ actions ($= 400$ subjects $* \ 10$ Walking Tasks) for the very first node in the ball triangle.

[23] For subjects in the *Randomization* studies, the realized stopping time and profit are contingent on the outcome of an independent randomization device. When constructing the realized stopping time τ, we attribute any remaining probability mass to the next risk repetition if the randomization device yields a stopping decision. Taking the contingency into account and ignoring randomizations along the walked paths, we find $\text{Mean}[\tau] = 3.10$, $\text{Std}[\tau] = 1.24$ and $\text{Skew}[\tau] = -0.26$ for subjects in *Randomization* studies.

FIGURE 14.—Aggregate risk-taking actions

Basic and *Path-dependence*

(a) Reverse decoloring

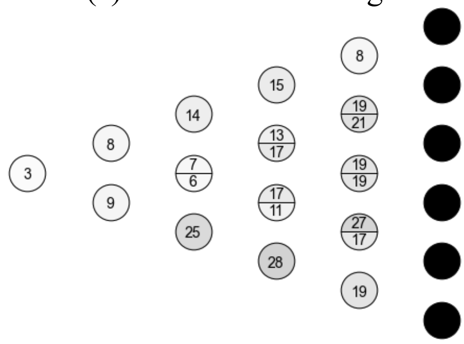

(b) Moments of τ and S_τ

	τ	S_τ
Mean[·]	3.7874	−0.0190
Std[·]	0.9832	0.4426
Skew[·]	−0.4956	0.0865

Randomization

(c) Reverse decoloring

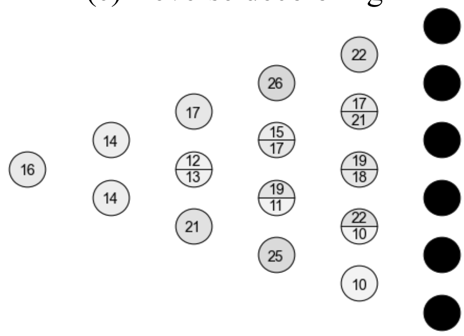

(d) Moments of τ and S_τ

	τ	S_τ
Mean[·]	2.8505	−0.0059
Std[·]	1.1717	0.3795
Skew[·]	−0.0332	−0.0276

Notes: Panel (a) of Figure 14 shows the reverse decoloring of $N = 215$ subjects that can stop or continue risk-taking during the Walking Tasks, whereas the panel (c) shows the reverse decoloring of $N = 185$ subjects that can randomize during the Walking Tasks. The right panel in each row shows the first three realized moments of the stopping time τ and profit distribution S_τ of subjects in the respective studies.

The reverse decoloring in panel (a) of Figure 14 shows—first—that subjects start risk-taking in 97% of all 2,150 Walking Tasks, and are thus not averse to taking a symmetric and fair risk if they are being offered to do so for up to five times. Second, the share of stopping actions increases with the risk repetition, since the balls are darker towards the right of the ball triangle.[24] Third, the average stopping probability is always below 30%, showing that the majority of subjects continues risk-taking at all balls. On average, subjects take the risk 3.79 out of 5 times (panel (b) of Figure 14). Fourth, we find no significant moderators of risk-taking actions. In particular, we do not observe differences in the average realized stopping time in

[24] On average, we find that 8% (13%, 17%, 20%) of all risk-taking actions are decisions to stop risk-taking at the second (third, fourth, fifth) repetition of the risk. Thus, subjects reach the fifth column of the ball triangle in 64.4% of all Walking Tasks, and always take the risk in 51.6% of all Walking Tasks.

the gain or loss domain (4.03 versus 3.91, $t(420) = 1.2319$), or systematically different reverse decolorings for paths that reach a ball from above or below.

Allowing subjects to use an independent randomization device (bottom row of Figure 14) leads to less risk-taking at early repetitions. Subjects in the *Randomization* studies take the first risk in 84% of the Walking Tasks, with a higher average stopping probability for the first three repetitions of the risk.[25] Subjects take the risk, on average, 2.85 out of 5 times, which is a significantly shorter mean realized stopping time than for subjects that cannot randomize ($t(329) = 7.4990$). The remaining findings are not affected by randomization. Figure B. 3 in the appendix replicates Figure 9 for the Walking Tasks, showing that subjects randomize often (88% of subjects at least once) with probabilities typically below 50% (85% of all randomized risk-taking actions).

Overall, we find that most subjects take a fair and symmetric binary risk often. Sequential risk-taking is largely buy-and-hold, with little evidence for moderators of this behavior. The average stopping probability decreases slightly with horizon.

IV.B Dynamic consistency

We now analyze if and how subjects deviate from their risk-taking plan when taking the risk sequentially. We split our within-subject analysis of dynamic consistency into two parts: First, we analyze subjects' deviation frequency. Second, we analyze deviation directions (i.e., to which risk-taking actions subjects deviate). Throughout the analysis of dynamic consistency, we focus on deviations along paths that are walkable under the precommitted risk-taking strategy, such that we disregard deviations that can only be observed after an earlier deviation within the same Walking Task and for which the plan has not been incentivized. We find (a) a high degree of dynamic consistency in that subjects follow their unconstrained risk-taking plan for around 86% of their risk-taking actions and (b) deviations, if observed, are mostly toward stopping in the gain domain and toward continuation in the loss domain.

[25] The average risk-taking action for the second and third repetition of the risk is 14% and 16%, respectively. For the first three repetitions, the average stopping probability is significantly higher for subjects that can randomize than for subjects that cannot randomize ($t(3{,}082) = 17.8383, t(3{,}574) = 6.4094, t(3{,}244) = 2.7103$), and insignificantly different afterwards (19% with $t(2{,}729) = 1.5419$ for the fourth repetition, and 19% with $t(2{,}283) = 1.0995$ for the fifth repetition). Degrees of freedom are estimated using the Welch–Satterthwaite equation.

Our measure of a subject's dynamic consistency (C_i) is inspired by Fischbacher, Hoffmann, and Schudy's (2017) measure of dynamic consistency as well as by Odean's (1998) measure of the disposition effect. We define the consistency of subject i as

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

C_i equals 1 if subject i always follows her plan, and 0 if she never follows her plan.[26] 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.

Figure 15 shows that subjects mostly follow their previously elicited plan when taking the risk sequentially. The distributions of the dynamic consistency measure C_i in Panel (a) of Figure 15 are concentrated above 0.70 in all studies. The average subject follows her risk-taking plan in 86.3% of all sequential risk-taking actions (median 90.9%), with a median dynamic concistency C_i above 0.89 for all studies (see Table C. 1 in Appendix C). For the median subject, dynamic consistency is not affected by the usage of path-dependence or randomization in the risk-taking plans, nor by randomization during the Walking Tasks. Although the median dynamic inconsistency is identical across studies, allowing for randomization leads to more heterogeneity in the dynamic consistency across subjects.[27]

[26] We similarly define dynamic consistency in the gain (loss) domain as the number of plan executions in the gain (loss) domain divided by the number of deviation opportunities in the gain (loss) domain. Note that our consistency measure does not take the strength of deviations into account, similar to the measures of Odean (1998) and Fischbacher, Hoffman, and Schudy (2017).

[27] Subjects that can use an independent randomization device are, on average, less dynamically consistent than subjects that cannot randomize (average C_i of 0.81 versus 0.91, $t(398) = 5.5547$). However, we do not observe differences in the median consistency for subjects that can or cannot randomize, such that the higher average C_i in the *Randomization* studies is due to outliers, as is evident from Panel (a) of Figure 15.

FIGURE 15.—Dynamic consistency across studies

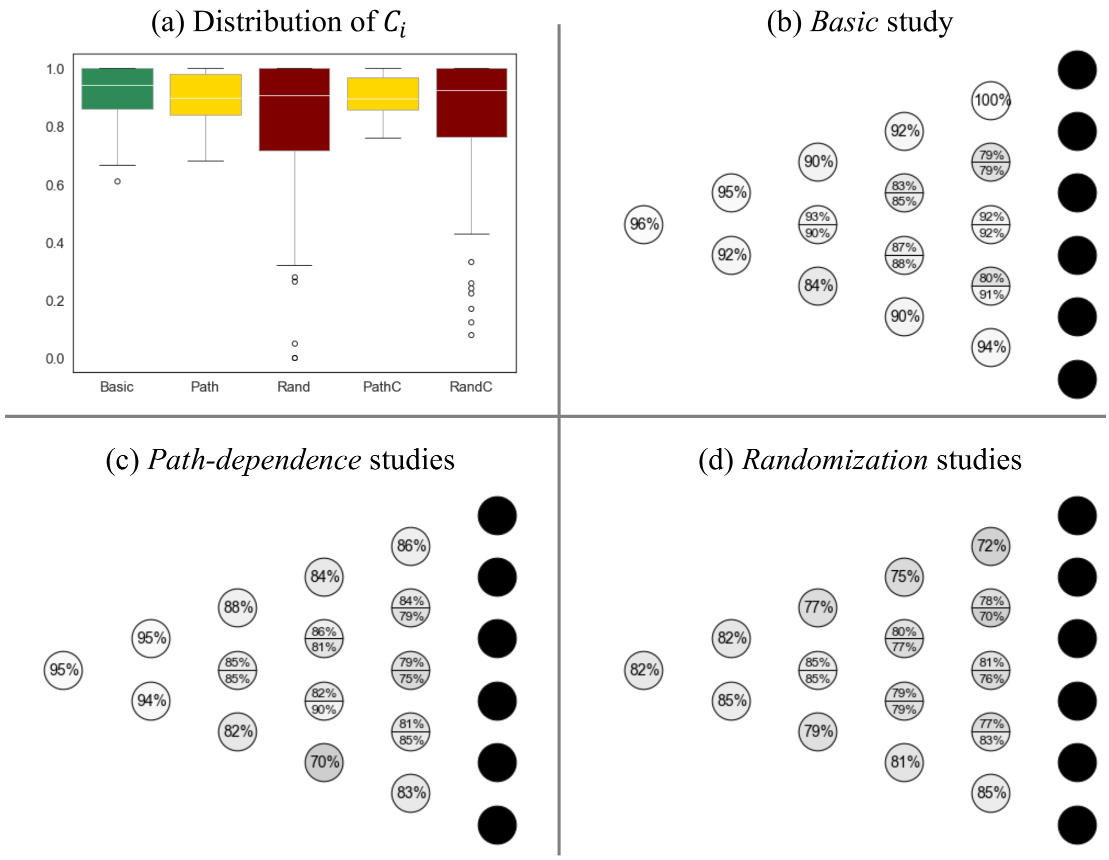

Notes: Panel (a) of Figure 15 shows the distribution of our dynamic consistency measure C_i, while panels (b) – (d) show the average shares of consistent risk-taking actions per ball for $N = 73$ $(142, 185)$ subjects in studies *Basic* (*Path-Dependence*, *Randomization*), respectively. We pool subjects for which the usage of path-dependence (randomization) is costless or costly.

Panels (b) – (d) show the share of consistent risk-taking actions at each ball in the *Basic*, *Path-Dependence*, and *Randomization* studies. Across all studies, the ballwise share of consistent risk-taking actions is above 70%. We do not observe differences in the ballwise dynamic consistency between *Basic* and *Path-Dependence* studies, but risk-taking is—on average—less dynamically consistent in *Randomization* studies (darker balls), especially at early balls. We find that dynamic consistency decreases with the repetition of the risk (darker balls towards the right), as 89% (89%, 84%, 81%, 80%) of risk-taking actions concerning the first (second, third, fourth, fifth) risk are dynamically consistent when pooling all subjects.[28]

[28] Further analysis shows that dynamic consistency does not vary much across gain or loss nodes, whether the previous risk was lost or won, or with the number of the Walking Task, see results of a logit-regression in Table C. 2. Task fixed-effects are insignificant in all specifications, such that we do not observe any order effects.

Overall, subjects largely follow their unconstrained risk-taking plan when taking the risk sequentially. If subjects can freely choose their risk-taking plan (as they can in our experiment), subsequent deviations are rare.

As a next step, we analyze the deviation directions of subjects that deviate from their risk-taking plan. We measure a *directed deviation* as the difference between the actual stopping probability (Walking Tasks) and the planned stopping probability (Decoloring Task), by $\Delta p = \text{Actual stopping probability} - \text{Planned stopping probability}$. Deviations to stopping have a positive sign ($\Delta p > 0$) and deviations to continuation have a negative sign ($\Delta p < 0$).

FIGURE 16.—Directed deviations at gain and loss balls

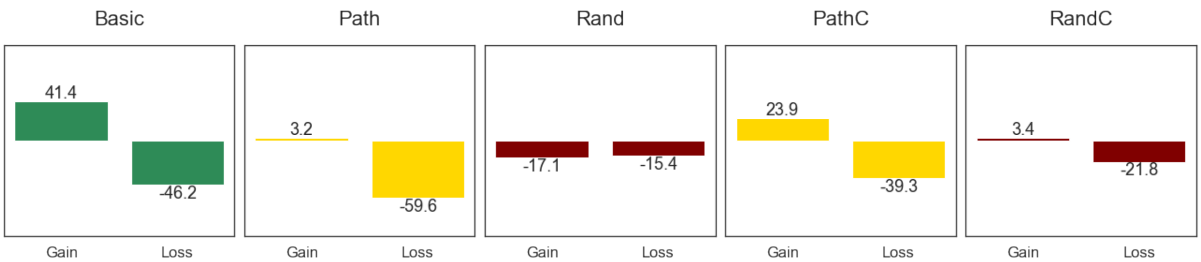

Notes: Figure 16 shows the average deviation Δp at gain and loss balls for each study.

On average, subjects deviate to continuation. Table C. 3 in the appendix shows that the average deviation Δp is negative for all studies except *Basic*, for which it is insignificantly positive. Figure 16 decomposes the directed deviation Δp into the average directed deviation at gain and loss balls. For all studies except *Randomization*, we find a higher average directed deviation Δp at gain balls than at loss balls, meaning that subjects deviate (more) to stopping than to continuing at gain balls than at loss balls (Table C. 3 reports t-Statistics). Indeed, subjects in all studies deviate to continuation at loss balls. In contrast, subjects in studies *Basic* and *Path-dependence*^{C} deviate to stopping at gain balls, while the deviation at gain balls is insignificantly different from zero for subjects in studies *Path-dependence* and *Randomization*^{C}. Subjects in our *Randomization* study deviate to risk-taking at gain balls.[29]

Figure B. 4 in the appendix further shows the average ballwise directed deviation Δp for each study and confirms the pattern observed above. Deviating subjects in study *Basic* deviate towards stopping for all gain balls, with an average ballwise directed deviation Δp ranging

[29] Since we observe that subjects are dynamically consistent at most balls, we also find that the type of the actual and reverse decoloring of subjects coincide for most subjects, such that subjects that plan to use a loss-exit (gain-exit, symmetric-exit) strategy also take the risk following a loss-exit strategy (gain-exit, symmetric-exit).

from $+12\%$ to $+75\%$. At loss balls, subjects deviate towards continuation with an average ballwise directed deviation Δp between -40% and -50%. Moreover, after losing the risk thrice (i.e., at the lowest ball in the fourth column with a loss of £-0.75), all four deviations are from stopping to continuation.[30] Similar patterns hold in the *Path-dependence* and *Randomization* studies, whereby subjects sometimes (*Path-dependence* studies) or always (*Randomization* studies) deviate to continuation at gain balls.

V. Discussion

Experimental design. We have proposed and implemented a method to elicit unconstrained stopping times for repeated risk-taking. Subjects design a risk-taking strategy under precommitment by decoloring the balls of a ball triangle that represents a standard symmetric random walk with $T = 5$ time steps (Decoloring Task). In addition to standard, time- and outcome-contingent risk-taking, we also allow some subjects to make use of path-dependence or randomization, as motivated by theory (Hu et al. 2022). The visualization as a ball triangle allows us to parsimoniously elicit risk-taking strategies without constraining subjects' strategies to certain strategy types (such as buy-and-hold or threshold strategies).[31] 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. The Decoloring Task thus constitutes an incentivized method for the elicitation of unconstrained stopping times. The incentive-compatibility of the Decoloring Task roots in the strategy method (Selten 1967).

We thereby contribute to earlier work that focused on sequential risk-taking without the prior elicitation of risk-taking strategies (e.g., Thaler and Johnson, 1990; Imas, 2016; Nielsen, 2019; Strack and Viefers, 2021) or on constrained strategy spaces (Fischbacher, Hoffmann, and Schudy 2017; Dertwinkel-Kalt and Frey 2022; Heimer et al. 2023, Antler and Arad, forthcoming). Perhaps most closely related is the experimental design by Magnani et al. (2022), who elicit unconstrained stochastic control plans to study the efficiency of dynamic portfolio

[30] We only observe one deviation after losing the risk four times, and this subject deviates from a plan to continue risk-taking to stopping, leading to $\Delta p = +100$ at the bottom ball in the fifth column.

[31] Subjects can either distinguish between 15 balls (*Basic* and *Randomization* studies) or 31 paths (*Path-dependence* studies). On average, subjects need less than three minutes to complete the Decoloring Task (124s in *Basic*, 177s in *Randomization*, and 190s in *Path-dependence* studies).

choice. Subjects in their experiment enter the wealth invested into the asset at each node of a binomial tree, proceeding sequentially through the binomial tree.[32]

Our Decoloring Task may have some advantages over a classical implementation of the strategy method for repeated risk-taking. Asking a subject what she would do in each possible contingency (like, "what would you do if you win twice and lose once?") seems more abstract. The order in which contingencies are presented may prime subjects, thereby affecting decisions, or be unnatural (e.g., if the contingencies were shown in a randomized order). Eliciting subjects' planned actions in a pre-specified order feels close 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 of actions by themselves when designing their strategy while also viewing each action as part of their complete risk-taking strategy. The order in which subjects make their choices for each contingency is therefore endogenous and each planned stopping/continuation action is visualized as part of the whole strategy, rather than being shown in isolation and one after the other.

In our implementation of the Decoloring Task, all risk-taking is about the simple risk of winning or losing a given amount with equal probability. In this first study of unconstrained stopping times under precommitment, we focussed on the dynamic aspects of risk-taking. Just like the large literature on static choice under risk abstracts from dynamic aspects, we abstracted from static aspects as much as possible by implementing the simplest risk that one may think of. We note that one could adapt the decoloring framework to study more general risks such as non-fair risks (e.g., Antler and Arad forthcoming), skewed risks (Ebert 2020), a multiplicative random walk (Strack and Viefers 2021), or non-stationary risks with time-changing moments (Crosetto and Filippin 2013; Nielsen 2019).[33] 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), to a consumer search problem (Stahl 1989), or other related applications.

---

[32] The total wealth available to subjects is changing with the outcomes of the underlying risks and with their investment decisions in the experiment by Magnani et al. (2022). Therefore, subjects proceed sequentially through the binomial tree either by forward or backward induction.

[33] In the dynamic version of the bomb risk elicitation task (BRET; Crosetto and Filippin 2013), subjects stop a process in which boxes are opened. If a box contains a "bomb," subjects' receive no earnings. Opening more boxes gives higher potential earnings, but increases the probability of opening a box that contains a "bomb." The moments of the outcome distribution thus change over time, a feature used by Nielsen (2019).

Strategy types. Although our experimental method allows for hundreds of behaviorally different risk-taking strategies, we find that the elicited risk-taking strategies map well into strategy classes previously analyzed and discussed in the theoretical stopping literature (Shiryaev 2007; Björk et al. 2021; He and Zhou 2022). Using an unsupervised K-means clustering, we identify clusters that can be described as stop-loss, buy-and-hold, never-start, and take-profit strategies. The unsupervised algorithm—that could identify arbitrary patterns in the data—maps unconstrained risk-taking strategies into interpretable strategy types. Our findings thus support theoretical and empirical work that restricts attention to these strategy types without considering general unconstrained risk-taking strategies. Depending on the setting and the puropose of a given study, these strategies may approximate repeated risk-taking behavior reasonably well.

Loss-exit strategies. Most subjects take the risk often, continue risk-taking when winning and stop risk-taking when losing. Less than ten percent of subjects do not start risk-taking, which is incompatible with risk-averse expected utility preferences (as the optimal stopping theorem shows that the expected profit induced by any risk-taking strategy is zero). Across all studies, we find that the relative majority of subjects' risk-taking strategies are loss-exit strategies in which subjects plan on stopping earlier when losses accumulate than when gains accumulate. Barberis (2012) shows theoretically that a CPT-agent takes risks repeatedly following a loss-exit strategy for a wide range of parameters. Intuitively, a loss-exit strategy induces a positively skewed outcome distribution, since it retains a small probability of a comparably large gain while limiting losses. A CPT-agent who overweights extreme outcomes that occur with small probability finds positively skewed risks attractive. Our results are consistent with empirical and experimental findings by Heimer et al. (2023) and Dertwinkel-Kalt and Frey (2022), who also observe frequent risk-taking and a majority of loss-exit strategies when subjects choose among threshold strategies.[34] We show that loss-exit strategies emerge endogenously even for unconstrained risk-taking strategies.

Threshold or trailing stop-loss. Allowing subjects to freely design any risk-taking strategy further allows us to analyze what *kind* of loss-exit strategy subjects choose. We find a prevalence of trailing stop-loss strategies over threshold stop-loss strategies. Trailing stop-loss

[34] Antler and Arad (forthcoming) find that subjects choose threshold strategies with a larger (in absolute terms) loss limit than gain limit, inducing gain-exit strategies rather than loss-exit strategies. The key difference to our study and the study by Heimer et al. (2023) is that the single risks used by Antler and Arad (forthcoming) are not fair and that subjects choose among a predefined set of of five threshold strategies. Antler and Arad (forthcoming, p. 9) provide a detailed discussion on the matter.

strategies are well known in the theory literature (often called Azéma-Yor stopping times). On an individual level, the ratio of trailing to threshold stop-loss strategies is roughly 3.5:1, ranging from 1.7:1 (*Randomization* study) to 4.8:1 (*Path-dependence*^{C} study). The relative evidence for trailing over threshold stop-loss may be important for brokers and policy-makers who wish to promote trading and commitment to loss limits.[35]

An important insight from our experiment is that subjects choose trailing stop-loss orders endogenously—without any restriction on the admissible strategies or without any suggestion given to subjects. The restriction of theoretical, empirical, or experimental analyses to threshold stop-loss strategies may thus, for example, force subjects to choose a risk-taking strategy that they would not have chosen otherwise. Moreover, in some theories of decision-making under risk, the attractiveness of any given strategy depends on the alternatives offered, such that exluding options from the choice set may have non-trivial effects on observed behavior.

Path-dependence. Theoretical work shows that a precommitted CPT-agent can benefit from using path-dependence or randomization since both options allow her to induce more positively skewed outcome distributions (He et al. 2017; Hu et al. 2022). Consistent with the theoretical prediction, 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 et al. 2022), the induced outcome distributions in the *Path-dependence* studies are not more positively skewed than in the *Basic* study. Moreover, subjects use path-dependent risk-taking 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) 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) stochastic control 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

---

[35] 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 (2020) shows 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.

knowing which path will realize, such that the outcome of the path-dependent risk-taking plan at any node is random at the time of designing the stopping time (see also the theoretical work by He et al. 2017; Hu et al. 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 risk-taking if it was up-down (or vice versa). The ex-ante probability of stopping at the path-dependent middle ball in the third colum is then 50%. This motivation for path-dependence in the dynamic case is similar to the motivation to randomize static decisions identified by Agranov, Healy, and Nielsen (forthcoming). In their experiment, a substantial fraction of subjects randomizes by varying their static choice over the course of 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. Consistent with our interpretation of path-dependence as a means to achieve randomization, we observe a higher demand for direct randomization than for path-dependence.

Randomization. We also provide an experimental analysis of randomization in a dynamic decision-making task, both for risk-taking under precommitment and for sequential risk-taking. 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 et al. 2018; Agranov and Ortoleva 2022),[36] but there are additional motives for randomization in the dynamic context. Randomized stopping times are the rule rather than the exception in theoretical work (Henderson et al. 2017; Hu et al. 2022). Although randomization is often used 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 than in the other studies. Second, subjects' stopping probabilities at gain nodes are close to the stopping probabilities at loss nodes. This behavior results in relatively more gain-exit strategies in the *Randomization* studies than in the other studies. Randomization is thus not systematically used to generate skewness (Henderson et al. 2017). Third, subjects typically randomize with stopping probabilities below

[36] 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 (forthcoming) observe that between 64% and 76% of subjects mix in at least one domain. 80% of subjects randomize at least once in our studies.

50%, which is markedly different from a ‘pull to the center’ effect. We currently have no explanation for this finding.

Dynamic consistency. In our experiment, subjects also take the risk sequentially without commitment. After we elicit unconstrained risk-taking strategies, subjects visually walk through their decolored ball triangle and can freely deviate from their risk-taking plan at any node along the realized path. Our within-subject design allows us to compare each subjects’ risk-taking plan to her risk-taking actions. We find that subjects follow their unconstrained risk-taking plan for 86% of their sequential risk-taking actions, thus showing an absolutely high degree of dynamic consistency (also relative to studies using threshold strategies, as discussed in the next paragraph). We noted earlier that many risk-taking strategies are trailing stop-loss strategies—which are not available in experiments focusing on threshold strategies. Subjects in those experiment thus had to choose a constrained optimal risk-taking plan to begin with, which may raise the frequency of deviations from the (constrained optimal) risk-taking plan when allowing for unconstrained risk-taking actions. We further find that the additional flexibility of the strategy or of sequential risk-taking introduced by path-dependence or randomization does not increase dynamic consistency above and beyond the flexibility offered by an unconstrained risk-taking plan. Nielsen and Rehbeck (2022) show that subjects are more likely to follow a canonical (static) choice axiom if they previously selected whether they want to follow this axiom or not. In combination with our results, the emerging pattern suggests that individuals are more likely to adhere to a plan, rule, or norm if they could previously choose the rule or express agreement therewith. By implementing constrained versions of the Decoloring Task and comparing the dynamic consistency in a constrained and unconstrained version of our experiment, one can quantify the effect of various restricted choice sets on dynamic consistency.

Deviation directions. Subjects who deviate from their risk-taking plan tend to take the risk more often than planned. Deviations to continuation are concentrated in the loss domain, while inconsistent subjects tend to stop risk-taking somewhat earlier in the gain domain. Stopping earlier at gain balls and continuing at loss balls is consistent with theoretical predictions (Barberis 2012; Ebert and Strack 2015; Hu et al. 2022), and with prior empirical as well as experimental findings (Barkan and Busemeyer 1999; Andrade and Iyer 2009; Magnani 2015; Imas 2016; Fischbacher et al. 2017; Dertwinkel-Kalt and Frey 2023; Heimer et al. 2023). In particular, Heimer et al. (2023) observe that threshold risk-taking strategies are of a loss-exit type, while sequential risk-taking behavior observed between-subject is of a gain-exit type. We find that within-subject—and despite a salient reminder of the subjects’ risk-taking strategy—

deviating subjects still cut gains and chase losses, thus corroborating the results in Heimer et al. (2023). Given the higher dynamic consistency of our subjects, we find overall less reversals of loss-exit plans to gain-exit actions. Subjects mostly follow their unconstrained risk-taking plan. This result speaks to the possibility that offering commitment to more general risk-taking strategies, in particular to trailing-stop-loss strategies, may help individuals to stick to their plan.

CONCLUSION

We propose a method for eliciting stopping times—each subjects' completely contingent plan for taking a risk repeatedly. We apply the method to study repeated risk-taking strategies under precommitment and focus on the simple risk of winning or losing a given amount with equal probability for up to five times. Although our experimental design allows for hundreds of behaviorally different risk-taking strategies, we find that individual risk-taking strategies map well to simple and interpretable strategy types, such as buy-and-hold and stop-loss strategies. Without any restriction on the admissible risk-taking strategies—and by applying an unsupervised algorithm that did not learn prespecified strategies—we find that individuals endogenously choose risk-taking strategies considered in earlier theoretical and empirical research and that are observed in the real world. Subjects plan on taking the risk often and design strategies that feature continue-when-winning and stopping-when-losing. Among these loss-exit strategies, trailing stop-loss strategies are used 3.5 times as often as threshold stop-loss strategies. Moreover, path-dependence and randomization are both used often, even if they are costly. The usage of path-dependence is more consistent with a revealed demand for randomization than with a direct preference for path-dependence, and—in line with this interpretation—we observe more usage of randomization than of path-dependence. Randomized stopping probabilities are thereby usually below 50%, shifting risk-taking strategies closer to buy-and-hold strategies than without randomization. We also analyze sequential risk-taking and find that subjects follow their plan for about 86% of their sequential risk-taking actions. Subjects show a high degree of dynamic consistency if their plan was elicited without any constraint.

Overall, subjects endogenously choose stop-loss plans even if they can freely design any risk-taking strategy. This observation supports theory and indicates that the restriction to these strategies consitutes a reasonable approximation to actually desired strategies. At the same time, allowing subjects to choose unconstrained risk-taking strategies delivers further insight into what subjects precisely want. We find that many subjects want path-dependence, randomization, and the trailing version of stop-loss strategies, and offering (more) flexible strategies may increase individuals' ability to stick to their plan beyond commitment devices that focus on restricted plans.

References

Agranov M, Healy PJ, Nielsen K (forthcoming) Stable randomization. Economic Journal

Agranov M, Ortoleva P (2022) Revealed Preferences for Randomization: An Overview. In: AEA Papers and Proceedings. pp 426–430

Agranov M, Ortoleva P (2017) Stochastic choice and preferences for randomization. Journal of Political Economy 125:40–68

Andrade EB, Iyer G (2009) Planned versus actual betting in sequential gambles. Journal of Marketing Research 46:372–383

Antler Y, Arad A (forthcoming) An Experimental Analysis of the Prize-Probability Tradeoff in Stopping Problems. Review of Economics and Statistics

Arthur D, Vassilvitskii S (2007) K-means++ the advantages of careful seeding. In: Proceedings of the eighteenth annual ACM-SIAM symposium on discrete algorithms. pp 1027–1035

Barberis N (2012) A model of casino gambling. Management Science 58:35–51

Barkan R, Busemeyer JR (1999) Changing plans: dynamic inconsistency and the effect of experience on the reference point. Psychonomic Bulletin & Review 6:547–554

Bell DE (1982) Regret in decision making under uncertainty. Operations research 30:961–981

Bishop CM (2006) Pattern recognition and machine learning. Springer

Björk T, Khapko M, Murgoci A (2021) Time-Inconsistent Control Theory with Finance Applications. Springer

Bordalo P, Gennaioli N, Shleifer A (2012) Salience Theory of Choice Under Risk. Quarterly Journal of Economics 127:1243–1285

Chen DL, Schonger M, Wickens C (2016) oTree—An open-source platform for laboratory, online, and field experiments. Journal of Behavioral and Experimental Finance 9:88–97

Crosetto P, Filippin A (2013) The "bomb" risk elicitation task. Journal of Risk and Uncertainty 47:31–65

Cubitt R, Starmer C, Sugden R (1998) Dynamic choice and the common ratio effect: An experimental investigation. Economic Journal 108:1362–1380

Dertwinkel-Kalt M, Frey J (2023) Optimal Stopping in a Dynamic Salience Model. Working paper available at SRRN

Dixit AK (1993) The art of smooth pasting. Taylor & Francis

Dixit AK, Pindyck RS (1994) Investment under uncertainty. Princeton University Press

Douglas BD, Ewell PJ, Brauer M (2023) Data quality in online human-subjects research: Comparisons between MTurk, Prolific, CloudResearch, Qualtrics, and SONA. PLoS ONE 18

Dwenger N, Kübler D, Weizsäcker G (2018) Flipping a coin: Evidence from university applications. Journal of Public Economics 167:240–250

Ebert S (2020) On Taking a Skewed Risk More Than Once. Working paper available at SRRN

Ebert S, Karehnke P (2022) Skewness Preferences in Choice under Risk. Working paper available at SRRN

Ebert S, Strack P (2015) Until the bitter end: on prospect theory in a dynamic context. American Economic Review 105:1618–33

Feldman P, Rehbeck J (2022) Revealing a preference for mixtures: An experimental study of risk. Quantitative Economics 13:761–786

Fischbacher U, Hoffmann G, Schudy S (2017) The causal effect of stop-loss and take-gain orders on the disposition effect. Review of Financial Studies 30:2110–2129

Gneezy U, Potters J (1997) An experiment on risk taking and evaluation periods. Quarterly Journal of Economics 112:631–645

Grosshans D, Zeisberger S (2018) All's well that ends well? On the importance of how returns are achieved. Journal of Banking & Finance 87:397–410

Gul F (1991) A Theory of Disappointment Aversion. Econometrica 59:667–686

Gupta N, Rigotti L, Wilson A (2021) The experimenters' dilemma: Inferential preferences over populations. Working paper available at arXiv

Haaland I, Roth C, Wohlfart J (2023) Designing information provision experiments. Journal of Economic Literature 61:3–40

He XD, Hu S, Obłój J, Zhou XY (2017) Path-dependent and randomized strategies in Barberis' casino gambling model. Operations Research 65:97–103

He XD, Zhou XY (2022) Who are I: Time inconsistency and intrapersonal conflict and reconciliation. In: Stochastic Analysis, Filtering, and Stochastic Optimization. Springer, pp 177–208

Heimer R, Iliewa Z, Imas A, Weber M (2023) Dynamic Inconsistency in Risky Choice: Evidence from the Lab and Field. Working paper available at SRRN

Henderson V, Hobson D, Tse A (2018) Probability weighting, stop-loss and the disposition effect. Journal of Economic Theory 178:360–397

Henderson V, Hobson D, Tse A (2017) Randomized strategies and prospect theory in a dynamic context. Journal of Economic Theory 168:287–300

Horton JJ, Rand DG, Zeckhauser RJ (2011) The online laboratory: Conducting experiments in a real labor market. Experimental Economics 14:399–425

Hu S, Obłój J, Zhou XY (2022) A Casino Gambling Model Under Cumulative Prospect Theory: Analysis and Algorithm. Management Science 69:2474–2496

Imas A (2016) The realization effect: Risk-taking after realized versus paper losses. American Economic Review 106:2086–2109

Johnson C, Baillon A, Bleichrodt H, et al (2021) Prince: An improved method for measuring incentivized preferences. Journal of Risk and Uncertainty 62:1–28

Kahneman D, Tversky A (1979) Prospect Theory: An Analysis of Decision under Risk. Econometrica 47:263–91

Karni E, Safra Z (1990) Behaviorally consistent optimal stopping rules. Journal of Economic Theory 51:391–402

Langer T, Weber M (2008) Does commitment or feedback influence myopic loss aversion?: An experimental analysis. Journal of Economic Behavior & Organization 67:810–819

Liu Y-J, Tsai C-L, Wang M-C, Zhu N (2010) Prior consequences and subsequent risk taking: New field evidence from the Taiwan Futures Exchange. Management Science 56:606–620

Loomes G, Sugden R (1982) Regret theory: An alternative theory of rational choice under uncertainty. Economic Journal 92:805–824

Machina MJ (1989) Dynamic consistency and non-expected utility models of choice under uncertainty. Journal of Economic Literature 27:1622–1668

Magnani J (2015) Testing for the disposition effect on optimal stopping decisions. American Economic Review (Papers & Proceedings) 105:371–75

Magnani J, Rabanal JP, Rud OA, Wang Y (2022) Efficiency of dynamic portfolio choices: An experiment. Review of Financial Studies 35:1279–1309

Merton RC, Samuelson PA (1992) Continuous-time finance

Nielsen K (2019) Dynamic risk preferences under realized and paper outcomes. Journal of Economic Behavior & Organization 161:68–78

Nielsen K, Rehbeck J (2022) When choices are mistakes. American Economic Review 112:2237–68

Odean T (1998) Are investors reluctant to realize their losses? Journal of Finance 53:1775–1798

Quiggin J (1982) A theory of anticipated utility. Journal of Economic Behavior & Organization 3:323–343

Rabin M (2002) Inference by believers in the law of small numbers. Quarterly Journal of Economics 117:775–816

Redelmeier DA, Tversky A (1992) On the framing of multiple prospects. Psychological Science 3:191–193

Selten R (1967) Die Strategiemethode zur Erforschung des eingeschränkt rationalen Verhaltens im Rahmen eines Oligopolexperiments. Beiträge zur experimentellen Wirtschaftsforschung 136–168

Shefrin H, Statman M (1985) The disposition to sell winners too early and ride losers too long: Theory and evidence. Journal of Finance 40:777–790

Shiryaev AN (2007) Optimal stopping rules. Springer

Snowberg E, Yariv L (2021) Testing the waters: Behavior across participant pools. American Economic Review 111:687–719

Stahl DO (1989) Oligopolistic pricing with sequential consumer search. American Economic Review 700–712

Strack P, Viefers P (2021) Too proud to stop: Regret in dynamic decisions. Journal of the European Economic Association 19:165–199

Thaler RH, Johnson EJ (1990) Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice. Management Science 36:643–660

Tversky A, Kahneman D (1992) Advances in prospect theory: Cumulative representation of uncertainty. Journal of Risk and Uncertainty 5:297–323

Tversky A, Shafir E (1992) The disjunction effect in choice under uncertainty. Psychological Science 3:305–310

Xu ZQ, Zhou XY (2013) Optimal Stopping Under Probability Distortion. Annals of Applied Probability 23:251–282

Yaari ME (1987) The dual theory of choice under risk. Econometrica 95–115

Appendix A. Experimental Details

The experiment and the associated instructions and comprehension questions are interactive. A tourist version of all five studies of the experiment will be made available. We now describe the most important aspects of the experimental design and its implementation.

A.1. Details on randomization

Before the start of the experiment, we exogenously and randomly assign subjects to the different studies *Basic*, *Path-dependence*, *Path-dependence*^{C}, *Randomization*, and *Randomization*^{C} in the (random) order of joining the experiment, and we give priority to studies with fewer complete or currently active subjects.

During the experiment, we implement different randomization procedures for (a) the selection of the payoff-relevant task, (b) the resolution of the fair, binary, and symmetric risk, and (c) the randomized stopping decision.

At the beginning of the experiment, subjects click one of eleven buttons to determine their payoff-relevant task. We implement a virtual coin toss to resolve the fair, symmetric and binary risk during the experiment. For all binary risks, subjects select their "winning side" for a coin toss as if asked to bet on "Heads" or "Tails" (see Figure A.1). Upon clicking the relevant button, we show subjects the outcome of the coin toss. If the coin toss outcome equals the chosen winning side, the risk is won. Otherwise, the subject loses the risk (see Figure A.1).

FIGURE A.1.—Resolution of the single risk

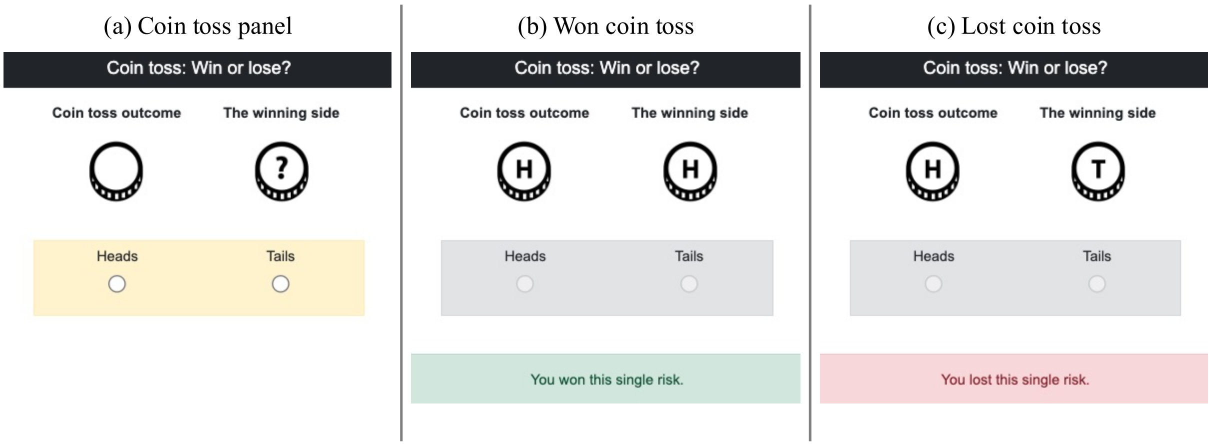

Notes: We resolve the single risk using a virtual coin toss that the subjects wins or loses with equal probability. Panel (a) shows the coin toss visualization and the input panel. The subject must choose the winning side, highlighted with the questionmark. If the coin toss outcome equals the winning side as chosen by the subject, she wins the single risk as in panel (b). If the coin toss outcome does not equal the winning side chosen by the subject, she loses the single risk as in panel (c).

Subjects in the *Randomization* studies can stop risk-taking with any probability on the grid 0, 10, ... , 90, 100. We saliently implement a different resolution mechanism—a card draw instead of a coin toss—to visually emphasize that the randomization device is independent of the single risk. The virtual card deck consists of ten white or black card (see Figure A.2). Risk-taking stops if the subject draws a black card, and continues if she draws a white card. The number of black cards in the card deck is such that the probability of drawing a black card equals the stopping probability. For example, if the subject stops risk-taking with 60 (20) percent probability, the card deck contains 6 (2) black cards.[37]

FIGURE A.2.—Resolution of randomized stopping

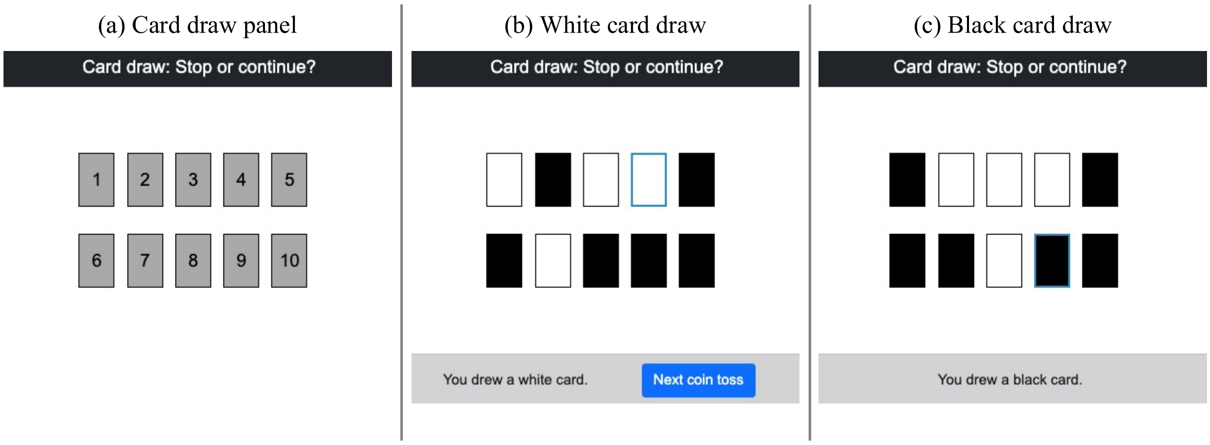

Notes: We resolve the randomized stopping using a virtual card deck with a total of ten cards. Panel (a) shows the input panel. The subject must click on a card. If the drawn card is white, as in panel (b), risk-taking continues. If the drawn card is black, as in panel (c), risk-taking stops.

A.2. Questionnaire and payment

The experiment ends with a short questionnaire and the payment. In the questionnaire, we ask subjects if they (subjectively) understood the experimental instructions and whether they had any issues with the experiment that are related to their internet connection or the display of the experiment. Moreover, we ask subjects if they believe that their bonus was computed fairly. Subjects can answer on a five-level Likert scale.

After the questionnaire, subjects learn their payment and the remaining risks are resolved if the Decoloring Task is payoff-relevant. If the bonus payment is determined by a Walking Task,

[37] The cards are visibly numbered from one to ten, and we assign a winning number to each card. If the winning number of the drawn card is smaller than or equal to the stopping probability, the card is black. Otherwise, the card is white.

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.

A.3. Sample properties

We programmed the experiment using oTree (Chen et al. 2016) and ran it via Prolific on November 16, 2022. Overall, 779 subjects started our experiment. 142 (18%) of them left before the first attention check. Out of the remaining 637 subjects, 66 (10%) failed the first attention check. Of the remaing 571 subjects, 44 (8%) left before the second attention check, and 84 of the remaining 527 subjects (16%) failed the second attention check. Of the remaining 443 subjects, 422 finished the experiment. 22 of these subjects had to be removed because they joined the experiment previously and failed an attention check. We only consider the data of the remaining 400 completes. The subject pool only contains a small share of students (11%), and a majority of subjects are in full-time employment (51%, see Table A. 1).

TABLE A. 1.—Descriptive sample statistics

		N	%
Gender	Male	200	50.00%
	Female	200	50.00%
Nationality	United Kingdom	321	80.25%
	United States	68	17.00%
	Ireland	9	2.25%
	Australia	1	0.25%
	New Zealand	1	0.25%
Student	Yes	44	11.00%
	No	341	85.25%
	N/A	15	3.75%
Employment	Full-Time	203	50.75%
	Part-Time	72	18.00%
	Unemployed	32	8.00%
	Not in paid work	51	12.75%
	Starting a job soon	1	0.25%
	Other	17	4.25%
	n/a	24	6.00%

Notes: Table A. 1 shows the number and share of subjects by gender, nationality, student status, and employment status.

Moreover, Table A. 2 shows that the mean age of our subjects is 39.4 years, and that they are experienced experimental subjects with an average of 870 completed studies on Prolific. On average, subjects took 30.6 minutes to complete the experiment, whereby the fastest subject took 13.5 minutes. The distribution of total earnings of the subjects is symmetric, and centered at the mean and median of £4.50, but we observe mass-points at earnings of £4.25 and £4.75.

TABLE A. 2.— Summary statistics of the sample and experimental procedure

	Mean	Median	Standard Deviation	Max	Min
Age (in years)	39.4	37.0	12.0	76.0	18.0
Duration (in min)	30.6	28.6	10.2	68.1	13.5
Total earnings (in GBP)	4.5	4.5	0.0	5.75	3.25
Studies completed on Prolific	870.4	639.5	864.0	4,836.0	0.0

Notes: Table A. 2 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.

Appendix B. Additional Figures

FIGURE B. 1.—Ballwise mean decolorings

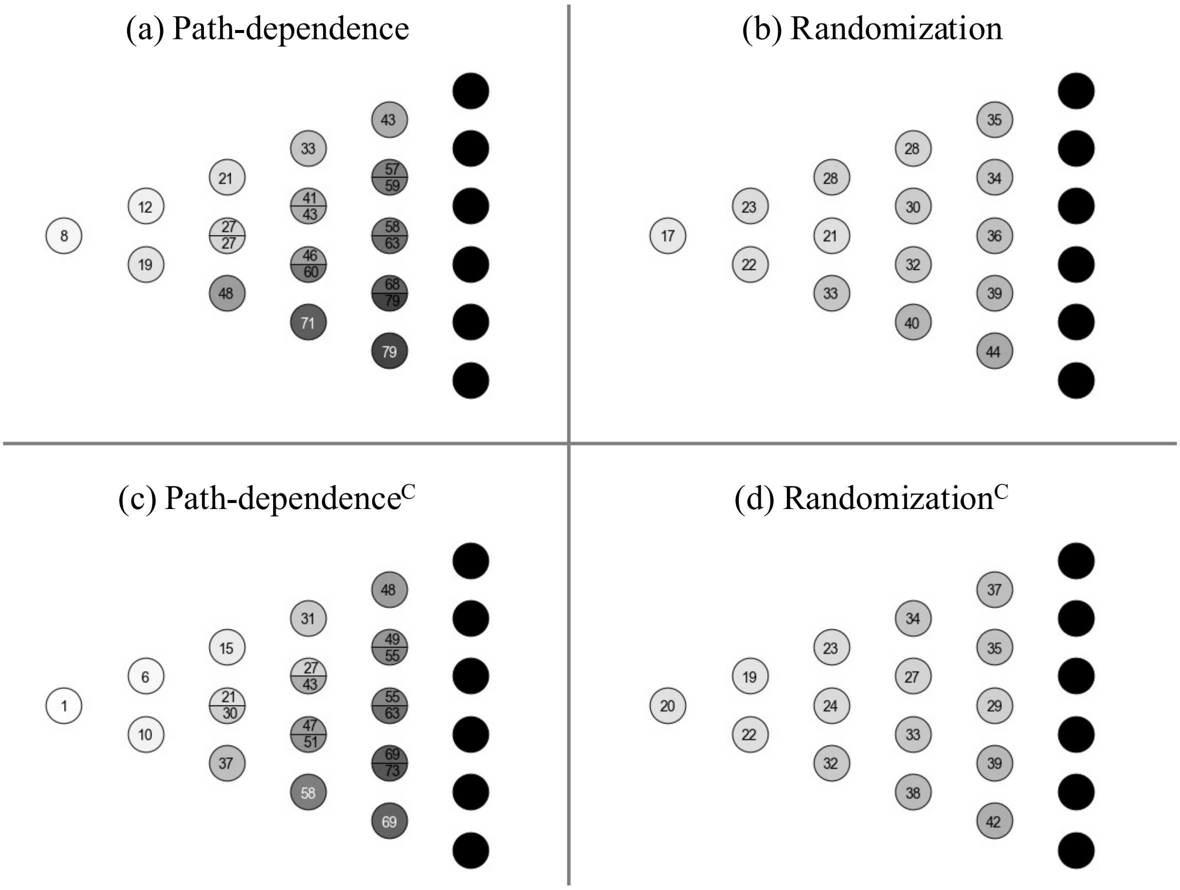

Notes: Figure B. 1 shows the ballwise mean decoloring of subjects in studies *Path-dependence* ($N = 75$), *Randomization* ($N = 88$), *Path-dependence*^{C} ($N = 67$), and *Randomization*^{C} ($N = 97$). The number inside each ball shows the average stopping probability at that node, with a darker greyscale corresponding to a higher stopping probability. We aggregate paths reaching a ball from above or below for *Path-dependence* studies.

FIGURE B. 2.—Prevalence of path-dependent and randomized balls when making these risk-taking options costly

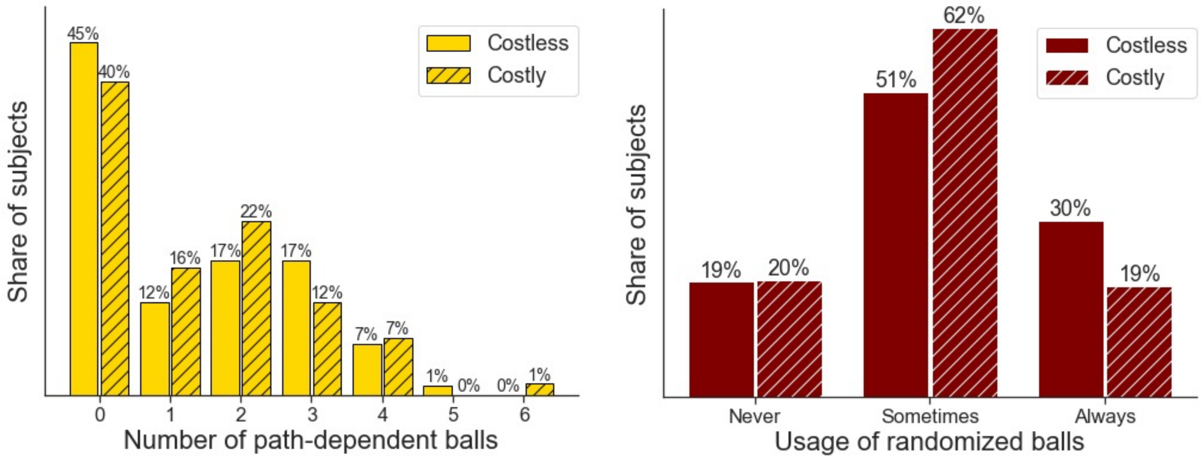

Notes: The left panel of Figure B. 2 shows the share of subjects for each number of path-dependently decolored balls for subjects in studies *Path-dependence* ($N = 75$) and *Path-dependence*^{C} ($N = 67$), while the right panel shows the share of subjects that never, sometimes or always (for all 15 balls) use randomization in studies *Randomization* ($N = 88$) and *Randomization*^{C} ($N = 97$).

FIGURE B. 3.—Prevalence and usage of randomization during the Walking Tasks

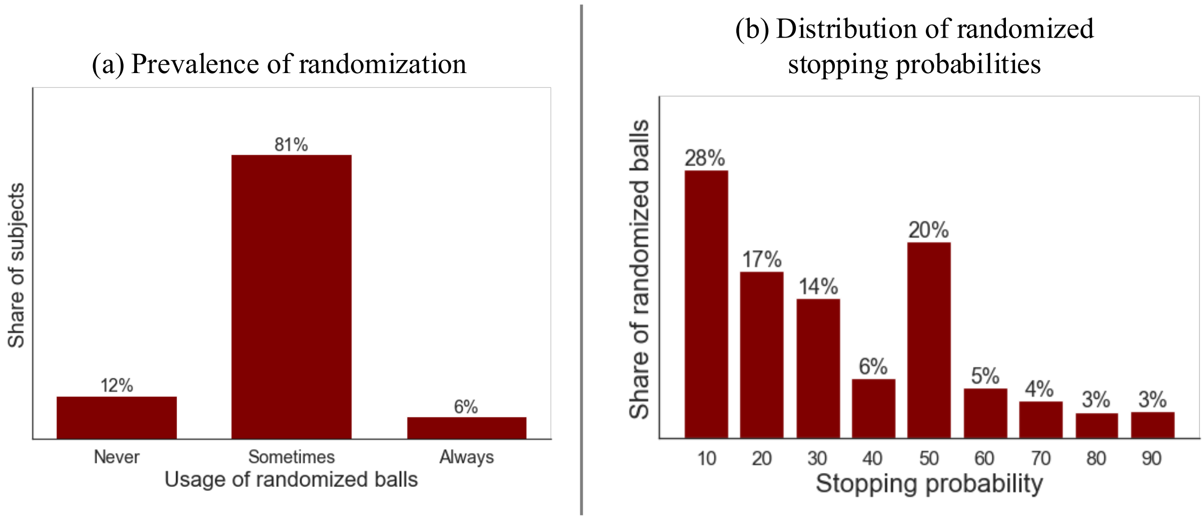

Notes: The left panel of Figure B. 3 shows the share of subjects that never, sometimes, or always use randomization. The right panel shows the share of each randomization probability among all randomized risk-taking actions.

FIGURE B. 4.—Ballwise plan deviation per study

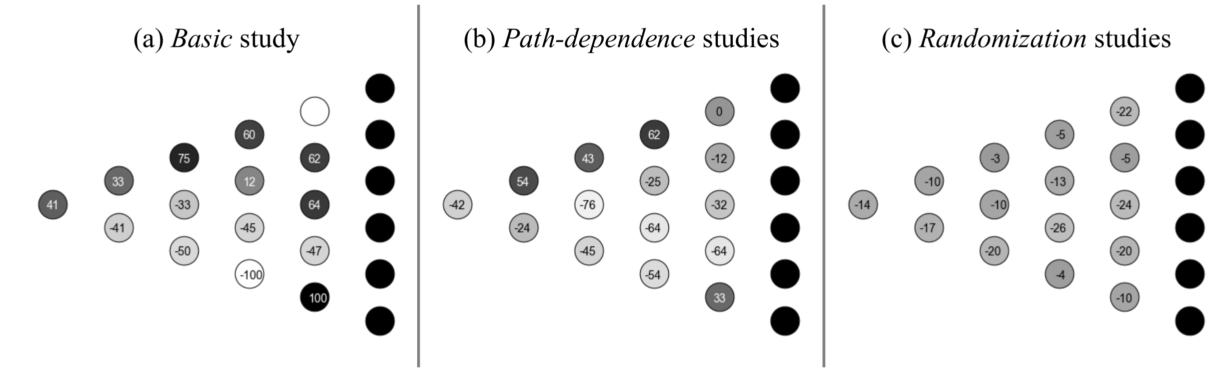

Notes: Figure B. 4 shows the average directed deviation Δp for subjects that deviate from their risk-taking strategy when taking the risk sequentially for each study.

Appendix C. Additional Tables

TABLE C. 1.—Summary statistics of the consistency measure C_i per study

	Basic	Path-dependent	Path-dependent^C	Randomization	Randomization^C
Mean	0.92	0.90	0.80	0.91	0.83
Standard deviation	0.09	0.08	0.26	0.07	0.23
Minimum	0.61	0.68	0.00	0.76	0.08
1st quartile	0.86	0.84	0.72	0.86	0.76
Median	0.94	0.89	0.91	0.90	0.92
3rd quartile	1.00	0.98	1.00	0.97	1.00
Maximum	1.00	1.00	1.00	1.00	1.00

Notes: Table C. 1 shows summary statistics of the dynamic consistency measure C_i. We only consider plan executions and deviation opportunities for nodes along paths that can be reached when following the subjects' precommitted risk-taking strategy.

TABLE C. 2.—Marginal effects of a logit-regression with heteroskedasticity-robust standard errors

	Deviation from risk-taking plan (0 / 1)			
	(1)	(2)	(3)	(4)
Gain node	0.0346		-0.0248	0.0315
	(0.007)		(0.013)	(0.019)
Loss node	0.0272		-0.0187	0.0015
	(0.007)		(0.013)	(0.012)
Risk repetition		0.0262	0.0200	
		(0.002)	(0.003)	
Risk repetition * Gain node			0.0150	
			(0.005)	
Risk repetition * Loss node			0.0102	
			(0.005)	
Last move				0.0132
				(0.015)
Last move * Gain node				-0.0380
				(0.023)
Last move * Loss node				-0.0154
				(0.025)
Observations	15,065	15,065	15,065	11,065
Study fixed-effects	Yes	Yes	Yes	Yes
Task fixed-effects	Yes	Yes	Yes	Yes
Pseudo R^2	0.0236	0.0348	0.0355	0.0148

Notes: Table C. 2 shows the marginal effects of a logit-regression with heteroskedasticity-robust standard errors. The dependent variable is a dummy equal to 1 if the subject deviates from her risk-taking plan and 0 otherwise. The explanatory variables control for whether the node is a gain or loss node, the repetition of the risk, and whether the last risk was won or lost (last move). We control for study and task-number fixed effects. The number of observations varies because we do not observe a prior outcome of the risk for node 0, such that these observations are excluded whenever we control for last move. Robust standard errors are in parentheses.

TABLE C. 3.—Average deviation from the risk-taking strategy elicited during the Decoloring Tasks

	Average Δp at			t-Statistic
	All nodes	Gain nodes	Loss nodes	(degrees of freedom)
Basic	2.72 (0.44, 256)	41.41 (4.50, 98)	-46.24 (-5.00, 92)	6.72 (190)
Path-dependence	-37.27 (-7.30, 321)	3.16 (0.31, 94)	-59.60 (-7.35, 98)	4.81 (192)
Path-dependence^C	-26.00 (-4.66, 299)	23.91 (2.35, 91)	-39.29 (-4.50, 111)	4.74 (202)
Randomization	-15.20 (-6.97, 651)	-17.08 (-4.89, 242)	-15.38 (-3.32, 142)	-0.29 (384)
RandomizationC	-12.12 (-5.41, 616)	3.35 (0.74, 175)	-21.81 (-5.73, 200)	4.27 (375)

Notes: Table C. 3 shows the average directed deviation Δp. A positive number indicates a deviation to stopping, while a negative number indicates a deviation to coninuing. We separately report the average Δp for gain and loss nodes, as well as the results of a t-test comparing both. In each cell, we write the t-statistic and degrees of freedom in parentheses.

TABLE C. 4.— Deviation-count based classification of elicited risk-taking strategies per study

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	11	30	19	2	3	3	5	2.73
Path-dependence	5	21	26	6	-	5	12	4.20
Path-dependence^C	4	19	23	1	2	6	12	4.75
Randomi-zation	3	5	19	2	1	3	55	1.67
Randomi-zation^C	3	8	30	5	4	2	45	2.67
Total	26	83	117	16	10	19	129	3.19

Notes: Table C. 4 shows the assignment of strategies to the classes of the 20 template strategies based on the lowest number of paths for which the elicited risk-taking strategy deviates from a template strategy for each study. We do not assign an elicited risk-taking strategy to a class if the most similar template strategy is not identified.

Appendix D. Machine-learning methodology and robustness

In this appendix, we complement the analysis of Section III.F and III.G. In the first part, we describe the unsupervised machine-learning procedure more closely and conduct robustness analyses that complement Section III.F. In the second part of the appendix, we use different methods to corroborate our findings from Section III.G.

D.1. Unsupervised machine-learning—Individual-level analysis

In Section III.F of the main text, we conduct an individual-level analysis of the elicited risk-taking strategies using a *K*-means clustering algorithm. In this section of the appendix, we first describe the *K*-means clustering and our implementation thereof in more detail. Afterwards, we highlight the robustness of our results by re-performing the analysis of Section III.F with a different number of clusters as well as by reducing the dimensionality of the clustering problem. The findings highlighted in Section III.F—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—400 31-tuples with the precommitted risk-taking action for each subpath of the ball triangle as entries—and partitions the data into $K \in \mathbb{N}^+$ clusters by minimizing the within-cluster variance, while simultaneously maximizing the between-cluster variance. The algorithm selects K random points from the dataset as the 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 thereby known to only converge on local minima (see Bishop 2006), such that the obtained clustering is sensitive to the initial location of the centroids. Whenever we implement the K-means algorithm, we therefore run 100 iterations with random initial centroids and 100 iterations using the k-means++ algorithm (Arthur and Vassilvitskii 2007).[38] Afterwards, we use a majority voting over the 200 iterations of the K-means clustering.

[38] Intuitively, the *k-means++* algorithm does not sample the initial centroids randomly but such that they are away 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.

The number of clusters K is a hyperparameter that needs to be chosen before performing the K-means clustering. Using standard methods[39] to determine the number of clusters, we find that $K = 9$ clusters lead to a partition of the data that provides a good data fit. We showed the analysis in Section III.F of the main text for $K = 6$ clusters because it is more parsimonious and leads to very similar conclusions as the analysis with $K = 9$ clusters.

Figure D. 1 shows the median decoloring for each of the nine clusters, while Table D. 1 shows the number of subjects, average moments of the stopping time τ and induced outcome distribution S_τ, and the number of strategies according to the Barberis (2012)-classification per cluster. Thus, Figure D. 1 replicates Figure 13 from the main text, and Table D. 1 replicates Table 1 with $K = 9$ clusters. The median strategy of clusters 1 and 2 corresponds to a buy-and-hold strategy, while the median strategy of cluster 3 corresponds to a never-start strategy, that of cluster 4 to a take-profit strategy, and that of clusters 5, 6, 7, and 8 to a stop-loss strategy. The median strategy of cluster 9 is more ambiguous and could be interpreted as stop-loss or buy-and-hold strategy. We support the interpretation using Table D. 1.

Combining the findings from Figure D. 1 and Table D. 1 leads to the cluster interpretation as in Table D. 2. Compared to the results in the main text, we find more subjets in clusters that we interpret as buy-and-hold strategies (169 compared to 121) and less subjects in clusters that can we interpret as take-profit strategies (34 compared to 79). The number of subjects in clusters that correspond to stop-loss strategies is almost identical to the main text (165 compared to 163). Overall, our findings from the individual-level analysis—especially the interpretability of unconstrained risk-taking strategies in terms of simpler strategy classes—seems robust to the number of clusters used.

[39] A standard elbow plot of the *intertia* (sum of the squared distance of a datapoint from it's cluster center) is inconclusive, since it is declining very smoothly with more clusters. However, we observe peaks of the clustering performance for around 9 clusters using the *silhouette* method (difference between the average intra-cluster distance and the average distance to the nearest other cluster) and the *BIC* of a *Gaussian mixture* model. Moreover, the BIC of a Gaussian mixture model also drops when using six clusters, such that six clusters also provide a reasonable separation of the data.

FIGURE D. 1.— Median risk-taking strategy per cluster

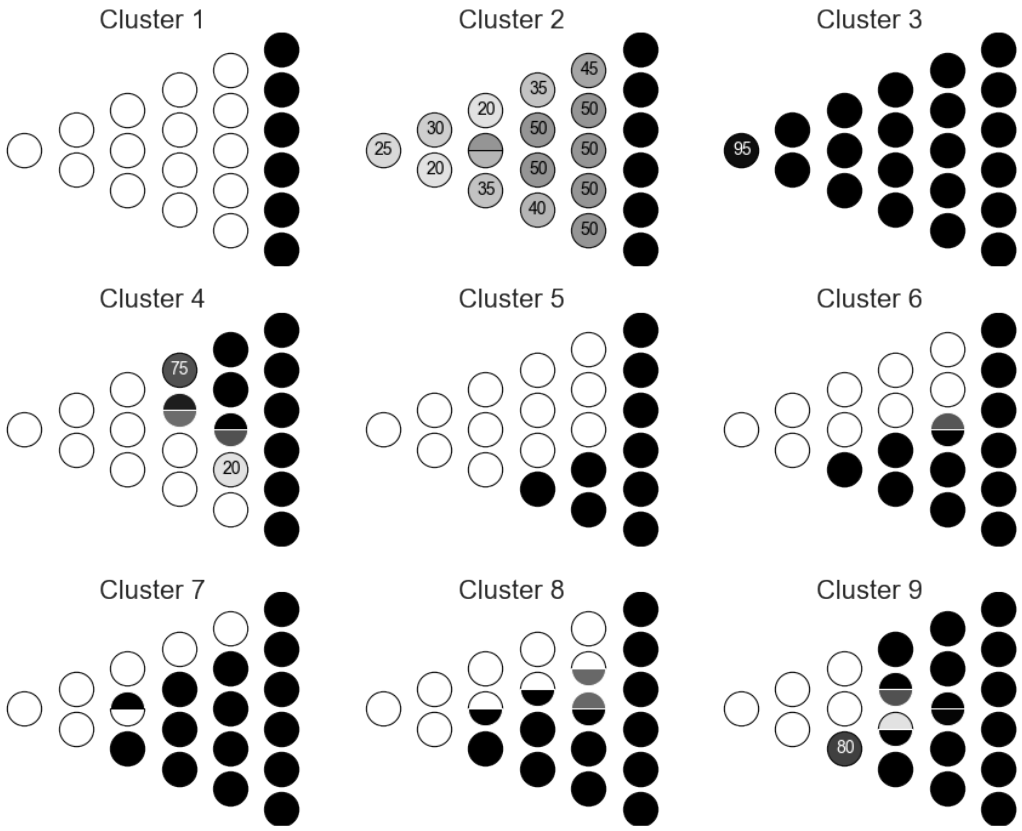

Notes: Figure D. 1 shows the median decoloring per cluster resulting from a K-means clustering with $K = 9$ clusters. To achieve stability in the cluster assignment, we use a majority voting over 200 iterations of the K-means clustering with 100 iterations for the random and k-means++ initialization, respectively.

TABLE D. 1.— Average moments of the stopping time τ and induced outcome distribution S_τ as well as the Barberis (2012) classification for each cluster

						Barberis (2012) classification		
Cluster	N	$\mathbb{E}(\tau)$	$\mathbb{E}(\tau\|\text{Gain})$	$\mathbb{E}(\tau\|\text{Loss})$	$\text{Skew}(S_\tau)$	Loss-exit	Gain-exit	Symmetric-exit
1	117	4.24	4.53	4.54	-0.02	32	36	49
2	52	2.09	2.86	2.79	0.00	24	16	12
3	32	0.67	0.80	0.75	0.07	7	5	20
4	34	3.49	3.23	4.28	-0.38	7	26	1
5	49	4.19	4.72	3.75	0.37	43	3	3
6	35	3.66	4.83	2.86	0.73	35	0	0
7	17	2.75	3.72	2.36	0.77	15	1	1
8	28	2.89	4.43	2.48	0.87	25	0	3
9	36	3.05	3.17	3.18	0.02	16	13	7
Total	400	-	-	-	-	204	100	96

Notes: Table D. 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. The last three columns show the number of subjects with a loss-exit, gain-exit, and symmetric-exit strategy in each cluster, based on each subject's expected stopping time conditional on stopping with a gain or loss (Barberis, 2012).

TABLE D. 2.— Assignment of clusters to economically meaningful strategy types

Strategy type	Clusters	N	Main text
Buy-and-hold	1, 2	169	121
Never-start	3	32	37
Take-profit	4	34	79
Stop-loss	5, 6, 7, 8, 9	165	163

Notes: Table D. 2 shows the assignment of the clusters to economically interpretable strategy types, based on the results from Figure D. 1 and Table D. 1. The last column shows the number of subjects per class as in the main text.

As a further robustness test, we replicate the analyses with $K = 6$ as well as a data-driven number of clusters—$K = 8$ in the following analysis—using a lower-dimensional representation of the risk-taking strategies.[40] In the main text and in the previous analysis, we represent the risk-taking strategies as 31-dimensional tuples with one precommitted stopping action for each (sub-)path. However, 400 data points, one for each subject, are sparse in a 31-dimensional space such that the clustering of the data could be affected by the high-dimensional representation. A projection of the risk-taking strategies onto a lower-dimensional space leads to more "dense" data that may be clustered more effectively.

We use a principal component analysis (PCA) to find a low-dimensional representation of the risk-taking strategies. The PCA finds the directions that best capture the variability in the data by performing an eigendecomposition of the covariance-matrix of all risk-taking strategies to find an orthonormal basis (see Bishop 2006). Overall, we find that the first four principal components retain about 70% of the total variance in the data.[41]

Figure D. 2 shows the projection of the 31-dimensional risk-taking strategies onto the first four principal components (grey dots). The elicited risk-taking strategies are uniformly distributed over the space spanned by the first four principal components. To judge if the PCA dimensions are meaningful, we project 300 simulated risk-taking strategies (see below for details) onto the same four principal components and find a meaningful separation of the simulated strategies: Simulated buy-and-hold strategies vary mostly along the first and third principal components, and the second principal component separates simulated stop-loss and take-profit strategies. The first and third principal component are mostly related to the repetition at which risk-taking stops, while the second principal component captures the gain or loss limit at which risk-taking stops. The fourth principal component is less interpretable.

[40] Using $K = 8$ clusters is suggested by the same data-driven methods as before, while we use $K = 6$ clusters to compare the results more closely to the main text.

[41] We use a spectral analysis to find the number of principal components that captures sufficient variance while being low-dimensional. We observe that the eigenvalues of the principal components decline somewhat sharply beyond the fourth principal component. Using a projection on seven principal components retains 80% of the variance.

FIGURE D. 2.—Projection of elicited risk-taking strategies on the first four principal components

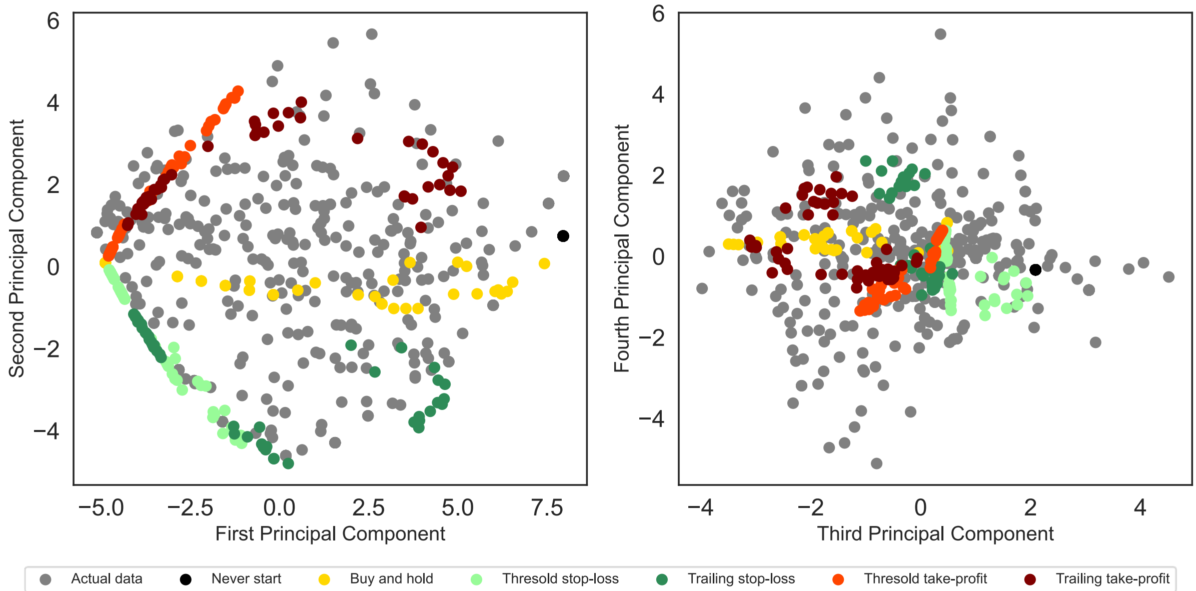

Notes: Figure D. 2 shows the projection of the elicited and simulated risk-taking strategies on the space spanned by the first four principal components. A PCA on the elicited risk-taking strategies shows that about 70% of the total variance can be retained by the first four principal components. In addition, we simulate 50 risk-taking strategies for each of the following strategy labels: never-start, buy-and-hold, threshold stop-loss, trailing stop-loss, threshold take-profit, and trailing take-profit. Thus, we simulate a total of 300 risk-taking strategies and project these strategies onto the same PCA space.

Figure D. 2 highlights the clustering challenge: The elicited risk-taking strategies are spread over the PCA space and more noisy than simulated risk-taking strategies, making a separation difficult. Figure D. 3 and Figure D. 4 show the median decoloring per cluster for $K = 6$ clusters (as in the main text) and for $K = 8$ clusters (determined using the same methods as in Footnote 39), respectively. We interpret the median risk-taking strategy and the cluster constituents as before; and show the cluster interpetration in Table D. 3. Clusters identified by a K-means clustering on PCA-transformed risk-taking strategies also yields economically interpretable strategy clusters. Moreover, Table D. 3 shows that the popularity of the strategy types is robust to the number of clusters and the PCA-transformation. For both, $K = 6$ and $K = 8$ clusters, we find that stop-loss strategies are more popular than buy-and-hold strategies, which are more popular than take-profit strategies. Across the different procedures, the lowest number of strategies is assigned to never-start clusters.

Overall, our behavioral findings are robust to using a different number of clusters, and to projecting the risk-taking strategies onto the space spanned by the first four principal components before performing the K-means clustering.

FIGURE D. 3.— Median risk-taking strategy per cluster for $K = 6$ clusters after a PCA transformation

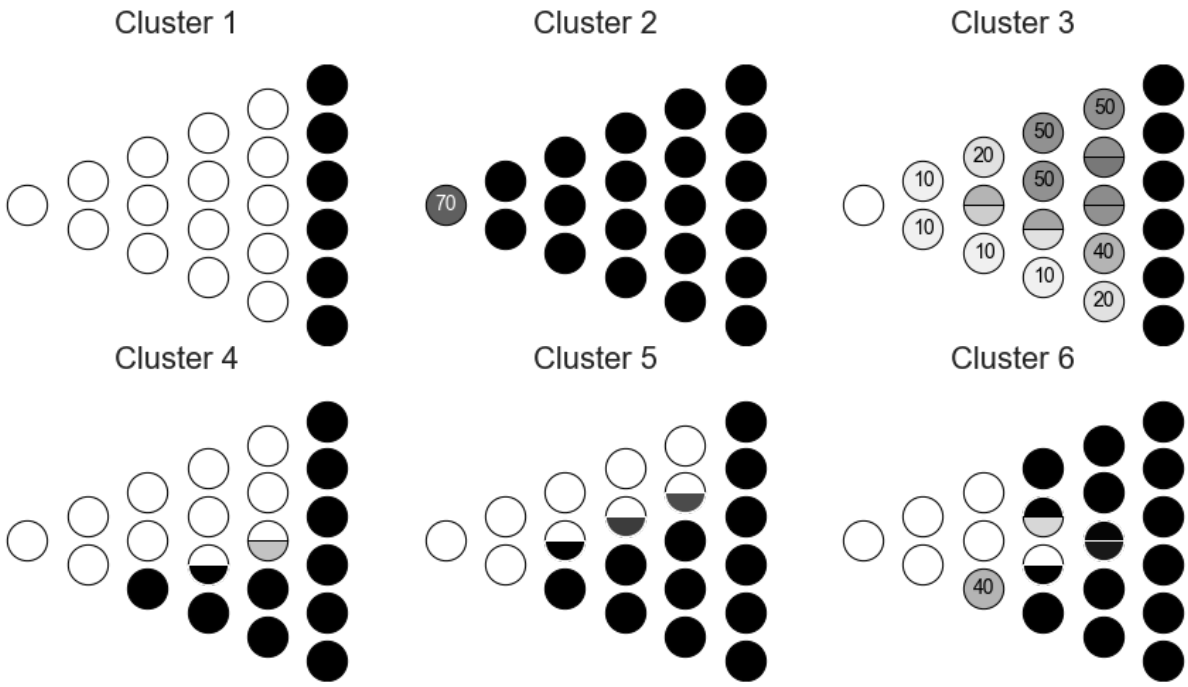

Notes: Figure D. 3 shows the median decoloring per cluster resulting from a K-means clustering with $K = 6$ clusters. To achieve stability in the cluster assignment, we use a majority voting over 200 iterations of the K-means clustering with 100 iterations for the random and k-means++ initialization, respectively. The risk-taking strategies are projected on the first four principal components before performing the K-means clustering.

FIGURE D. 4.— Median risk-taking strategy per cluster for $K = 8$ clusters after a PCA transformation

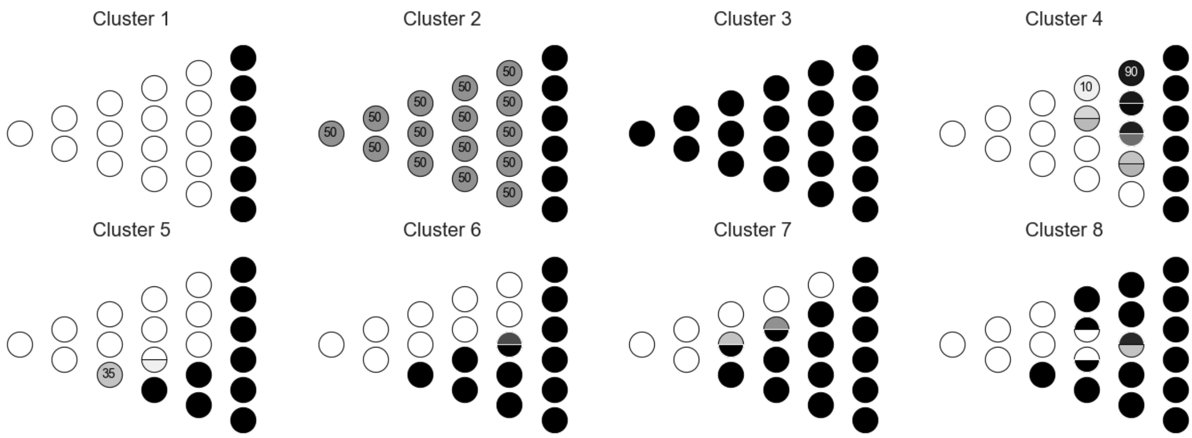

Notes: Figure D. 4 shows the median decoloring per cluster resulting from a K-means clustering with $K = 8$ clusters. To achieve stability in the cluster assignment, we use a majority voting over 200 iterations of the K-means clustering with 100 iterations for the random and k-means++ initialization, respectively. The risk-taking strategies are projected on the first four principal components before performing the K-means clustering.

TABLE D. 3.— Assignment of clusters to economically meaningful strategy types

	$K = 6$		$K = 8$		
Strategy type	Clusters	N	Clusters	N	Main text
Buy-and-hold	1	121	1, 2	154	121
Never-start	2	35	3	30	37
Take-profit	3	73	4	48	79
Stop-loss	4, 5, 6	171	5, 6, 7, 8	168	163

Notes: Table D. 3 shows the assignment of the clusters to economically interpretable strategy types, based on the results from Figure D. 3 and Figure D. 4. The last column shows the number of subjects per class as in the main text.

D.2. Supervised machine-learning—Differentiating trailing and threshold strategies

In the second part of this appendix, we complement the analysis in Section III.G of the main text. After in the previous Appendix D.1 we highlighted that the interpretation of individual risk-taking strategies in terms of buy-and-hold, never-start, take-profit, and stop-loss strategies is robust towards using more clusters and towards projecting the data onto a lower-dimensional space, we now analyze what kind of stop-loss strategies subjects use. As before, we aim at differentiating trailing and threshold stop-loss strategies. 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 in this second part of the appendix builds on supervised machine-learning 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 in two ways: First, the main text used a simple count-based deviation measure (see Section III.G), while we use another metric (cosine-similarity) and three supervised machine-learning algorithms in this appendix. Second, we defined a comprehensive set of all time-and outcome-contingent risk-taking strategies in the main text. Because supervised algorithms learn from labeled data, we generate a different set of template strategies. Regardless of the method and of the construction of the labeled strategies, we find that trailing stop-loss strategies are more popular than threshold stop-loss strategies and thus corroborate the findings from Section III.G.

Time-and outcome- contingent template strategies. We construct a comprehensive set of all time-and outcome-contingent strategies for the following strategy types, which is also used in Section III.G of the main text:

- *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.
- *Thresold 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. The upper level-number of balls (that is, one, two, three balls) in each column are white, and the remaining balls in each column are black.
- *Threshold* and *trailing take-profit* strategies are defined as the mirror image of the trailing or threshold stop-loss strategies.

Figure D. 5 shows all labeled template strategies.

FIGURE D. 5.—All Basic risk-taking strategies for each class label

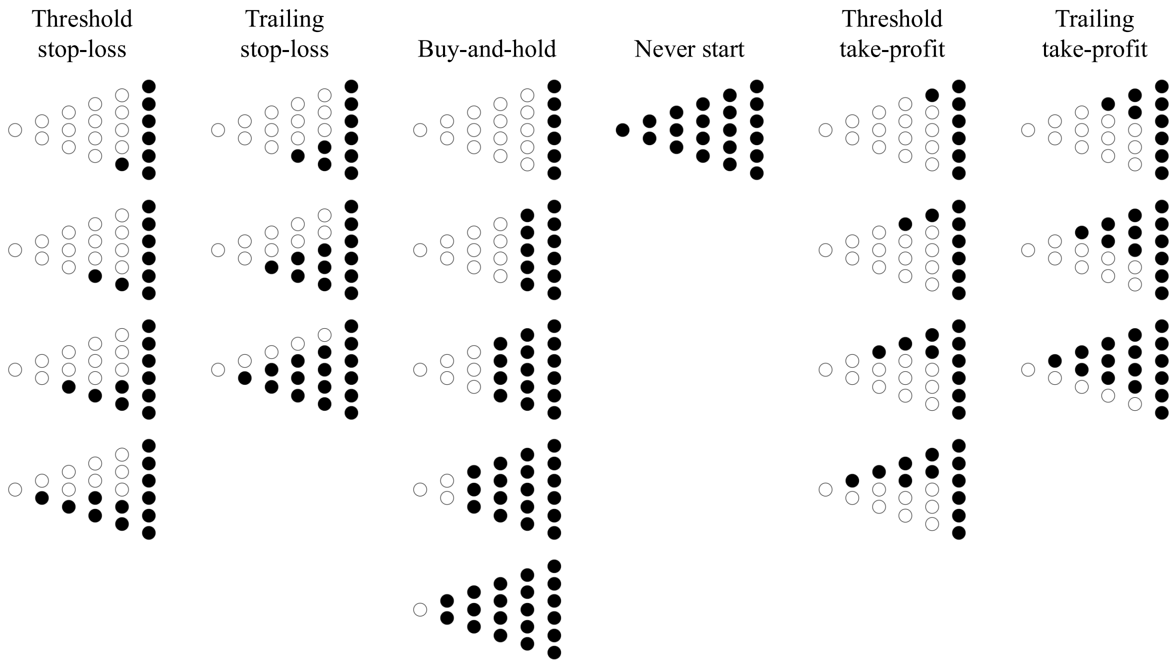

Notes: Figure D. 5 shows the comprehensive set of 20 *Basic* template strategies, with 4 threshold stop-loss, 3 trailing-stop-loss, 5 buy-and-hold, 1 never-start, 4 threshold take-profit, and 3 trailing take-profit strategies.

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 now complement the measure by using the more granular cosine-similarity.[42] 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 D. 6 shows the share of elicited strategies whose similarity to the most similar simulated risk-taking 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 around 50% of the elicited risk-taking strategies can be classified if the threshold is 0.9. Panel (b) of Figure D. 6 shows the classification 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

---

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

more common than threshold stop-loss strategies (27% to 7% for a threshold of 0.6, and 12% to 2% for a threshold of 0.9). Additionally, Table D. 4 shows the classification separately for each study, whereby we do not impose a similarity threshold. On the study level, 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 1.45:1 (*Randomization*^{C}) to 23.00:1 (*Path-dependence*^{C}). Overall, using a cosine-similarity measure instead of the change count used in the main text corroborates our finding that trailing stop-loss strategies are more popular than threshold stop-loss strategies.

FIGURE D. 6.—Share of strategies exceeding the similarity threshold as a function of the threshold and exemplary classifications

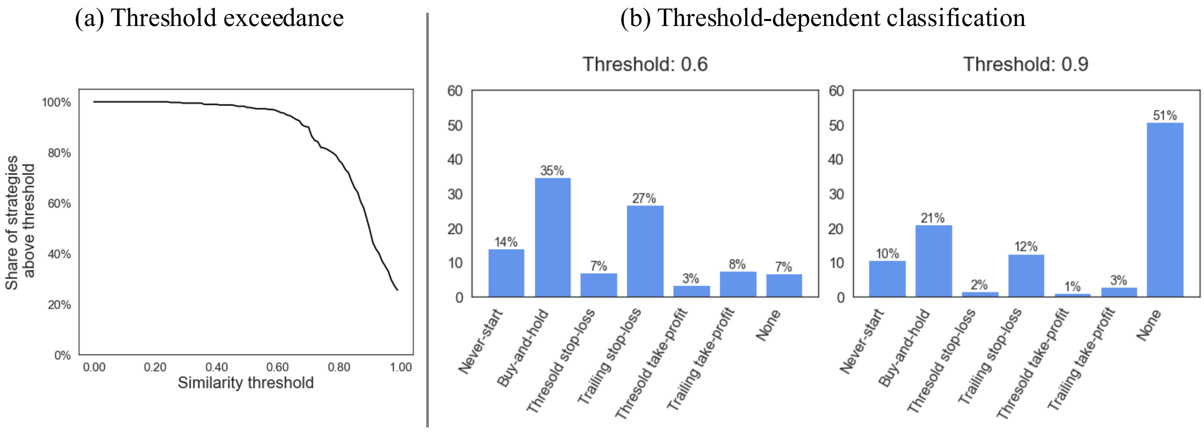

Notes: Panel (a) of Figure D. 6 shows the share of elicited risk-taking strategies whose cosine-similarity to the most similar template strategy exceeds the threshold parameter on the x-Axis. We classify a risk-taking strategy as belonging to a class $c \in$ {never-start, buy-and-hold, stop-loss, trailing stop-loss, take-profit, trailing take-profit} if the similarity to the most similar strategy exceeds the threshold. Panel (b) shows the classification that results from using thresholds 0.6 and 0.9, respectively.

TABLE D. 4.—Cosine-similarity based classification of elicited risk-taking strategies

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	10	30	23	2	-	3	5	3.00
Path-dependence	3	25	33	6	-	3	5	8.33
Path-dependence^C	1	23	28	1	-	10	4	23.00
Randomi-zation	5	14	31	24	3	9	2	2.80
Randomi-zation^C	11	16	26	26	10	7	1	1.45
Total	30	108	141	59	13	32	17	3.60

Notes: Table D. 4 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 two template strategies from different classes are most-similar to the elicited risk-taking strategy.

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-tuple. 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 the precommitted stopping decision at a certain ball 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.[43]
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.

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

The classification of a strategy is eventually obtained as a majority vote from many decision trees.

Table D. 5 below shows the classification obtained from each of the three supervised machine-learning algorithms. For each algorithm, we find that most elicited risk-taking strategies are classified as buy-and-hold strategies, followed by trailing stop-loss strategies. Moreover, 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). The ratio of trailing stop-loss to threshold stop-loss ranges from 3.16 (logistic regression) to 3.96 (decision tree), which is comparable to the results reported in the main text.

TABLE D. 5.—Class assignment for different supervised machine learning algorithms

Class	Logistic regression	Decision tree	Random forest
Never-start	45	46	23
Buy and hold	130	181	185
Threshold stop-loss	37	27	32
Trailing stop-loss	117	107	110
Threshold take-profit	23	8	9
Trailing take-profit	48	31	41
Trailing stop-loss / Threshold stop-loss	3.16	3.96	3.44

Notes: Table D. 5 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 stop-loss to threshold stop-loss strategies.

Simulated template strategies. The classification methods used in this appendix learn an assignment of features of the labeled training data to the strategy classes, and then use the learned assignment to classify the elicited risk-taking strategies. Thus, supervised learning procedures are sensitive to the labeled training data—the template strategies. Therefore, we now construct a different set of labeled template strategies and show the robustness of our results when calibrating the algorithms on the new training data.

We simulate 50 risk-taking strategies for each of the strategy classes discussed above, including never-start, buy-and-hold, threshold stop-loss, trailing stop-loss, threshold take-profit and trailing take-profit strategies. Instead of defining a comprehensive set of risk-taking strategies, we randomly draw a stop-loss level $SL \in \{-0.25, -0.50, \dots, -0.75, -1.00\}$, a stop repetition $SR \in \{1, 2, 3, 4, 5\}$, and a stopping probability $SP \in \{30, 40, \ldots, 70, 80\}$. Figure D. 7 shows three exemplary simulated risk-taking strategies for each strategy class.

FIGURE D. 7.—Three exemplary simulated strategies for each class label

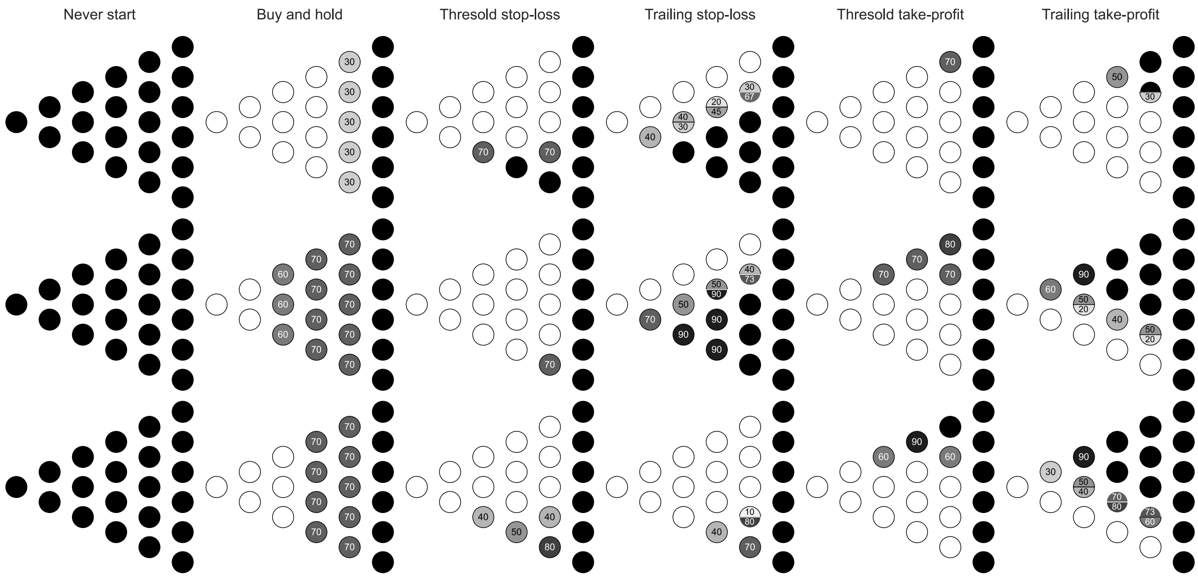

Notes: Figure D. 7 shows three exemplary simulated risk-taking strategies for each of the class labels never-start, buy-and-hold, stop-loss, trailing stop-loss, take-profit and trailing take-profit.

We again use cosine-similarity (Figure D. 8) and the supervised machine-learning algorithms logistic regression, decision tree, and random forest (Table D. 6) to classify the elicited risk-taking strategies. The previous results are robust towards using a different set of labeled training data, as we find that most strategies are classified as buy-and-hold strategies, and more strategies are classified as trailing stop-loss than threshold stop-loss with a ratio ranging from 2.76 (decision tree) to 4.70 (random forest).

Overall, the findings discussed in Section III.G of the main text are robust towards using (a) different algorithms including three supervised machine-learning methods, and (b) to using a different training set of the algorithms that includes even randomized risk-taking strategies.

FIGURE D. 8.—Share of strategies exceeding the similarity threshold as a function of the threshold and exemplary classifications

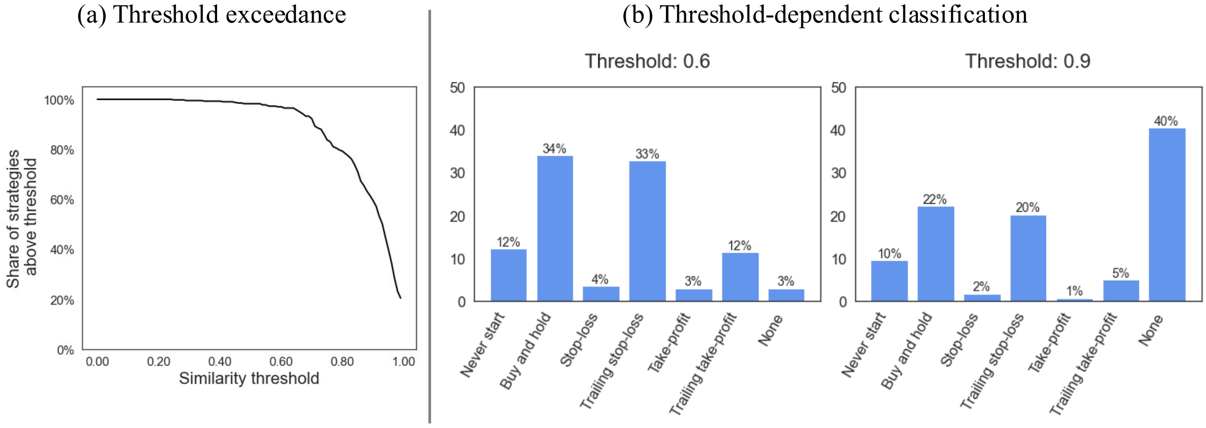

Notes: Panel (a) of Figure D. 8 shows the share of elicited risk-taking strategies whose cosine-similarity to the most similar template strategy exceeds the threshold parameter on the x-Axis. We classify a risk-taking strategy as belonging to a class $c \in$ {never-start, buy-and-hold, stop-loss, trailing stop-loss, take-profit, trailing take-profit} if the similarity to the most similar strategy exceeds the threshold. Panel (b) shows the classification that results from using thresholds 0.6 and 0.9, respectively.

TABLE D. 6.—Class assignment for different supervised machine learning algorithms

Class	Logistic regression	Decision tree	Random forest
Never-start	45	46	28
Buy and hold	130	176	186
Stop-loss	38	37	23
Trailing stop-loss	117	102	108
Take-profit	23	8	11
Trailing take-profit	47	31	44
Trailing stop-loss / Threshold stop-loss	3.08	2.76	4.70

Notes: Table D. 6 shows the assignment of strategies to the six strategy classes using three different supervised machine learning procedures. The bottom rows show the proportion of trailing stop-loss to threshold stop-loss strategies.