

Optimal Stopping in a Dynamic Saliency Model

Markus Dertwinkel-Kalt and Jonas Frey

Discussion Paper #2023/01

DFG Research Unit: “Consumer Preferences, Consumer Mistakes, and Firms’ Response” (FOR 5392)

Optimal Stopping in a Dynamic Saliience Model*

Markus Dertwinkel-Kalt

Jonas Frey

September 2023

Abstract

We study dynamic choice under risk through the lens of salience theory. We derive predictions on salient thinkers' gambling decisions and strategy choices. We test our model experimentally and find support for all of our predictions. We also detect a strong correlation between static and dynamic choices, suggesting that salience theory can coherently explain risky choice in both static *and* dynamic contexts. Our results help to understand when people sell assets, stop gambling, enter the job market or retire.

JEL-Classification: D01; D81; D90.

Keywords: Saliience Theory; Skewness Seeking; Behavioral Stopping.

*The experiments reported in this paper were pre-registered in the AEA RCT registry as trials AEARCTR-0005359 and AEARCTR-0007398, respectively. First of all, we thank Mats Köster for his continuous support throughout the duration of the project. We also thank the editor, Masaki Aoyagi, two anonymous referees as well as Sebastian Ebert, Nicola Gennaioli, Katrin Gödker, Paul Heidhues, Alex Imas, Joshua Miller, Frank Schlütter, Philipp Strack, and Peter Wakker for helpful comments and suggestions. We gratefully acknowledge financial support by the German Research Foundation (DFG project No. 462020252). Dertwinkel-Kalt: University of Münster & Max Planck Institute for Collective Goods, markus.dertwinkel-kalt@wiwi.uni-muenster.de. Frey: Said Business School, University of Oxford, & University of Münster, jonas.frey.dphil2017@said.oxford.edu.

1 Introduction

Dynamic decisions under risk are ubiquitous. Examples include the decisions when to sell an asset, when to stop gambling, when to enter the job market or to retire, when to buy a flight ticket or a durable good, and when to stop searching for a house or a spouse. An extensive behavioral literature of choice under risk — going back to Kahneman and Tversky (1979) — has singled out the skewness of the underlying probability distribution as one important driver of risk attitudes. In particular, people typically seek (positively) skewed and avoid negatively skewed risks. The corresponding literature on skewness seeking has traditionally focused on static decisions in order to explain, for instance, why people buy expensive insurance and lottery tickets at the same time. In dynamic problems, skewness seeking even matters when the underlying risk or stochastic process is symmetric (e.g., Barberis, 2012; Ebert and Strack, 2015; Ebert, 2020), because the decision maker can select a skewed return distribution through the choice of her stopping strategy. Hence, skewness seeking may play an even bigger role in dynamic than in static problems.

Since skewness seeking in static and dynamic decisions plausibly stems from the same cognitive mechanism, the existing literature has adapted static models of choice under risk — such as expected utility theory (EUT) or cumulative prospect theory (CPT; Tversky and Kahneman, 1992) — to dynamic decision making. We follow this approach by studying, both theoretically and experimentally, the dynamic implications of Bordalo *et al.*'s (2012) salience theory. In salience theory, skewness seeking in static problems originates from the idea that outcomes, which are extreme relative to a reference point — such as the relatively large upside of a right-skewed risk — are particularly salient. These outcomes attract a disproportionate amount of attention and their probabilities are overweighted (Bordalo *et al.*, 2012). Thus, salience theory exhibits the core prediction that agents seek skewness (Dertwinkel-Kalt and Köster, 2020). At the same time, the salience distortions are bounded, which places a tight limit on the degree of probability weighting and therefore also on skewness seeking. Models that imply much stronger skewness seeking than salience theory (such as CPT) could make reasonable predictions in static settings but make unreasonably extreme predictions in some important dynamic settings (Ebert, 2015). Hence, salience theory represents a natural candidate for modelling skewness seeking in dynamic choices.

We therefore apply salience theory to dynamic contexts, and derive, test and support salience theory's predictions for dynamic choice under risk. We also document a strong relation between skewness seeking in static and in dynamic setups. This suggests that there is a common mechanism that drives static and dynamic choice under risk, highlighting the importance of developing a model that can explain both static and dynamic decisions.

In Section 3, we adapt salience theory to analyze dynamic choice under risk at the hand of standard optimal stopping problems. We ask when a naïve¹ salient thinker stops an Arithmetic

¹Non-linear probability weighting implies that an agent's optimal strategy at time t might no longer be optimal at some later point in time (e.g., Machina, 1989). Optimal stopping behavior under salience theory and other behavioral models, thus, depends on whether the agent is aware of this time-inconsistency (i.e., the agent is sophisticated) or not (i.e., the agent is naïve). We assume that the agent is naïve (as Ebert and Strack, 2015), which is also supported by our experimental results.

Brownian Motion (ABM) with a non-positive drift and a finite expiration date. Expected utility theory with a concave utility function cannot explain gambling when the agent loses money in expectation, that is, when the ABM’s drift is negative. With a specific stopping strategy in mind, a *salient thinker*, however, inflates the probabilities of those realizations that “differ most” from his current wealth level (*contrast effect*), which might render gambling attractive. Adopting the naïve decision rule proposed in the literature (e.g., Ebert and Strack, 2015), we assume that the (naïve) salient thinker continues to gamble as long as he can find at least one stopping strategy that is more attractive to him than stopping.

In Section 4, we derive our theoretical results. Unlike an expected-utility agent with a concave utility function, a naïve salient thinker gambles even when he loses money in expectation. At the same time, unlike in CPT (Ebert and Strack, 2015), a naïve salient thinker stops any ABM with a sufficiently negative drift. In a next step, we restrict the choice set to all *stop-loss and take-profit strategies*² to learn more about how a naïve salient thinker plans to stop, and how he will revise this plan over time. These additional restrictions allow for more interesting experimental predictions of salience theory. First, a salient thinker chooses a particular subset of stop-loss and take-profit strategies, which give rise to a right-skewed distribution of returns; so-called *loss-exit strategies* (see, e.g., Barberis, 2012; Heimer *et al.*, 2023). A loss-exit strategy is defined as a stop-loss and take-profit strategy for which the stop-loss threshold is closer to the current value of the process than the take-profit threshold, so that — by the contrast effect — stopping at a gain is more salient than stopping at a loss. Second, a naïve salient thinker does not necessarily follow his initial plan, but might instead revise his strategy over time. In particular, salience theory can rationalize stopping behavior that is consistent with the well-known disposition effect (e.g., Shefrin and Statman, 1985; Odean, 1998; Barberis, 2012).

Section 5 presents a laboratory experiment that is designed to test our salience-based predictions on stopping behavior. Participants in the experiment have to decide when to stop ABMs with different non-positive drifts. Subjects stop the process by defining an upper and a lower bound and the process is stopped if it reaches either bound. If a process is stopped, subjects can either sell it or restart it by moving the bounds. This design allows us to test whether subjects choose loss-exit strategies (i.e., strategies with a salient upside) and whether they revise their initial strategies as predicted by the model. We validate our approach of adapting the static salience model to an optimal stopping problem by further eliciting skewness seeking in static choices. Generalizing results from Dertwinkel-Kalt and Köster (2020), we show that, for a fixed expected value and variance, a salient thinker chooses a binary lottery over the safe option paying its expected value with certainty if and only if the lottery’s skewness exceeds some threshold. If salience is indeed the psychological mechanism driving skewness seeking in general, it should coherently explain revealed attitudes toward skewness in such static choices as well as in the optimal stopping problems.

²A stop-loss and take-profit strategy is characterized by a stop-loss threshold below the current value of the process and a take-profit threshold above the current value of the process at which the process will be stopped. These strategies are often proposed by retail banks to their customers (see, e.g., the brokerage data by Heimer *et al.*, 2023) and have attracted much attention in the related literature (e.g., Xu and Zhou, 2013; Ebert and Strack, 2015; Fischbacher *et al.*, 2017; Heimer *et al.*, 2023). Important for us, these strategies allow agents to obtain skewed return distributions even if the underlying stochastic process is symmetric.

In Section 6, we present our experimental results. First of all, we find that subjects select skewed return distributions: for the median subject, more than 70% of all chosen strategies are loss-exit strategies. Furthermore, 93% of the subjects revise their initial strategy at least once, and actual behavior is reminiscent of the disposition effect.

We next examine how sensitive subjects' stopping decisions are to the drift of the process. While it might seem obvious that fewer subjects should gamble with more negative drifts, both EUT and CPT predict that this is not the case. Risk averse EUT agents should not gamble for any non-positive drift. In contrast, subjects following the commonly used CPT-models as analyzed by Ebert and Strack (2015) should always gamble regardless of how negative the drift is. We find that most subjects start gambling a fair process with a drift of zero, but then stop before reaching the expiration date. Moreover, subjects stop the earlier, the more negative the drift of the process. Interpreted through the lens of our model, these results indicate heterogeneity in the strength of salience distortions of our subjects: Around 95% of the subjects reveal sufficiently strong salience distortions that they start the fair process, but for only 60% of the subjects salience distortions are so strong that also gambling with the most negative drift is attractive.

We also find a positive correlation between static and dynamic skewness seeking, which is both, statistically and economically, significant: subjects that reveal stronger skewness seeking in static choices also have a larger propensity to choose loss-exit strategies in the dynamic ones. Overall, our experimental results suggest that a good model of dynamic risk-taking without commitment should include moderately strong skewness seeking. Moreover, it should be based on a mechanism that can also explain skewness seeking in static choices. Our salience theory model fulfils these criteria and is therefore able to coherently explain choices in static and dynamic problems.

In Section 7, we show that several popular models of static choice under risk struggle to explain the dynamic evidence from our experiment because they either predict too strong skewness seeking or no skewness seeking at all. Arguably, alternative models with the right strength of skewness seeking that allows to explain our data can be developed, but we are not aware of any such model among the commonly used ones. Moreover, as salience theory can jointly explain static and dynamic data, we regard it as a prime candidate for a unified model of static and dynamic choice under risk.

We conclude in Section 8 by discussing further applications of our findings.

2 Related Literature on Behavioral Stopping

Our paper is related to a large theoretical (e.g., Machina, 1989; Karni and Safra, 1990; Barberis, 2012; Ebert and Strack, 2015, 2018; Henderson *et al.*, 2017; He *et al.*, 2019; Strack and Viefers, 2021) as well as a growing experimental (e.g. Imas, 2016; Nielsen, 2019; Strack and Viefers, 2021; Heimer *et al.*, 2023) literature on behavioral stopping. On the one hand, we add to the theoretical literature by providing the first study of behavioral stopping in salience theory. On the other hand, we contribute to the experimental salience literature (for a survey, see Bordalo *et al.*, 2022) by testing salience theory's predictions on behavioral stopping as well as by investigating whether it can coherently explain both static and dynamic choice under risk.

Most existing theoretical work on behavioral stopping deals with the implications of non-linear probability weighting for dynamic gambling, with a focus on the behavior predicted by cumulative prospect theory (henceforth: CPT, see Machina, 1989; Karni and Safra, 1990; Barberis, 2012; Xu and Zhou, 2013; Ebert and Strack, 2015, 2018; Henderson *et al.*, 2017, 2018; He *et al.*, 2019). This focus can be explained by the fact that non-linear probability weights imply (empirically relevant) time-inconsistent preferences (e.g., Machina, 1989). Predicted behavior depends, in particular, on whether or not the agent is *naïve* about his time-inconsistency. A naïve agent will revise his strategy throughout time, while a (fully) sophisticated agent foresees her intention to adjust certain strategies and chooses only strategies she will actually follow through with (e.g., Karni and Safra, 1990). With time-inconsistent preferences also the question of whether the agent can commit to a strategy becomes important. The literature has studied the stopping behavior of naïve agents without commitment (e.g., Barberis, 2012; Ebert and Strack, 2015) as well as with partial or full commitment (e.g., Xu and Zhou, 2013; Henderson *et al.*, 2017; He *et al.*, 2019).

For our purpose of testing salience theory’s predictions, the setups with naïvete, but without commitment are relevant (Barberis, 2012; Ebert and Strack, 2015): naïvete is a more plausible assumption than sophistication (for instance, sophisticates should not start gambling with non-positive drifts), and only the absence of commitment allows to test whether predictions on time-inconsistent behavior hold true. In the seminal paper by Barberis (2012), it is numerically shown that in finite discrete time setups naive CPT agents without commitment mostly choose loss-exit (as compared to gain-exit) strategies and start to gamble (at least for a wide range of parameters), but revise their strategies, so that *ex post* they exhibit gain-exit behavior. The reason is that close to the expiration date, agents cannot choose strongly skewed return distributions anymore, and therefore exit earlier than intended.³ In continuous time setups, this mechanism is not at work: agents can always choose strongly skewed return distributions and, as a consequence, naive CPT-agents never stop (Ebert and Strack, 2015). This also holds true with an indefinite end date: for most empirically relevant cumulative prospect theory parameter values, a naive agent does not stop with probability one at any loss level (He *et al.*, 2019). This never-stopping prediction can only be avoided by allowing for randomized stopping strategies and thereby some form of commitment (Henderson *et al.*, 2017), or by imposing different functional forms for dynamic choices than those typically used for static choices (Duraj, 2020; Huang *et al.*, 2020). We show in how far the predictions of salience theory differ from those by other models such as CPT, and show in particular that dynamic salience theory does not yield the (too) extreme never-stopping prediction that Ebert and Strack (2015) have derived for CPT.

We also contribute to the small, but growing experimental literature on behavioral stopping (e.g., Imas, 2016; Nielsen, 2019; Strack and Viefers, 2021). Unlike us, these papers focused on the question in how far stopping decisions are path-dependent; in particular, in how far the realization of previous gains and losses affects behavior. Closest related to us is the contemporaneous paper by Heimer *et al.* (2023), who study optimal stopping behavior using a process consisting of repeated (fair) coin tosses. Similarly to us, they focus on stop-loss and take-profit strategies,

³Heimer *et al.* (2023) provide direct experimental evidence for this prediction.

and they find, both in laboratory experiments as well as in observational brokerage data, that subjects *ex ante* choose loss-exit strategies, but then deviate by revealing disposition-effect-like behavior. Both findings are also reflected in our data. In contrast to Heimer *et al.* (2023), our paper is focused on salience theory and establishes a novel link between static and dynamic skewness seeking.

3 A Dynamic Version of Salience Theory of Choice under Risk

3.1 Static Model

Consider an agent who has to choose from some set \mathcal{C} that contains exactly two non-negative random variables (or *lotteries*), X and Y , with a joint cumulative distribution function (CDF) $F : \mathbb{R}_{\geq 0}^2 \rightarrow [0, 1]$. A state of the world here refers to a tuple of outcomes, $(x, y) \in \mathbb{R}_{\geq 0}^2$. We denote the state space by $\mathcal{S} \subseteq \mathbb{R}_{\geq 0}^2$. If a random variable is degenerate, we call it a *safe* option.

According to salience theory of choice under risk (Bordalo *et al.*, 2012), the agent is a *salient thinker*, who evaluates a random variable by assigning a subjective probability to each state of the world $s \in \mathcal{S}$ that depends on the state's objective probability and on its salience. The salience of a state is assessed by a so-called *salience function*, which is defined as follows:

Definition 1 (Salience Function). *We say that a symmetric, bounded, and twice differentiable function $\sigma : \mathbb{R}_{\geq 0}^2 \rightarrow \mathbb{R}_{> 0}$ is a salience function if and only if it satisfies the following two properties:⁴*

1. *Ordering. Let $x \geq y$. Then, for any $\epsilon, \epsilon' \geq 0$ with $\epsilon + \epsilon' > 0$, we have*

$$\sigma(x + \epsilon, y - \epsilon') > \sigma(x, y).$$

2. *Diminishing sensitivity. For any $x > y$ and any $\epsilon > 0$, we have*

$$\sigma(x + \epsilon, y + \epsilon) < \sigma(x, y).$$

We say that a given state of the world $(x, y) \in \mathcal{S}$ is the more salient the larger its salience value is. The ordering property implies that a state of the world is the more salient the more the attainable outcomes in this state differ. In this sense, ordering captures the well-known *contrast effect* (e.g., Schkade and Kahneman, 1998), whereby large contrasts (in outcomes) attract a great deal of attention. Diminishing sensitivity reflects *Weber's law* of perception, and it implies that the salience of a state decreases if the outcomes in this state uniformly increase. Hence, diminishing sensitivity describes a *level effect*, according to which a given contrast in outcomes is less salient at a higher outcome level, thereby qualifying the contrast effect.

A salient thinker is intrinsically (weakly) risk-averse but may, depending on the salience of outcomes, sometimes behave as if being risk-seeking. He evaluates monetary outcomes via a strictly increasing, (weakly) concave, and twice differentiable value function $v : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$,

⁴Bordalo *et al.* (2012) also allow for random variables with negative outcomes and add a third property to ensure that diminishing sensitivity (with respect to zero) reflects to the negative domain: by the *reflection* property, for any $w, x, y, z \geq 0$, it holds that $\sigma(x, y) > \sigma(w, z)$ if and only if $\sigma(-x, -y) > \sigma(-w, -z)$.

and forms an “expectation” by assigning each state of the world a (probability) weight that depends on how a given option compares to the alternative at hand in this state. More specifically, a salient thinker behaves as if maximizing a *salience-weighted utility*, which is defined as follows:

Definition 2. *The salience-weighted utility of a random variable X evaluated in $\mathcal{C} = \{X, Y\}$ equals*

$$U^s(X|\mathcal{C}) = \int_{\mathbb{R}_{\geq 0}^2} v(x) \cdot \frac{\sigma(v(x), v(y))}{\int_{\mathbb{R}_{\geq 0}^2} \sigma(v(s), v(t)) dF(s, t)} dF(x, y),$$

where $\sigma : \mathbb{R}_{\geq 0}^2 \rightarrow \mathbb{R}_{> 0}$ is a salience function that is bounded away from zero.

Since the salience-weighted probabilities are normalized so that they sum up to one (e.g., Bordalo *et al.*, 2012; Dertwinkel-Kalt and Köster, 2020), a salient thinker’s valuation of a safe option $x \in \mathbb{R}_{\geq 0}$ is undistorted and given by $v(x)$, irrespective of the properties of the alternative option.

3.2 Dynamic Model

Stochastic process. Following Ebert and Strack (2015, 2018), we model an agent’s wealth via a Markov diffusion. Specifically, we consider an *Arithmetic Brownian Motion* (ABM),

$$dX_t = \mu dt + \nu dW_t,$$

with an initial value $X_0 = x$, a constant drift $\mu \in \mathbb{R}$ and a constant volatility $\nu \in \mathbb{R}_{> 0}$, as well as a standard Brownian Motion $(W_t)_{t \in \mathbb{R}_{\geq 0}}$.

To make the theory testable in the context of an incentivized lab experiment, we deviate from Ebert and Strack (2015, 2018) in two ways: First, we assume that the process is non-negative, and absorbing in zero. Second, we allow for a finite *expiration date* $T \in \mathbb{R}_{> 0} \cup \{\infty\}$.

Stopping strategies. As in Ebert and Strack (2015), we represent the set of stopping strategies by the set of stopping times, where each stopping time τ refers to a deterministic plan of when to stop the process. The central feature of a stopping time is that it is based on past information only: that is, any τ is adapted to the natural filtration $(\mathcal{F}_t)_{t \in \mathbb{R}_{\geq 0}}$ of the process $(X_t)_{t \in \mathbb{R}_{\geq 0}}$. For a fixed expiration date $T \in \mathbb{R}_{> 0} \cup \{\infty\}$, choosing a stopping time $\tau \leq T$ (with probability one) implements a random wealth level X_τ with a cumulative distribution function denoted by F_τ .

For our first results, we do not impose any restrictions on the set of deterministic stopping times that the agent can choose from. But, to learn more about the strategies that are attractive to a salient thinker and the role of skewness for stopping behavior, we derive additional results under the assumption that the agent is restricted to choose a *threshold stopping time* $\tau_{a,b}$ — defined as the first leaving time of the interval (a, b) — that implements a random wealth level $X_{T \wedge \tau_{a,b}}$. The set of threshold stopping times represents the set of stop-loss and take-profit strategies, which are often proposed by retail banks to their customers (see., e.g., the brokerage data by Heimer *et al.*, 2023) and which have attracted much attention in the behavioral and financial literature (e.g., Xu and Zhou, 2013; Ebert and Strack, 2015; Fischbacher *et al.*, 2017; Heimer *et al.*, 2023).

Solution concept under salience theory. Any form of non-linear probability weighting — whether it is salience-driven or mechanical — implies that an agent’s optimal strategy at time t might no longer be optimal at some later point in time (e.g., Machina, 1989). Thus, optimal stopping behavior under salience theory depends on whether or not the salient thinker is aware of this time-inconsistency. We follow Ebert and Strack (2015) in assuming that the agent is naïve about his time-inconsistency. As we think of salience effects as unconscious distortions of perception, we regard the assumption of naïvete as sensible. In Section 5 and Appendix B, we further discuss how to experimentally test this assumption within the salience framework.

As in Ebert and Strack (2015), we assume that “at every point in time the naïve [salient thinker] looks for a [...] strategy τ that brings him higher [salience-weighted utility] than stopping [...]. If such a strategy exists, he holds on to the investment — irrespective of his earlier plan.” Assuming that the naïve salient thinker continues to gamble if and only if he strictly prefers to do so, the decision rule then reads as follows.

Definition 3 (Continuation Rule). *Let $x_t \in \mathbb{R}_{\geq 0}$ be the current wealth level at time $t \in [0, T]$. A naïve salient thinker continues at time t if there exists a stopping time τ , such that $U(X_\tau | \{X_\tau, x_t\}) > u(x_t)$, that is, if the salient thinker finds a strategy that gives him a strictly higher salience-weighted utility than stopping at time t . Otherwise, the naïve salient thinker stops at time t .*

Our decision rule imposes the additional assumption that a naïve salient thinker evaluates each stopping strategy in isolation: at any point in time, the *consideration set* — that is, the set of strategies that the agent compares when making his stopping decision — includes a single strategy to continue with, X_τ , and the alternative to stop right now, x_t ; the consideration set thus is assumed to be $\{X_\tau, x_t\}$. Since salience theory is a model of context-dependent behavior to derive testable predictions, it is necessary to impose some assumption on the consideration set. With infinitely many strategies to choose from, we regard the above specification as plausible. Moreover, our experimental design (see Section 4 for details) highlights a single strategy at a time, so that subjects likely evaluate this strategy in isolation. Still, one might argue that previously chosen strategies affect the perception of whatever strategy is considered next. Without any guidance on how the consideration set changes over time, however, it is impossible to provide a comprehensive analysis.⁵ To tie our hands, we pre-registered our assumptions on the consideration set before running the experiment.

⁵When restricting attention to stop-loss and take-profit strategies, one could argue that the previously chosen lower bound a_p of a stop-loss and take-profit strategy provides a “reference point” for the newly selected lower bound a_n , and that the previously chosen upper bound b_p provides a “reference point” for the newly chosen upper bound b_n in the sense that the respective salience weights are $\sigma(v(a_n), v(a_p))$ and $\sigma(v(b_n), v(b_p))$. Then, conditional on not stopping the process, subjects would always adjust the upper threshold by more than the lower threshold, as otherwise the lower threshold would be salient. While this prediction is inconsistent with the data that we present later on, there might be other specifications of the consideration that are consistent with our experimental findings.

4 Stopping Behavior of a Naïve Salient Thinker

4.1 Motivating Example

To illustrate the salience mechanism, consider a salient thinker with a linear value function, $v(x) = x$, who decides when to stop a fair process with zero drift that does not expire ($T = \infty$). Suppose that the agent adopts a stop-loss and take-profit strategy, which can be represented by a threshold stopping time $\tau_{a,b}$ with a being the lower and b being the upper threshold. Such a strategy induces a binary lottery, $X_{\tau_{a,b}} = (a, p; b, 1 - p)$, over wealth. Because the process has a drift of zero, at time t , the expected value of following this stop-loss and take-profit strategy is $\mathbb{E}_t[X_{\tau_{a,b}}] = x_t$. Does the salient thinker ever stop?

It is immediate to see that a salient thinker with a linear value function chooses a binary lottery with upside payoff b , downside payoff a , and expected value x_t over the safe option paying x_t if and only if the lottery's upside b is assigned a larger salience weight than the lottery's downside a , that is, if and only if $\sigma(b, x_t) > \sigma(a, x_t)$. As a consequence, whenever $\sigma(b, x_t) > \sigma(a, x_t)$, following the stop-loss and take-profit strategy represented by $\tau_{a,b}$ is more attractive to the salient thinker than stopping at time t . Since $\sigma(b, x_t) > \sigma(x_t, x_t)$ due to ordering, and since the salience function is continuous, we can always find a stopping time $\tau_{a,b}$ — with a close enough to the current wealth level x_t — that the salient thinker prefers to stopping at time t . Hence, he never stops. It is easily verified that the result remains to hold for a finite expiration date. All missing proofs are provided in Appendix A.

Proposition 1. *Fix an initial wealth level $x \in \mathbb{R}_{>0}$ and expiration date $T \in \mathbb{R}_{>0} \cup \{\infty\}$. A naïve salient thinker with a linear value function does not stop a process with zero drift at any positive level of wealth.*

4.2 Main Theoretical Result

We are interested in how general the never-stopping result derived in the previous subsection is. By Definition 3, a salient thinker continues (or starts) to gamble if he can find a strategy that gives him strictly higher utility than not gambling. We will now show that there are two reasons why a salient thinker cannot find such a strategy and hence stops before the expiration date: either the drift of the process is sufficiently negative, or the salient thinker is intrinsically risk-averse. More precisely, while a naïve salient thinker with a linear value function also holds processes with a slightly negative drift until the expiration date, a salient thinker with a sufficiently concave value function does not start even a fair process, and this holds irrespective of his intrinsic risk-aversion (i.e., irrespective of how concave his value function is). This last prediction distinguishes salience theory from models like CPT (see Ebert and Strack, 2015), and it constitutes our main theoretical result.

Consider an ABM with an arbitrary drift $\mu \in \mathbb{R}$ and volatility $\nu \in \mathbb{R}_{>0}$. By Definition 3, a naïve salient thinker does not start to gamble if and only if, for any stopping time $\tau \leq T$,

$$\int_{\mathbb{R}_{\geq 0}} (v(z) - v(x)) \sigma(v(z), v(x)) dF_{\tau}(z) \leq 0, \quad (1)$$

where F_{τ} denotes the CDF of the induced wealth level X_{τ} . Fixing the initial value $X_0 = x$, we

define an *auxiliary utility function* $\tilde{u}(z) := (v(z) - v(x))\sigma(v(z), v(x))$, which is strictly increasing and differentiable in $z \in \mathbb{R}_{\geq 0}$. By construction, the condition derived in Eq. (1) is equivalent to

$$\int_{\mathbb{R}_{\geq 0}} \tilde{u}(z) dF_{\tau}(z) \leq \tilde{u}(x).$$

In words, for any fixed initial value $X_0 = x$, a naïve salient thinker does not start if and only if an EUT-agent with a utility function $\tilde{u}(\cdot)$ does not start. The main step in proving that a naïve salient thinker does not start any ABM with a sufficiently negative drift, is to derive a bound on how risk-seeking a salient thinker can ever be.

Our first result approximates a salient thinker's willingness to take risk by that of an EUT-agent with an exponential utility function. More precisely, there is an EUT-agent with exponential utility who takes up some risks that a salient thinker certainly avoids, thereby imposing a bound on the salient thinker's willingness to take risk. Given that we can approximate a salient thinker's willingness to take risk by that of an EUT-agent with an exponential utility function, we can apply Proposition 1 in Ebert and Strack (2015) to show that a naïve salient thinker does not start any process with a sufficiently negative drift.

Theorem 1. *For any expiration date $T \in \mathbb{R}_{>0} \cup \{\infty\}$, any initial wealth level $x \in \mathbb{R}_{>0}$ and any volatility $\nu \in \mathbb{R}_{>0}$, there exists some $\tilde{\mu} \in \mathbb{R}$, such that a naïve salient thinker does not start any process with a drift $\mu < \tilde{\mu}$.*

Building on Theorem 1, we observe that an intrinsically risk-averse salient thinker may not even start a fair process; whether he does so depends on his intrinsic risk aversion (i.e., the concavity of his value function). We thus obtain the following corollary to the preceding theorem.

Corollary 1. *Depending on the concavity of their value function, salient thinkers may start or not start a process with zero drift.*

Theorem 1 and Corollary 1 allow us to distinguish between salience theory and its main alternative models (as discussed in detail in Section 7): EUT with a concave value function as well as reference-dependent preferences without probability weighting predict that a process with a non-positive drift is not started, while CPT and models of disappointment aversion predict that such a process is always started (and even never stopped before the expiration date). Salience theory permits for (an arguably more realistic) heterogeneity in gambling behavior.

That salience theory produces more realistic predictions than CPT is due to the boundedness of the salience function; if the salience function was unbounded, similar predictions as in CPT would prevail. While Bordalo *et al.* (2012) did not psychologically motivate the boundedness of the salience function they assumed, it is in line with the well-known fact that humans have difficulties interpreting numbers outside of the range they commonly experience, namely, very small and very large numbers (see Resnick *et al.*, 2017, for a review). Salience distortions stem from large payoff contrasts that attract attention. But when humans cannot well understand the difference in magnitude between two large numbers, then it is natural to assume that increasing a large contrast even further does not induce a (strong) behavioral reaction and, therefore, should also not distort salience weights much further (as salience theory wants to well-describe

actual behavior). This is precisely the effect that the boundedness of the salience function produces.⁶

4.3 Gambling an (Un)Fair Process with Stop-Loss and Take-Profit Strategies

To learn more about the behavior of a naïve salient thinker, we restrict the choice set to all stop-loss and take-profit strategies, and consider only processes with a non-positive drift, $\mu \in \mathbb{R}_{\leq 0}$. First, we characterize the type of stop-loss and take-profit strategies that is attractive to a salient thinker. Second, using the additional structure, we derive a stronger result on the limits of naïve gambling. Third, we show that salience theory can rationalize the disposition effect; that is, the tendency to rather stop processes that have increased in value than those that have decreased in value (e.g., Shefrin and Statman, 1985; Odean, 1998; Weber and Camerer, 1998; Imas, 2016).

The role of skewness in naïve gambling. When referring to skewness, we use the most conventional definition of skewness, whereby skewness $S[X]$ of a lottery X is defined by the third standardized central moment

$$S[X] := \mathbb{E} \left[\left(\frac{X - \mathbb{E}[X]}{\sqrt{\text{Var}[X]}} \right)^3 \right]. \quad (2)$$

We can then define right- and left- skewness as well as loss-exit strategies, a subset of stop-loss and take-profit strategies that give rise to right-skewed return distributions.

Definition 4. Lottery X is called right-skewed (or, equivalently, positively skewed) if $S(X) > 0$, left-skewed (or, equivalently, negatively skewed) if $S(X) < 0$, and symmetric otherwise.

Definition 5 (Loss-Exit Strategy). A loss-exit strategy at wealth level x_t is a stop-loss and take-profit strategy $\tau_{a,b}$ (with b denoting the upper and a the lower threshold) such that $b - x_t > x_t - a$.

A loss-exit strategy derives its name from combining a relatively large upside with a moderate downside (so that the upper threshold is further away from the current wealth level than the lower threshold), which makes it likely to stop at a loss when the process of the drift is non-positive. To show that such a strategy induces a right-skewed return distribution, suppose the underlying process has a non-positive drift and no expiration date. Then, a loss-exit strategy induces a binary lottery that is positively skewed (this directly follows from the fact that the upper threshold is reached with less than 50% probability, see Lemma 1 in Dertwinkel-Kalt and Köster, 2020). As a finite expiration date induces a very complicated CDF (see Lemma 2 in the Appendix), we cannot analytically prove that this holds true also with a finite expiration date; we can, however, back this claim with the help of numerical simulations, where we simulate

⁶Interestingly, a similar argument can also motivate the fact that probability distortions are unbounded in prospect theory. In standard prospect theory, agents mechanically overweight small probabilities, and in cumulative prospect theory, they overweight the probabilities of extreme events, which are also small for most probability distributions. When humans are bad at distinguishing between the magnitude of two small probabilities, it makes sense that shrinking an already small probability further does not affect the subjective probability by much. Hence, probabilities in prospect theory become more and more overweighted as they get smaller.

repeatedly playing loss-exit strategies in our setup with a finite expiration date and calculate the skewness of the resulting empirical distribution (see Appendix A.6).

Contrast and level effects together imply that a salient thinker adopts a stop-loss and take-profit strategy only if it is a loss-exit strategy. To see why, consider again the case without an expiration date, and assume a drift of zero and a linear value function. In this case, any stop-loss and take-profit strategy is associated with a threshold stopping time $\tau_{a,b}$ and, because the process has zero drift, it induces a binary lottery $X_{\tau_{a,b}} = (a, p; b, 1 - p)$ with an expected value of $\mathbb{E}[X_{\tau_{a,b}}] = x$. A salient thinker, thus, adopts a stop-loss and take-profit strategy only if the upside of the corresponding binary lottery is salient. If $b - x_t \leq x_t - a$, then by the contrast and level effects, the downside state where a is realized is more salient than the upside state where b is realized, which makes this stop-loss and take-profit strategy unattractive to a salient thinker. Conversely, due to the level effect of the salience function, $b - x_t > x_t - a$ does not imply that the lottery's upside b is more salient than the downside a , so that a salient thinker does not find every loss-exit strategy attractive. All arguments extend to processes with a negative drift and to a setup with a finite expiration date, as well as to our salience model where the value function is not linear, but weakly concave. We obtain:

Proposition 2. *If a salient thinker does not stop a process, he always chooses a loss-exit strategy.*

As loss-exit strategies induce positively skewed return distributions (see our argumentation after Definition 5), this proposition implies that a salient thinker is *skewness seeking*. Specifically, while she does not find every strategy leading to a positively skewed outcome distribution or even every loss exit strategy attractive, every strategy that she does find attractive is positively skewed.⁷

A stronger result on the limits of naïve gambling. Using Proposition 2, we can strengthen our result on the limits of naïve gambling: a salient thinker, who is restricted to choose a stop-loss and take-profit strategy, does not start if and *only if* the drift falls below some threshold.

To fix ideas, let us get back to the case of no expiration date, so that any stop-loss and take-profit strategy induces a binary lottery $X_{\tau_{a,b}} = (a, p; b, 1 - p)$ over wealth. For any such strategy, the probability $p = p(a, b, \mu)$, with which the downside of the corresponding binary lottery is realized, monotonically decreases in the drift of the process. Hence, an increase in the drift μ improves the distribution induced by a stop-loss and take-profit strategy in terms of first-order stochastic dominance. By Proposition 1 in Dertwinkel-Kalt and Köster (2020), a salient thinker's certainty equivalent is monotonic with respect to first-order stochastic dominance shifts. This implies that, if a salient thinker is willing to gamble according to stopping time $\tau_{a,b}$ for a drift μ' , then this stopping time is still more attractive than not starting for any drift $\mu > \mu'$. In sum, a naïve salient thinker does not start if and *only if* the drift falls below some threshold.

What happens if we allow for a finite expiration date instead? Because the drift of the process affects the probability of stopping before the expiration date, it is not clear, in general, whether

⁷Notably, in salience theory, it is not the case that when choosing among two positively skewed lotteries, the more skewed lottery is always preferred (which has been formally proven in Corollary 2 in Dertwinkel-Kalt and Köster, 2020).

the distribution of $X_{T \wedge \tau_{a,b}}$ improves in terms of first-order stochastic dominance as the drift increases. For loss-exit strategies, however, an increase in the drift does improve the distribution of $X_{T \wedge \tau_{a,b}}$ in terms of first-order stochastic dominance (Lemma 2 (d) in Appendix A). Hence, by Proposition 2, we can again invoke Proposition 1 in Dertwinkel-Kalt and Köster (2020) to establish that a salient thinker’s gambling behavior is monotonic in the drift of the process.

Proposition 3. *For any expiration date $T \in \mathbb{R}_{>0} \cup \{\infty\}$, any initial wealth level $x \in \mathbb{R}_{>0}$ and any volatility $\nu \in \mathbb{R}_{>0}$, there exists some constant $\tilde{\mu} \in \mathbb{R}$, such that a naïve salient thinker — who is restricted to choose a stop-loss and take-profit strategy — does not start if and only if the drift of the process satisfies $\mu \leq \tilde{\mu}$.*

Proposition 3 differs from Theorem 1 in two aspects: on the one hand, the class of strategies that we consider is more restrictive as we only consider stop-loss and take-profit strategies here; but on the other hand, we obtain in Proposition 3 not just an “if”, but an “if and only if” statement.

Salience theory and the disposition effect. Even if only stop-loss and take-profit strategies are available, so that *planned* behavior is path-independent, salience theory can explain *actual* behavior consistent with the disposition effect; i.e., the tendency to rather stop when the process has increased in value than decreased in value. Our salience-based explanation of the disposition effect is similar in spirit to the CPT-based one by Barberis (2012): it is not the exact path of the process, but the current wealth level that affects a salient thinker’s disposition to stop.

To establish an intuition for when a salient thinker is likely to stop, let us again abstract from an expiration date. A naïve salient thinker stops at time t if and only if, for any $\epsilon, \epsilon' > 0$,

$$\frac{\sigma(v(x_t - \epsilon), v(x_t))}{\sigma(v(x_t + \epsilon'), v(x_t))} \times \frac{v(x_t) - v(x_t - \epsilon)}{v(x_t + \epsilon') - v(x_t)} \geq \frac{1 - p}{p}, \quad (3)$$

where $p = p(\epsilon, \epsilon', \mu)$ denotes the probability of stopping at a loss relative to the current wealth level. Because of the constant drift, the right-hand side of (3) is independent of the current wealth level x_t (see Lemma 1 in Appendix A). If the left-hand side of (3) is increasing in x_t , the salience model, thus, predicts a *disposition effect*: in this case, stopping becomes more likely after the process has increased in value and less likely after it has decreased in value. If the left-hand side of (3) is decreasing in x_t , however, salience theory predicts the exact opposite behavior. In sum, salience theory can rationalize, but does not predict the disposition effect.

While in general we stay agnostic regarding the functional forms of salience and value functions, it could still be interesting to see whether common salience specifications could explain the disposition effect or not. So far, the salience literature has adopted the specification proposed in Bordalo *et al.* (2012), which uses a linear value function and salience function

$$\sigma(x, y) = \frac{|x - y|}{|x| + |y| + \theta} \quad (4)$$

with $\theta > 0$. In Appendix A.5 we examine whether this salience specification predicts the disposition effect and find that this is not the case. Hence, in order for salience theory to predict

the disposition effect, other salience specifications are needed: either other salience functions or other value functions (e.g., piece-wise linear value function reflecting loss aversion, as used in the Online Appendix of Bordalo *et al.*, 2012).

5 An Experiment on Dynamic Gambling Behavior

In this section, we present and discuss our experimental design. A translated version of the experiment is available at: <https://os-experiment-archive.herokuapp.com/demo>.⁸

5.1 Experimental Design

We conducted a pre-registered lab experiment in which subjects had to repeatedly decide at which price to sell different assets. Subjects made their selling decisions in (approximately) continuous time, and they could hold each asset for a maximum duration of 10 seconds. If a subject did not sell an asset within 10 seconds, it was automatically sold at the price reached at the expiration date. We set the initial price of each asset to $x = 100$ Taler, an experimental currency that was converted into € at a ratio of 10:1 at the end of the experiment.

The price of an asset followed an ABM with a drift $\mu \in \{0, -0.1, -0.3, -0.5, -1, -2\}$ and a volatility $\nu = 5$. The price was updated every tenth of a second (i.e. $T = 100$), with the price changes being drawn from a normal distribution with mean μ and variance ν^2 .⁹ Hence, although the implemented price paths are not truly continuous, the incentives provided to the subjects are exactly the same as in the continuous-time model introduced in Section 3. Moreover, while using a discrete number of time periods is necessary for implementation, the process looked smooth, and subjects could not know how many discrete steps it consisted of.

As it is illustrated in Figure 1, we restricted the choice set to all stop-loss and take-profit strategies: at every point in time, subjects could choose an upper and a lower stopping threshold. Once the price of the asset reached either threshold, subjects could decide whether to sell the asset at this price or to adjust the thresholds and continue the process (see the lower left panel); that is, the strategies were non-binding to rule out any form of commitment. Subjects could pause the process at any point in time to adjust the thresholds (see the upper right panel). But, importantly, subjects could set only one upper threshold and one lower threshold at a time, and thus observed a stopping strategy in isolation. At the beginning, the upper and lower thresholds were centered symmetrically around the initial price (see the upper left panel). To start the process, subjects had to move each threshold at least once. Before starting the process, subjects could decide to sell the asset immediately (see also the upper left panel).

Overall, subjects made six selling decisions, one decision for each of the drift parameters. The order of the drifts was randomized at the subject level. It is important for our analysis to make sure that subjects understand before the start of the decision round that a non-positive drift of a process means that they will, on average, not gain money from gambling with this

⁸The experimental design, including the fully specified salience model and its predictions, was pre-registered in the AEA RCT registry as trial AEARCTR-0005359.

⁹Notice that the drift of an ABM is additive over time. To help subjects understand what the drift of a process is, we thus presented them aggregated drifts per second (i.e., 10μ) in the experiment.



Figure 1: Screenshots of the decision screen for the process with zero drift (in German). The text above the chart mentions the drift for this round (“The practice rounds are over now - now it’s getting serious. Please make your selling decision. The drift in this round is 0.”). The red lines indicate the upper and lower stopping thresholds. The blue button in the upper left panel says “Sell Immediately”. The button in the upper right panel allows subjects to pause the process. The buttons in the lower left panel say “Sell” or “Adjust the bounds”. The lower right panel shows the final selling price.

process. If they do not understand this, they might start gambling (even if they have EUT preferences), because they expect to earn money; then they would stop as soon as they have learned from observing the process during the decision round that they are losing money over time.

We have two approaches to rule out such learning about the drift during the decision round: First, at the beginning of each round, we inform the subjects about this round’s drift. If subjects fully understand the implications of a given drift, this would be sufficient to rule out learning. However, subjects may not understand the meaning of a drift, despite our explanations in the instructions. Therefore, we also let subjects watch the development of three sample paths from the underlying process for 10 seconds each. Moreover, we show them an overview of ten additional sample paths of the process (see Figure 2). These thirteen sample paths that subjects see before making a decision (which were randomly drawn at the subject level, meaning that different subjects saw different sample paths of the same underlying process) should give them a quite good understanding of the process and its drift. Thus, we feel confident that substantial learning during the decision round while watching the change of the process is not a driver of our results.

After completing the six stopping problems, subjects faced a series of twelve (static) choices between a binary lottery and the safe option paying the lottery’s expected value with certainty.



Figure 2: Screenshots of the sampling screens for the process with zero drift (in German).

We used two sets of lotteries with the exact same expected value (either €30 or €50) and the exact same variance, but different levels of skewness (see Table 2 in Appendix C for an overview). The order of the lotteries was randomized at the subject level. Finally, subjects answered five questions of a modified cognitive reflection test (CRT; closely aligned to Primi *et al.*, 2016), and the five financial literacy questions proposed by Lusardi and Mitchell (2011). All additional questions are listed in Appendix C.

At the end of the experiment, for each subject, one of the six selling decisions was randomly drawn by the computer to be payoff-relevant. We randomly selected one subject in each session for whom, in addition, one of the twelve static choices was randomly chosen to be payoff-relevant. Subjects were further rewarded for correctly answering CRT and financial literacy questions (1 Taler per question), and they received an additional €4 for their participation.

We conducted 5 sessions with a total number of $n = 158$ subjects. The sessions took place in January 2020 in the experimental laboratory at the University of Cologne. The experiment was conducted using the software oTree (Chen *et al.*, 2016) and participants were invited via ORSEE (Greiner, 2015). The experiment lasted for around 45 minutes on average. Subjects earned on average slightly less than €15, with earnings ranging from €4 to €117.

5.2 Implementation and Discussion of the Design

In this subsection, we provide additional information on the implementation of the experiment, and we discuss in how far our design choices are essential given the objectives of our study.

Explanation of the process. To make the definition of the process easily accessible for subjects, we followed a mostly visual approach. In particular, we did not confront subjects with the differential equation that defines an ABM. Instead we simply told subjects the following:

“In this experiment you will see assets of varying profitability. How profitable an asset is in the long run is described by the drift of the asset. The drift denotes the average change in the value of the process per second.

A positive drift implies that the asset will increase in value in the long run, while a negative drift implies that the asset will decrease in value in the long run. Notice that the value of the asset varies. Hence, even an asset with a negative drift sometimes increases in value.”

To get some understanding of the process and its drift, subjects were presented with exemplary sample paths from three processes with different drifts.¹⁰ Subjects were told that the processes they would see in the experiment differ *only* in their drift. In particular, we told them that all processes have in common that they are non-negative and absorbing in zero.

Finally, to make sure that subjects really understood the stochasticity of an ABM (without confusing them by introducing a formal notion of variance), we told them that

“Independent of the drift, the value of the asset can, in principle, become arbitrarily large. The probability that the asset’s value indeed becomes very large is the smaller the more negative the drift is. But even an asset with a very negative drift can attain a very large value.”

This may raise the concern that subjects could think (at least if they did not carefully read the previous part of the instructions, stating that a negative drift gives, on average, a decrease in value) that even assets with negative drifts are on average a profitable investment. In this case, the total of thirteen sample paths that subjects see for each drift before making their selling decision should give subjects a rough understanding of the expected value of the process.

We regard this part of the instructions as particularly important since the predictions of salience theory rely on the assumption that subjects understand the potential skewness induced by stop-loss and take-profit strategies with a large upper stopping threshold. A translation of the full screen-by-screen instructions is provided in Appendix C.

Features of the process. To make our theory testable, we deviate from Ebert and Strack (2015) in two ways: First, since it is impossible to implement a process that can run forever with probability one, we implemented — similar as Heimer *et al.* (2023) — a finite expiration date. Alternatively, we could have implemented a random termination rule, according to which, at time t , the asset is automatically sold with probability $\omega_t \in [0, 1]$. A finite expiration date makes a theoretical analysis of stopping behavior feasible, while with a random termination rule the probability distribution associated with a given stop-loss and take-profit strategy would not be tractable anymore. A finite expiration date is also easier to explain to the subjects, which we regard — given the complexity of the experiment — as a major advantage. Second, to ensure incentive-compatibility, we make the process absorbing in zero.¹¹ We further restrict the drift of the process to be non-positive because processes with a positive drift do not allow us to separate between the predictions of different models such as EUT, CPT, and salience theory.

Duration of the Process. A potential concern of our experimental design is that the process only runs for 10 seconds, which is significantly shorter than the time horizon of our motivating examples such as stock trading or job search. However, during the 10 seconds processes already change considerably, and this rather short time span makes it easy to visually follow the development of the process. Notably, skewness both intuitively and theoretically affects short- and

¹⁰The sample paths we used in the instructions are exemplary for these processes in the sense that the final values after 10 seconds are 120 (for $\mu = 2$), 100 (for $\mu = 0$), and 80 (for $\mu = -2$), respectively. All subjects saw the exact same sample paths in the instructions.

¹¹In principle, we could have implemented losses up to the size of an endowment that subjects received at the beginning of the experiment. But even then we would have needed to bound the process from below.

long-run behavior likewise. Hence, abstracting from concerns about reaction times (which are discussed below), skewness effects can be studied with short and long processes alike. Notably, subjects can pause the process at any time and thus slow down their decision making. We view the short duration of the process even as an advantage as it allows subjects to stay focused on the task which would become increasingly difficult for longer horizons.

(Approximately) Continuous time. It is not feasible to implement a truly continuous process. Instead, we update the process every tenth of a second by drawing from a normal distribution with mean μ and variance ν^2 . This way the problem of whether to stop after s seconds in our experiment is equivalent to that of stopping an ABM with drift μ and volatility ν at $10 - s$ seconds before the expiration date.¹² Moreover, as the process looked smooth, subjects could never know whether and when only a few time periods were left. This design feature is crucial because it allows subjects to select a strongly skewed return distribution even near the expiration date.¹³ This way our experimental implementation fits well to the continuous time process in our model. Our experimental setup also closely approximates many real live situations where investor or gamblers can always choose strongly skewed return distributions.

Restriction of the choice set. Subjects could choose among all stop-loss and take-profit strategies.¹⁴ This design choice was made based on both practical and theoretical considerations. First, we need an experimental design that allows us to learn something about the actual strategies that subjects choose. When simply providing subjects with a STOP-button, so that they could implement any strategy, we would not learn anything beyond realized stopping times. Stop-loss and take-profit strategies are not only easy to elicit but also enable subjects to choose highly skewed return distributions. This allows us to study the role of skewness in stopping problems. Second, stop-loss and take-profit strategies are highly relevant in practice, which is reflected in the large interest that this type of stopping strategy has attracted in the economics literature (e.g. Xu and Zhou, 2013; Ebert and Strack, 2015; Fischbacher *et al.*, 2017; Heimer *et al.*, 2023).

Non-binding strategies and costless adjustments. We allowed the subjects to costlessly adjust the stop-loss and take-profit thresholds over time: subjects could stop the process at any time, adjust one or both thresholds, and then continue the process. Moreover, the chosen strategies were non-binding in the sense that, once the price of the asset reached either threshold, subjects could decide whether to really sell the asset at this price or whether to adjust the thresholds and continue the process. Again we made both design choices for practical and theoretical reasons.

¹²Subjects could really implement any combination of stop-loss and take-profit thresholds they like. For conciseness, assume that $X_t = 100$, and a subject would like to stop either at 110 or 99. If in the next step the process was updated to say 98, the subject is still paid according to his chosen stop-loss threshold of 99 (unless he revises his strategy to continue gambling). The same is true in case the process “jumps” above the take-profit threshold.

¹³This is not possible in the discrete-time setup in Barberis (2012) or the experimental setup in Heimer *et al.* (2023).

¹⁴We intentionally designed the decision screen, where subjects set a single upper and a single lower threshold (see Figure 1), in a way that makes it hard for subjects to visualize a strategy that does not fall into the class of stop-loss and take-profit strategies. But, even if subjects adopted other strategies, the test of our main theoretical result — namely, Theorem 1 — would still be valid, as here we did not impose any restriction on the choice set.

First, if either strategy adjustments were costly or if the strategies were binding, subjects could partially commit to a strategy. While the commitment effect of costly strategy adjustments is obvious, binding strategies introduce partial commitment when subjects anticipate that they will not be able to adjust their strategy fast enough; namely, before the process hits a threshold. In our main real-world examples, such as selling an asset and gambling in a casino, investors or gamblers have (at best) very limited commitment power (as also demonstrated by Heimer *et al.*, 2023, using brokerage data). Third, since subjects have a non-zero reaction time, non-binding strategies reduce noise in measuring the *intended* stopping time. Preventing this kind of noise seems particularly important, as it would be asymmetric — making stopping disproportionately more likely than non-stopping — and hard to model. This improves the fit between our experiment and our model, where stopping results from the unavailability of an attractive threshold stopping strategy. In the model, the agent chooses infinitely quick in continuous time, which is not feasible for the experimental subjects. However, upon the process hitting one of the thresholds subjects can take as much deliberation time as they need to figure out whether they want to continue gambling.

Importantly, even though strategy adjustments are costless, the exact thresholds are important and should be carefully set by the subject right from the beginning. The stop-loss threshold, for instance, gives a lower bound on the value that the process can reach and, therefore, should not be set below the level that the subject would not want to undercut. Likewise, the take-profit threshold should not be set too high, as otherwise moderate gains cannot be cashed in. Since the value of the process changes in (almost) continuous time, but subjects are not able to adjust the thresholds in continuous time, choosing the “right” thresholds to begin with is important. Subjects could, however, start with bounds that are tighter than the ones they actually want to play with, and plan to adjust them once one bound is hit. As this, however, involves an extra effort without any benefit, we would not think such behavior is a dominant force in our experiment.

Indicators of naïvete. When assuming a fixed expiration date and restricting the choice set to all stop-loss and take-profit strategies, we cannot interpret adjustments of the initial strategy as time-inconsistent behavior and thus as an indication of naïvete, since the remaining time until the expiration date conveys payoff-relevant information. Looking at processes with a non-positive drift, however, allows us to test the naïvete assumption within the salience framework.

A sophisticated salient thinker differs from her naïve counterpart in that she anticipates her future selves to act in a different way than her present self does, which she takes into account when making her stopping decision. A sophisticated salient thinker who lacks commitment then behaves as if she was playing a game with her future selves (e.g., Karni and Safra, 1990). To solve this game, we adopt the equilibrium concept of Ebert and Strack (2018), according to which a given stopping strategy constitutes an equilibrium if and only if at every point in time it is optimal to follow this strategy, taking as given that all future selves will do so.

As we show in Appendix B, a sophisticated salient thinker, who lacks commitment and chooses from the set of all stop-loss and take-profit strategies, does not start any process with a non-positive drift. Consequently, (partial) naïvete is a necessary assumption to rationalize gambling in the context of our experiment within the salience framework.

5.3 Experimental Predictions

We now state the precise predictions of salience theory that guided our experimental design. We slightly deviate from our pre-registration, which was based on the salience model with a linear value function: due to the weakly concave value function, Prediction 2 differs from what was pre-registered, while the remaining predictions are identical to the pre-registered ones.

At its core, our dynamic salience model is built to explain skewness-seeking behavior in dynamic choices under risk. Since subjects are restricted to play stop-loss and take-profit strategies in our experiment, we hypothesize (based on Proposition 2) that they play loss-exit strategies. This prediction is not shared by EUT, but by other models on skewness seeking such as CPT and models on disappointment aversion.

Prediction 1. *Conditional on not selling the asset, subjects choose a loss-exit strategy.*

The main theoretical contribution of our model, relative to existing models on skewness effects, is that it can rationalize gambling with moderately negative expected values while ruling out that agents accept any sufficiently skewed gamble regardless of how unfair it is. Therefore, our model also makes two predictions on the relation between the drift of the process and the subjects' decisions to start gambling.

Prediction 2. *If $\mu = 0$, subjects might start to gamble, and they might stop before the expiration date.*

Prediction 3. *The share of subjects selling the asset immediately monotonically decreases in the drift.*

By Corollary 1, salient thinkers may or may not start to gamble. Due to the complicated cumulative distribution function that emerges from our process with a finite expiration date and the play of stop-loss and take-profit strategies (see Lemma 2 in den Appendix), we cannot formally prove that in-between stopping is possible. But with the help of simulations, we can give suggestive evidence that salient thinkers that start to gamble need not gamble until the expiration date. The reason for this is that hitting some fixed bounds becomes less likely as time passes, making the process's distribution more symmetric and therefore less attractive for a salient thinker. We thus obtain Prediction 2. Prediction 3 directly follows from Proposition 3. These predictions distinguish our model from EUT with a concave utility function, which does not yield Prediction 2. It also distinguishes our model from CPT as modelled by Ebert and Strack (2015) as well as from models on disappointment aversion, both of which do neither yield Prediction 2 nor Prediction 3. Therefore, we regard it as important to test these predictions, even if they are very intuitive.

Lastly, since we extend a theory of "static" choice under risk to a dynamic setup, we are interested in the empirical relationship between a subject's attitude toward static and dynamic risks. If salience is indeed the psychological mechanism underlying our results, it should coherently explain behavior revealed in static and dynamic choices. As we show in Appendix D, a salient thinker chooses a binary lottery, with a fixed expected value and a fixed variance, over the safe option paying its expected value if and only if the lottery's skewness exceeds a certain threshold. By Proposition 2, this preference for positive skewness is also what drives a salient thinker's stopping behavior. We therefore classify both static and dynamic choices into being

skewness-seeking or not. We say that a static choice is skewness-seeking if the subject chooses a right-skewed lottery over the safe option or the safe option over a left-skewed or symmetric lottery (see Table 2 in Appendix C for the exact lotteries and classification). We further classify a stopping strategy as being skewness-seeking if it is a loss-exit strategy and, thus, induces a right-skewed return distribution. Based on Proposition 2, we expect a positive correlation between the share of skewness-seeking choices in static and dynamic decisions.

Prediction 4. *The share of skewness-seeking choices by a subject in the static decisions is positively correlated with the share of loss-exit strategies this subject chooses in the dynamic decisions.*

This prediction is not shared by any of the alternative models that we discuss throughout this paper, either because they do not predict skewness seeking (as EUT and reference-dependent preferences without probability weighting) or because they cannot explain the heterogeneity in gambling behavior that we observe (as it is the case for CPT and models of disappointment aversion).

6 Experimental Results on Dynamic Gambling Behavior

We first describe our data. Subsequently, we present our main experimental results, as well as evidence on how subjects revise their strategies over time, and we discuss in how far this speaks to the salience mechanism that drives our predictions on stopping behavior. Finally, we present exploratory results on disposition-effect-like behavior, and on the role of cognitive skills.

6.1 Data and Descriptive Statistics

For all subjects and all processes, our data includes the choice whether to start the process as well as all chosen stop-loss and take-profit strategies (consisting of an upper bound and a lower bound). We also record the times when each strategy was chosen and the value of the process at each point in time. From these values, we can calculate at which time, if ever, subjects stopped a process, as well as the distance between the two bounds and the process at each point in time.

Table 1 shows descriptive statistics for the data from our experiment. We can see that the share of subjects who sell the asset immediately increases as the drift becomes more negative. The share of subjects who do not sell the asset before it expires and the average time the asset was held decreases as the drift becomes more negative. Moreover, the upper bounds are, on average, further away from the current value of the process than the lower bounds for all drifts. The termination value also decreases with the drift except for the comparison between the drifts -1 and -2. The higher termination value for a drift of -2 is likely driven by the higher fraction of subjects selling the asset immediately (i.e., at a value of 100).

Drift:	0.0	-0.1	-0.3	-0.5	-1.0	-2.0
% Sold Immediately	5.70%	9.49%	13.29%	22.15%	27.22%	41.14%
% Never Sold	18.99%	8.86%	9.49%	5.06%	3.80%	1.27%
Termination Value	95.45 (40.19)	92.34 (38.99)	87.98 (33.63)	81.89 (32.59)	73.97 (34.03)	82.74 (26.70)
Stopping Time	6.31 (3.72)	5.03 (3.69)	3.77 (3.72)	2.99 (3.59)	2.04 (3.16)	0.74 (1.85)
Upper Bound	136.92 (27.99)	133.57 (28.47)	124.97 (29.51)	124.93 (25.22)	122.64 (26.63)	122.14 (26.80)
Lower Bound	71.82 (31.23)	75.88 (33.82)	68.40 (31.38)	71.62 (33.15)	66.42 (32.24)	77.75 (30.04)
Distance Lower Bound to Value	19.02 (18.85)	16.80 (17.80)	19.38 (20.61)	18.52 (20.33)	19.99 (22.52)	14.94 (21.07)
Distance Upper Bound to Value	33.76 (27.74)	28.86 (27.02)	26.74 (22.40)	24.55 (25.98)	25.47 (27.02)	24.99 (31.30)

Table 1: The table shows descriptive statistics for our experimental data. The values without parentheses are the means. The values in parentheses are the standard deviations. Each column shows data for one drift. “% Sold Immediately” is the percentage of subjects that sold the asset immediately and hence never started gambling. “% Never Sold” is the percentage of subjects who held the asset until the expiration date. The stopping time is either the time at which a subject sells the asset after it hits one of the bounds or it equals the value of 10 seconds, the time after which the asset process expires and the asset is sold automatically. The values for the upper and lower bounds include one data point for the initial bounds set by a subject before he can start the process and a data point for each time a bound was adjusted. Similarly, the variables “Distance Upper/Lower Bound to Value” include the (absolute) distance between a bound and the current value of the process every time a bound is adjusted.

6.2 Main Test of our Salience Predictions

First, we show that, consistent with Prediction 1, a majority of subjects initially chooses a loss-exit strategy. This result on initial strategies holds across all the different drifts that we considered (see Figure 12 in Appendix E).

Result 1 (a). *Conditional on not selling immediately, 65% of initial strategies are loss-exit strategies.*

We also perform a t-test with standard errors clustered at the subject level and confirm that the share of subjects who initially choose loss exit strategies is significantly above 50%, implying that, on average, subjects are skewness seeking.

When aggregating all the strategies a subject has chosen throughout the experiment (including both initial and revised strategies), we observe that a majority of the subjects predominantly chooses loss-exit strategies and that 17% of the subjects pick exclusively loss-exit strategies (see Figure 3 for the distribution across all subjects). This gives the second part of our result on chosen strategies:

Result 1 (b). *For the median subject, 73% of all strategies chosen throughout the experiment are loss-exit strategies.*

Overall, these results suggest that the majority of subjects are skewness seeking as predicted by our model. The fact that not all selected strategies are right-skewed can partially be explained

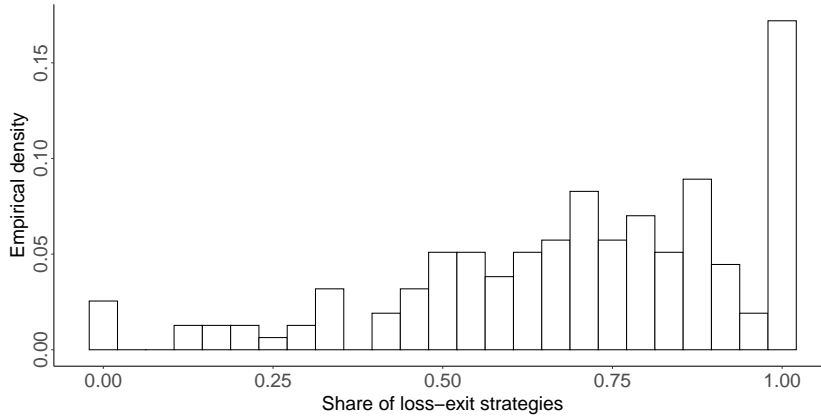


Figure 3: The figure depicts the empirical distribution of the share of loss-exit strategies across subjects. The share is calculated on the subject level by taking all strategies, including both initial and revised strategies, aggregating across different drifts, and determining the percentage of those strategies that are loss-exit strategies.

by inherent noise in experimental data collection, and partially by subject heterogeneity (meaning that not all subjects are salient thinkers).

Our next result — as depicted in Figure 4 — is a monotonic relationship between the drift of the process and a subject’s stopping behavior. Specifically, subjects do gamble (even if the drift is negative), but their behavior is sensitive to the drift of the process. At the hand of Figure 4, we will successively discuss the results corresponding to Predictions 2 and 3.

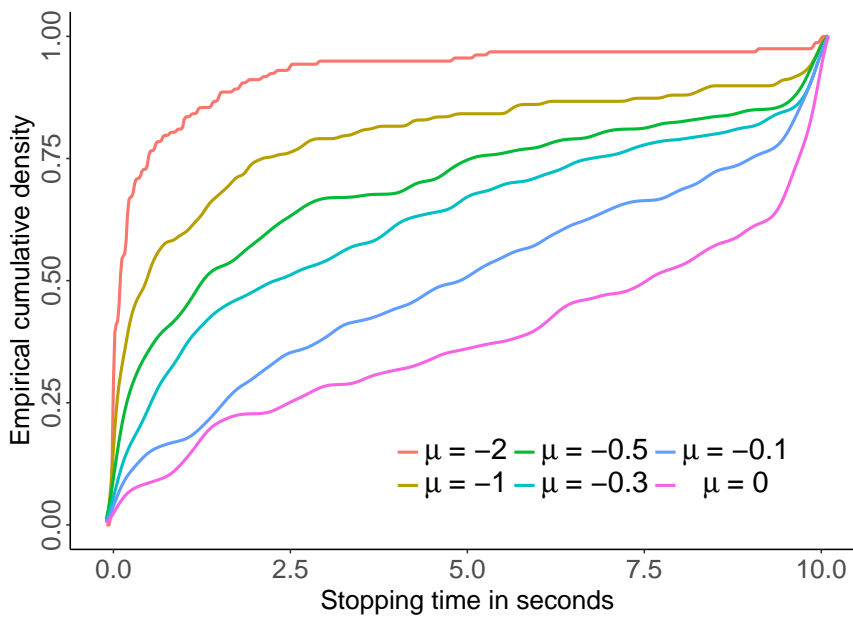


Figure 4: The figure depicts the smoothed empirical cumulative distribution functions of stopping times, one for each of the different drifts.

To address Prediction 2, we look into stopping behavior for the fair process with zero drift. Around a fifth of all subjects hold the asset with a drift of zero until the expiration date, while only about 5% of all subjects sell the asset with a drift of zero immediately. Moreover, 65% of the subjects hold this asset for more than 5 of the maximal 10 seconds. We summarize:

Result 2. *The median subject holds the fair asset with zero drift for 7.65 out of 10 seconds, and around 19% of the subjects hold the fair asset until the expiration date.*

While salience theory does not make a precise prediction on when subjects stop a fair process with zero drift (see Prediction 2), Result 2 is clearly inconsistent with both, EUT with a concave utility function — which predicts that subjects sell the asset immediately — as well as with CPT as modelled by Ebert (2015) — which predicts that subjects will hold the asset until the expiration date.

Next, we investigate how the drift affects a subject’s decision whether to start a process and, after starting, when to stop it. The share of subjects selling immediately monotonically decreases in the drift of the process (see the right panel of Figure 10 in Appendix E). Drift-sensitive stopping behavior is consistent with Prediction 3: estimating a linear probability model indicates that increasing the drift by one unit reduces the average probability of selling immediately by 17.1 p.p. (p -value < 0.001 , standard errors clustered at the subject level).

Result 3 (a). *The share of subjects selling immediately monotonically decreases in the drift of the process.*

This result is also clearly inconsistent with EUT and CPT, both of which predict that the share of subjects who sell the asset immediately is constant in the drift, either because subjects should always sell immediately (EUT) or always gamble until the expiration date (CPT).

Figure 4 further shows that not only the share of subjects selling the asset immediately is monotonic in the drift, but that the whole distribution of stopping times shifts upward in the sense of first-order stochastic dominance as the drift increases.¹⁵

Result 3 (b). *Subjects stop earlier for processes with more negative drifts.*

These results provide valuable insights into the strength of skewness effects. In our model, agents are intrinsically risk averse, so that they will not start a process with a non-positive drift unless they can implement a strategy that induces a skewed outcome distribution. When deciding whether to start a process with a drift of zero, subjects face a trade-off between the variance of the return distribution — which they dislike — and the skewness of the return distribution — which they can select with the right bounds and which they like. We find that, interpreted through the lens of our model, for 95% of subjects skewness seeking is strong enough to make them gamble. But the more negative the drift of the process is, the stronger subjects’ skewness seeking needs to be to render gambling attractive. Hence, our finding that the share of subjects who start to gamble monotonically decreases in the drift of the process, substantiates heterogeneity in the strength of skewness seeking.

We further look into whether subjects hold the process until it expires. Conditional on starting, most subjects do not hold the processes until the expiration date. Even for the fair process only 19% do so, which is consistent with salience theory but conflicts with models such as CPT that predict that agents will never stop gambling regardless how negative the drift of the process is.

¹⁵This is only violated for the processes with a drift of $\mu = -0.1$ and $\mu = -0.3$ in very few points, so that these violations are not even visible in the smoothed CDFs depicted in Figure 4.

Finally, we study whether skewness seeking in static and dynamic decisions is related. As depicted in Figure 5, subjects behave quite consistently in the static and the dynamic decision problems. To test for the link between static and dynamic skewness seeking, we regress the share of loss-exit strategies amongst all strategies chosen throughout the six dynamic problems on the share of skewness-seeking choices in the twelve static problems. We find a positive and statistically significant correlation, which gives our fourth result:

Result 4. *The share of skewness-seeking choices by a subject in the twelve static decisions is positively correlated with the share of loss-exit strategies this subject chooses in the six selling decisions.*




Figure 5: The figure depicts the relationship between static and dynamic gambling behavior. We further provide the estimated slope-coefficient of the depicted linear regression, which is significantly larger than zero. The data points are scattered for illustrative purposes.

One might be concerned that Result 4 conflicts with the “discrepancy” between static and dynamic risk taking documented in Heimer *et al.* (2023). Heimer *et al.* (2023), however, compare the willingness to take risk when a fair coin is flipped once (i.e., in a static choice) — which induces a *symmetric* distribution of returns — and when a fair coin is flipped repeatedly (i.e., in a dynamic choice) — in which case the right stopping strategies allow to create very *right-skewed* distributions of returns. In other words, while we study the relationship between skewness seeking in static and dynamic decisions, Heimer *et al.* (2023) look at the difference in behavior between static and dynamic problems that results from the fact that the latter enables subjects to choose a skewed distribution of returns.

6.3 On the Salience-Mechanism: Frequency and Direction of Strategy Adjustments

Consistent with our model, strategy revisions are ubiquitous and follow precise patterns. Altogether, (i) more than 93% of the subjects (147 out of 158) revised their initial strategy in at least one of the six selling tasks, (ii) conditional on starting, subjects adjust their strategies 1.6 times per task, and (iii) about 70% of the strategy adjustments happen in an attempt to prolong

gambling after the process has hit one of the previously chosen thresholds. Moreover, if a subject chooses a loss-exit strategy and the process hits a threshold, the subject is — conditional on not stopping the process — more than six times as likely to again choose a loss-exit instead of a gain-exit strategy (see the left table of Figure 13 in Appendix E),¹⁶ which is consistent with Prediction 1.

Conditional on not selling the asset immediately, around 45% of the processes are stopped “later” than when the subject initially planned to stop the process; that is, 45% of the processes pass (at least) one of the initial thresholds without being stopped. Notably, the share of processes being stopped later than initially planned monotonically increases in the drift of the process, from 20% (for the most negative drift) to around 54% (for zero drift). This suggests that subjects revise their strategies, as predicted by our model, and the fact that this behavior is more pronounced for processes with a less negative drift is again in line with our salience model’s prediction that subjects are sensitive to the drift in a “reasonable” way.

We further observe that 35% of the processes fall below the initial stop-loss threshold, but only 12% of the processes rise above the initial take-profit threshold. Taken together these results indicate exactly the type of strategy revisions that our model (and also the model by Ebert and Strack, 2015, in an extreme form) builds on in order to explain excessive gambling: subjects choose loss-exit strategies and thereby positively skewed return distributions, and then adjust strategies as soon as the stop-loss threshold is hit in order to continue gambling with a newly chosen loss-exit strategy. In sum, the findings on strategy adjustments indicate that our model gives a quite accurate description of the mechanism underlying our main experimental results.

6.4 On the Disposition to Stop and the Role of Cognitive Skills

Subjects reveal a disposition effect. As alluded to before, more processes fall below the initial stop-loss threshold (namely, 35%) than rise above the initial take-profit threshold (namely, 12%). Keeping the asset “too long” (compared to the subject’s initial plan) when the process has decreased in value rather than increased in value is reminiscent of the disposition effect, whereby assets are rather sold in the gain domain and rather held in the loss domain.

Another test for the disposition effect is to compare the likelihood of selling assets that have gained a particular amount to that of selling assets that have lost exactly the same amount: by the disposition effect, the former assets should be more likely to be sold than the latter, which is precisely what we find. Those subjects, who have revised their initial strategy for a respective process at least once, are, on average, more likely to sell a process at value $100 + x$ than one at value $100 - x$ (see Figure 6). To make selling decisions comparable, we here consider only processes with a drift of zero, for which gains and losses are equally likely.¹⁷

As we discussed in Section 4.3, the disposition effect is consistent with salience theory, but not with the standard specification of the salience model that the salience literature has adopted

¹⁶Strategy adjustments conditional on not hitting a threshold follow a similar pattern. Suppose that in the moment of pausing the process the currently played strategy is a loss-exit strategy. Then, all of our subjects have selected another loss-exit strategy and no one has switched to a gain-exit strategy (right table in Figure 13 in Appendix E).

¹⁷A similar picture also arises if all selling decisions including those for processes with a negative drift are taken into account; due to losses being more likely for negative drifts, however, the interpretation of the respective findings is less clear, which is why we focus on the fair processes here.

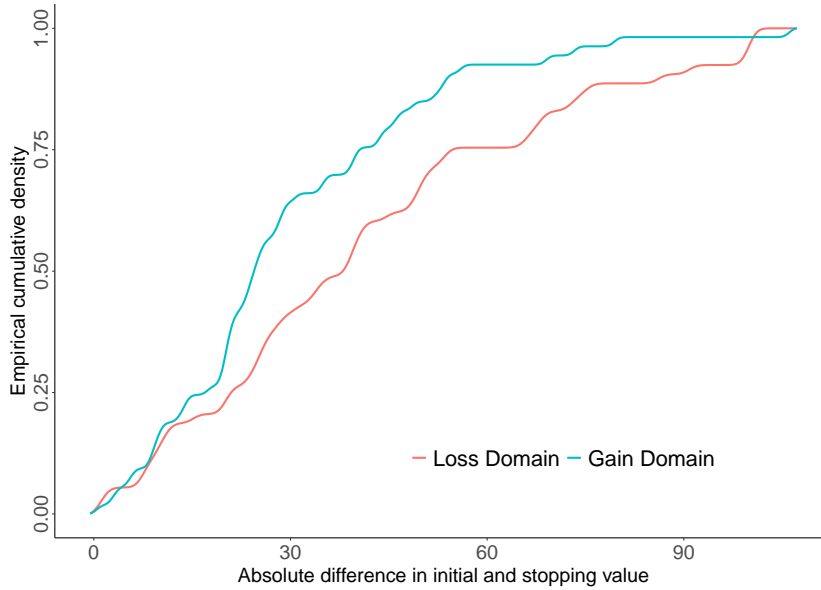


Figure 6: The figure depicts the smoothed empirical CDF of stopping at a given distance to the initial value of the process, separately for processes that have gained and that have lost in value. We consider only fair processes with a drift of zero for which the initial strategy was adjusted at least once.

so far.

Cognitive skills matter. Below-median subjects in terms of cognitive skills — as measured by the sum of correct answers to the modified CRT and the financial literacy questions — are particularly likely to gamble in our experiment. For instance, for the process with a drift of zero, the share of below-median subjects holding the asset until the expiration date is twice as large as the share of above-median subjects doing so (see Figure 11 in Appendix E). Notably, both the below- and above-median subjects support Prediction 3: both are responsive to a change in the drift of the process.

6.5 Limitations of the model in explaining the data

While the majority of subjects chooses mostly strategies inducing right-skewed return distributions, strategies inducing left-skewed return distributions are, unlike what our model predicts, also chosen. This can be explained by noise in the data and by heterogeneity in subjects' susceptibility to salience.

Another issue is that, despite the good fit of our model and our data, we cannot conclusively show that salience indeed is the mechanism that drives our results. Given the strong correlation between skewness effects revealed in static and in dynamic choices, there arguably is one cognitive mechanism that drives all of these skewness effects, but this doesn't have to be salience. In any case, we think of the model as a useful "as if" model.

7 Discussion of Alternative Models

7.1 Expected Utility Theory

In order to explain basic findings in choice under risk — such as an aversion toward symmetric mean-preserving spreads — EUT needs to assume a strictly concave utility function (see, e.g., Bernoulli, 1738; Rothschild and Stiglitz, 1970). Under this assumption, however, EUT predicts that all assets with a non-positive drift will be immediately sold, and it, thus, cannot explain why subjects start to gamble in our experiment (see Result 2 in Section 6).

In order to rationalize Result 2 via EUT, we would need to assume that the utility function is convex over at least some range around the initial value of the asset. But, even if we would allow for a completely flexible utility function, which switches back-and-forth from being concave to being convex, EUT cannot explain the skewness-dependence of risk attitudes, as elicited in the static choices between a binary risk and its expected value: here, subjects seek, for *different* outcome levels, sufficiently right-skewed risks, but avoid left-skewed risks (see Figure 14 in Appendix E). While EUT could, in principle, rationalize this behavior for *one* outcome level via a utility function that is concave first and then becomes convex, it cannot do so for multiple outcome levels, as the inflection point from concave to convex would have to change with the outcome level. Salience theory, in contrast, predicts skewness-dependent risk attitudes for *any* outcome level (see Appendix D and Dertwinkel-Kalt and Köster, 2020), and is thus consistent with the data. Moreover, EUT — in contrast to salience theory — does in general not explain why subjects prefer loss-exit strategies over gain-exit strategies (Result 1 in Section 6). In sum, EUT cannot coherently explain our findings on static and dynamic risk attitudes.

7.2 Cumulative Prospect Theory

Abstracting from a finite expiration date, Ebert and Strack (2015) have shown that, under empirically weak assumptions on the probability weighting function, a CPT-agent will never stop an ABM, irrespective of how negative its drift is. This stark never-stopping result follows from the fact that the preference for positive skewness induced by common CPT-specifications is so strong that the naïve CPT-agent can always find a stop-loss and take-profit strategy that is more attractive than not starting. As we numerically show in Appendix F, at the example of the representative CPT-agent proposed by Tversky and Kahneman (1992),¹⁸ the never-stopping result extends to processes with a finite expiration date. Consequently, common specifications of CPT can neither rationalize the fact that subjects stop a process with zero drift before the expiration date (i.e., Result 2) nor that stopping behavior is sensitive to the drift of the process (i.e., Result 3).¹⁹ As a consequence of this never-stopping result, CPT is also inconsistent with the disposi-

¹⁸It is easily verified that the stark never-stopping result extends to finite expiration dates also for other common CPT-specifications. But, for expositional convenience and in line with the related literature (Barberis, 2012; Heimer *et al.*, 2023), we focus on the representative CPT-agent based on the estimates by Tversky and Kahneman (1992).

¹⁹CPT belongs to the class of rank-dependent utility models (see, e.g., Quiggin, 1982), which do not assume in general, however, that behavior is reference-dependent and affected by loss aversion. As the never-stopping result of CPT does not rely on either reference-dependence or loss aversion, it extends to a larger class of models within the RDU-family (as shown by Duraj, 2020). But, due to the flexibility of rank-dependent utility models, we do not obtain general predictions regarding the stopping behavior of an RDU-agent in our setup.

tion effect in a setting like ours (or the one by Ebert and Strack, 2015, as they argue). CPT can, however, also account for Result 1: because a CPT-agent overweights the tails of a probability distribution, he likes the right-skewed distribution generated by loss-exit strategies (this has been also shown in Barberis, 2012; Ebert and Strack, 2015; Heimer *et al.*, 2023).²⁰

7.3 Reference-Dependent Preferences without Probability Weighting

Barberis and Xiong (2009, 2012) propose an explanation of the disposition effect based on a version of prospect theory *without* probability weighting, according to which gains and losses are experienced at the level of an individual asset in the moment of selling it.²¹ Moreover, Barberis and Xiong (2012) derive results that are seemingly similar to the drift-sensitivity of a naïve salient thinker that we establish in this paper. This apparent similarity, however, is driven by the different setup that they analyze: to establish their result, Barberis and Xiong assume, in particular, that (1) upon selling an asset the agent can immediately reinvest his wealth in another asset, (2) when selling an asset the agent pays positive transaction costs, and (3) the time-horizon is sufficiently long for discounting to play an important role. Our experimental design shares neither of these features, so that their results cannot be applied to our setting. Using a stylized version of the model by Barberis and Xiong (2012), we demonstrate in the following that their *realization-utility* approach, which has found some experimental support (e.g., Imas, 2016), cannot account for our experimental findings.

Without loss of generality, we abstract from a finite expiration date and from discounting. Adapting the model in Barberis and Xiong (2012) to our setup, we assume that the agent’s utility is given by the sum of an asset’s net present value and her realization utility from selling the asset, where the latter is given by a (piece-wise) linear function $u(\cdot)$ defined as follows: $u(x) = x - r$ if $x \geq r$ and $u(x) = \lambda(x - r)$ if $x < r$ for some loss-aversion parameter $\lambda \geq 1$ and a reference point $r = x_0$.²² The agent’s utility derived from selling the asset at time t is equal to

$$\underbrace{X_t}_{\text{net present value}} + \underbrace{u(X_t)}_{\text{realization utility}}$$

Now consider a threshold stopping time $\tau_{a,b}$ with $a < x_0 < b$, and denote by $p = p(a, b, x_0)$ the probability that the process is stopped at the stop-loss threshold a . The agent sells the asset immediately if and only if, for any such threshold stopping time, it holds that

$$\underbrace{pa + (1 - p)b}_{\text{expected net present value}} + \underbrace{p\lambda(a - x_0) + (1 - p)(b - x_0)}_{\text{expected realization utility}} \leq x_0,$$

²⁰The most striking difference between salience theory and alternative approaches — such as EUT and CPT — is that it predicts behavior to be context-dependent in the sense that the evaluation of a given option depends on the alternatives at hand. We ran an additional (pre-registered) experiment that documents context-dependent stopping behavior in line with salience theory. To focus on our main results, and not to disrupt the flow of the main text, we decided to relegate this additional experiment to Appendix G, however.

²¹Barberis and Xiong (2009) show that other, more common reference point specifications (such as annual gains and losses) do not allow CPT to explain the disposition effect.

²²Precisely, the case of $\lambda = 1$ refers to Eq. (7) in Barberis and Xiong (2012), while $\lambda > 1$ corresponds to Eq. (18) in their paper.

or, equivalently,

$$2(1-p)(b-x_0) \leq (1+\lambda)p(x_0-a). \quad (5)$$

A sufficient condition for Eq. (5) to hold is that $(1-p)(b-x_0) \leq p(x_0-a)$ or, equivalently, $\mathbb{E}[X_{\tau_{a,b}}] \leq x_0$, which is satisfied for any process with a non-positive drift. We conclude that an agent with realization utility à la Barberis and Xiong (2012) would immediately sell any asset in our experiment; that is, their model can neither account for Result 2 nor Result 3.²³

More generally, the preceding analysis highlights that some form of non-linear probability weighting is necessary to explain our results on skewness seeking, not only in the dynamic selling decisions, but also in the static choices, which we analyze in Appendix E (see Figure 14). The former point is made in an informal way also in Heimer *et al.* (2023). Adding non-linear probability weighting to the model by Barberis and Xiong (2012) would yield a model that is essentially equivalent to the ones studied in Barberis (2012) or Ebert and Strack (2015), which we have already discussed in detail in the previous subsection.

7.4 Disappointment Aversion

Gul (1991) proposes a theory of disappointment aversion to explain the Allais paradox, in particular, the certainty effect.²⁴ The model can, in principle, rationalize skewness seeking and thereby gambling in the context of our experiment (Duraj, 2020, Proposition 4). But, as we will formally argue in the following, under the assumptions necessary to explain skewness seeking, it also predicts that subjects will not stop a process with zero drift before the expiration date, which is inconsistent with Result 2.

If we abstract from a finite expiration date (i.e., if $T = \infty$ holds), a disappointment-averse agent values the random variable induced by a threshold stopping time $\tau_{a,b}$ at

$$V(X_{\tau_{a,b}}) = \frac{p(1+\beta)}{1-p+p\beta}u(a) + \frac{1-p}{1-p+p\beta}u(b),$$

where u is a classical utility function and $\beta > -1$ captures the agent's disappointment aversion.

As illustrated in Gul (1991), we need to assume $\beta > 0$ in order to rationalize puzzling behavior like the Allais paradox. But, given that $\beta > 0$, the only way to rationalize a preference for sufficiently right-skewed risks is to assume a convex utility function $u(\cdot)$. Precisely, with a concave utility function, the disappointment-averse agent would reject any fairly priced risk, and he would thus sell any asset with a non-positive drift immediately, which contradicts both our results on dynamic (i.e., Result 2) and static choices (see Figure 14 in Appendix E).

So, let us assume not only that $\beta > 0$, but also that $u(\cdot)$ is convex. As in our experiment, we

²³As the setup in Barberis and Xiong (2012) shows substantial differences to our setup—for instance, asset selling goes along with substantial transaction costs—this is no contradiction to their Figure 1. This does also not change if we use a different variant of their model, namely one where we drop the “expected net present value” term. In that case, only gain-loss utility prevails, and as losses loom larger than gains a symmetric process (or one with a negative drift) cannot be attractive.

²⁴Disappointment aversion is a special case of cautious expected utility (Cerrei-Vioglio *et al.*, 2015), which is so flexible, however, that it can explain basically any kind of stopping behavior, including the stark never-stopping result predicted by CPT (see Proposition 6 in Duraj, 2020).

assume that the agent can only choose stop-loss and take-profit strategies. A disappointment-averse agent stops a process with zero drift at time t , if and only if, for any stopping time $\tau_{a,b}$,

$$\frac{\frac{u(b)-u(x_t)}{b-x_t}}{\frac{u(x_t)-u(a)}{x_t-a}} \leq 1 + \beta.$$

Since $u(\cdot)$ is convex by assumption, the left-hand side of the preceding inequality is strictly increasing in b (and strictly decreasing in a). Again since $u(\cdot)$ is convex, for any fixed $a \geq 0$, the left-hand side approaches infinity, as b approaches infinity. But this implies that, for any fixed $\beta > 0$, we can find a finite b , such that the above inequality is violated. Consequently, a disappointment-averse agent with a convex utility function never stops a process with zero drift, which contradicts the fact that a large majority of subjects stop the process with zero drift before the expiration date (i.e., Result 2). All the preceding arguments carry over to the case of a finite expiration date. In sum, we conclude that a model of disappointment aversion cannot coherently explain the findings on skewness seeking in static and dynamic settings.

8 Conclusion

While we find that people take up symmetric gambles if they can obtain skewed return distributions through the choice of their stopping strategies, theoretical considerations suggest similar behavior in case the underlying process is skewed itself. On the one hand, even if the underlying process is negatively skewed, the return distribution associated with the “right” loss-exit strategy is again positively skewed (Ebert, 2020). And if the process itself is positively skewed — which is indeed the case in many real-world applications — our results are likely to be amplified.

A first example refers to processes underlying many casino gambles (as discussed in Ebert and Strack, 2015, Section V) and many asset values, which are not symmetric, but positively skewed. Skewness seeking, as modelled by salience theory, then suggest that consumers gamble or over-invest all the more, as the skewness created with their stopping strategies is exacerbated by the skewness of the process. As an alternative example, we could think of teenagers or young adults who decide whether to pursue the career of a professional athlete, actor, or musician. While the probability of actually making it to the professional level is small, it requires substantial investments of time and other resources to take the shot at becoming a superstar. A teenager who practices excessively for a particular sport, for instance, might as a result neglect school or studies, thereby lowering the attainable wage in the likely case that he fails to become a professional athlete. Now suppose that, as suggested by our model, this teenager adopts the following strategy: each year, he hopes for a breakthrough, but plans to quit on sports and instead study otherwise. This strategy generates a positively skewed return distribution, which can be particularly appealing due to the skewness that is inherent to the process of becoming a superstar. After each failure, however, the teenager revises his plans and decides to try it for *one more year*, as this way he can again experience a right-skewed distribution of returns. This idea of excessively pursuing a career is not only consistent with our model, but it is also supported

by empirical studies (e.g., Choi *et al.*, 2022; Grove *et al.*, 2021). A similar type of argument applies to the classical problem of searching for a job, one of our introductory examples from the classical stopping literature. Here, skewness seeking can explain why people pass on too many mediocre jobs, thereby forgoing pay over a longer time horizon, in the hope of finding one of very few outstanding jobs with excellent pay. Also in this example the skewness of the return distribution that results from the chosen stopping strategy is complemented by the skewness of the process itself. In sum, skewness seeking can explain time-inconsistent behavior in trying to reach an elusive goal.

References

- BARBERIS, N. (2012). A model of casino gambling. *Management Science*, **58** (1), 35–51.
- and XIONG, W. (2009). What drives the disposition effect? An analysis of a long-standing preference-based explanation. *Journal of Finance*, **64** (2), 751–784.
- and — (2012). Realization utility. *Journal of Financial Economics*, **104** (2), 251–271.
- BERNOULLI, D. (1738). Specimen theoriae novae de mensura sortis. *Comentarii Academiae Scientiarum Imperialis Petropolitanae*, **5** (2), 175–92 [translated by L. Sommer (1954) in *Econometrica*, **22**, 23–36].
- BORDALO, P., GENNAIOLI, N. and SHLEIFER, A. (2022). Saliency. *Annual Review of Economics*, **14** (1), 521–544.
- , NICOLA GENNAIOLI and SHLEIFER, A. (2012). Saliency theory of choice under risk. *Quarterly Journal of Economics*, **127** (3), 1243–1285.
- CERREIA-VIOGLIO, S., DILLENBERGER, D. and ORTOLEVA, P. (2015). Cautious expected utility and the certainty effect. *Econometrica*, **83** (2), 693–728.
- CHEN, D., SCHONGER, M. and WICKENS, C. (2016). oTree - an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, **9**, 88–97.
- CHOI, D., LOU, D. and MUKHERJEE, A. (2022). Superstar firms and college major choice. https://www.dropbox.com/s/00ce456x151nuwx/Education_ChoiLouMukherjee.pdf?dl=0, unpublished.
- COX, D. R. and MILLER, H. D. (1977). *The Theory of Stochastic Processes*. Chapman and Hall Ltd.
- DEWINKEL-KALT, M. and KÖSTER, M. (2020). Saliency and skewness preferences. *Journal of the European Economic Association*, **18** (5), 2057–2107.
- DHAMI, S. (2016). *The Foundations of Behavioral Economic Analysis*. Oxford University Press.
- DURAJ, J. (2020). Optimal stopping with general risk preferences. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2897765, unpublished.

- EBERT, S. (2015). On skewed risks in economic models and experiments. *Journal of Economic Behavior & Organization*, **112**, 85–97.
- (2020). On taking a skewed risk more than once. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3731565, unpublished.
- and STRACK, P. (2015). Until the bitter end: On prospect theory in a dynamic context. *American Economic Review*, **105** (4), 1618–1633.
- and — (2018). Never, ever getting started: On prospect theory without commitment. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2765550, unpublished.
- FISCHBACHER, U., HOFFMANN, G. and SCHUDY, S. (2017). The causal effect of stop-loss and take-gain orders on the disposition effect. *Review of Financial Studies*, **30** (6), 2110–2129.
- GREINER, B. (2015). Subject pool recruitment procedures: organizing experiments with orsee. *Journal of the Economic Science Association*, **1** (1), 114–125.
- GROVE, W. A., JETTER, M. and PAPPS, K. L. (2021). Career lotto? labor supply in a superstar market. *Journal of Economic Behavior & Organization*, **183**, 362–376.
- GUL, F. (1991). A theory of disappointment aversion. *Econometrica*, **59**, 667–686.
- HE, X., HU, S., OBLOJ, J. and ZHOU, X. (2019). Optimal exit time from casino gambling: Strategies of pre-committed and naive gamblers. *SIAM Journal on Control and Optimization*, **57**, 1854–1868.
- HEIMER, R., ILIWA, Z., IMAS, A. and WEBER, M. (2023). Dynamic inconsistency in risky choice: Evidence from the lab and field. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3600583.
- HENDERSON, V., HOBSON, D. and ALEX, S. (2017). Randomized strategies and prospect theory in a dynamic context. *Journal of Economic Theory*, **168**, 287–300.
- , — and — (2018). Probability weighting, stop-loss and the disposition effect. *Journal of Economic Theory*, **178**, 360–397.
- HUANG, Y.-J., NGUYEN-HUU, A. and ZHOU, X. Y. (2020). General stopping behaviors of naïve and noncommitted sophisticated agents, with application to probability distortion. *Mathematical Finance*, **30** (1), 310–340.
- HUBER, J., PAYNE, J. W. and PUTO, C. (1982). Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of Consumer Research*, **9**, 90–98.
- IMAS, A. (2016). The realization effect: Risk-taking after realized versus paper losses. *American Economic Review*, **106** (8), 2086–2109.
- KAHNEMAN, D. and TVERSKY, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, **47** (2), 263–291.

- KARNI, E. and SAFRA, Z. (1990). Behaviorally consistent optimal stopping rules. *Journal of Economic Theory*, **51** (2), 391–402.
- KÓSZEGI, B. and RABIN, M. (2007). Reference-dependent risk attitudes. *American Economic Review*, **97** (4), 1047–1073.
- KOTHİYAL, A., SPINU, V. and WAKKER, P. P. (2011). Prospect theory for continuous distributions: A preference foundation. *Journal of Risk and Uncertainty*, **42** (3), 195.
- LUSARDI, A. and MITCHELL, O. (2011). *Financial Literacy and Planning: Implications for Retirement Wellbeing*. Oxford University Press.
- MACHINA, M. J. (1989). Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature*, **27** (4), 1622–1668.
- NIELSEN, K. (2019). Dynamic risk preferences under realized and paper outcomes. *Journal of Economic Behavior & Organization*, **161**, 68–78.
- ODEAN, T. (1998). Are investors reluctant to realize their losses? *Journal of Finance*, **53** (5), 1775–1798.
- PRIMI, C., MORSANYI, K., CHIESI, F., DONATI, M. A. and HAMILTON, J. (2016). The development and testing of a new version of the cognitive reflection test applying item response theory (irt). *Journal of Behavioral Decision Making*, **29** (5), 453–469.
- QUIGGIN, J. (1982). A theory of anticipated utility. *Journal of Economic Behavior & Organization*, **3**, 323–343.
- RESNICK, I., NEWCOMBE, N. S. and SHIPLEY, T. F. (2017). Dealing with big numbers: Representation and understanding of magnitudes outside of human experience. *Cognitive science*, **41** (4), 1020–1041.
- REVUZ, D. and YOR, M. (1999). *Continuous martingales and Brownian motion*. Springer-Verlag, Berlin.
- ROTHSCHILD, M. and STIGLITZ, J. E. (1970). Increasing risk: I. A definition. *Journal of Economic Theory*, **2** (3), 225–243.
- SCHKADE, D. A. and KAHNEMAN, D. (1998). Does living in California make people happy? A focusing illusion in judgments of life satisfaction. *Psychological Science*, **9** (5), 340–346.
- SHEFRIN, H. and STATMAN, M. (1985). The disposition to sell winners too early and ride losers too long: Theory and evidence. *Journal of Finance*, **40** (3), 777–790.
- SHREVE, S. E. (2004). *Stochastic Calculus for Finance II: Continuous-Time Models*. Springer Science & Business Media.
- STRACK, P. and VIEFERS, P. (2021). Too proud to stop: Regret in dynamic decisions. *Journal of the European Economic Association*, **19** (1), 165–199.

- TVERSKY, A. and KAHNEMAN, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, **5** (4), 297–323.
- WAKKER, P. P. (2010). *Prospect theory: For risk and ambiguity*. Cambridge University Press.
- WEBER, M. and CAMERER, C. F. (1998). The disposition effect in securities trading: An experimental analysis. *Journal of Economic Behavior & Organization*, **33** (2), 167–184.
- XU, Z. and ZHOU, X. (2013). Optimal stopping under probability distortion. *Annals of Applied Probability*, **23** (1), 251–282.

Appendix A: Proofs

A.1: Preliminary Results on Arithmetic Brownian Motions

Fix an initial wealth level $x \in \mathbb{R}_{>0}$ and an expiration date $T \in \mathbb{R}_{>0}$. Throughout this section, we take the perspective of period $t = 0$ and consider a threshold stopping time $\tau_{a,b}$ with $a < x < b$. Our first result describes the distribution $X_{T \wedge \tau_{a,b}}$ conditional on stopping before expiration.

Lemma 1. *If $\mu \neq 0$, then, for any threshold stopping time $\tau_{a,b}$ with $a < x < b$, we have*

$$\mathbb{P}_0[X_{\tau_{a,b}} = a] = \frac{\exp(-(2\mu/\nu^2)b) - \exp(-(2\mu/\nu^2)x)}{\exp(-(2\mu/\nu^2)b) - \exp(-(2\mu/\nu^2)a)}. \quad (6)$$

If $\mu = 0$, then $\mathbb{P}_0[X_{\tau_{a,b}} = a] = \frac{b-x}{b-a}$. In particular, an increase in the drift of the process improves the distribution of $X_{\tau_{a,b}}$ in terms of first-order stochastic dominance.

Proof. Fix some $a, b \in \mathbb{R}_{\geq 0}$ with $a < x < b$. For any threshold stopping time $\tau_{a,b}$, we have

$$\mathbb{P}_0[X_{\tau_{a,b}} = a] = \frac{\Psi(b) - \Psi(x)}{\Psi(b) - \Psi(a)},$$

where $\Psi : \mathbb{R} \rightarrow \mathbb{R}$, $z \mapsto \Psi(z) = \int_0^z \exp(-\int_0^y 2\frac{\mu}{\nu^2} dv) dy = \int_0^z \exp(-2\frac{\mu}{\nu^2}y) dy$ is a strictly increasing *scale function* (e.g., Revuz and Yor, 1999, pp. 302). For any $\mu \neq 0$, we obtain

$$\Psi(z) = \int_0^z \exp\left(-2\frac{\mu}{\nu^2}y\right) dy = \frac{\nu^2}{2\mu} [1 - \exp(-(2\mu/\nu^2)z)],$$

while for $\mu = 0$, we have $\Psi(z) = \int_0^z 1 dy = z$, which yields the claim. The last part of the lemma follows from taking the partial derivative of the right-hand side of Eq. (6) with respect to μ . \square

Our second result derives the probability of reaching the expiration date, and describes several properties of the distribution of $X_{T \wedge \tau_{a,b}}$ conditional on stopping at the expiration date.

Lemma 2. (a) *The probability of stopping at the expiration date equals*

$$\mathbb{P}_0[\tau_{a,b} \geq T | X_0 = x] = \int_a^b q(y, T | X_0 = x) dy,$$

where the integrand is given by

$$q(y, T | X_0 = x) = \frac{2 \exp\left(\frac{\mu(y-x)}{\nu^2} - \frac{T}{2} \frac{\mu^2}{\nu^2}\right)}{(b-a)} \sum_{n=1}^{\infty} \left\{ \sin\left(\frac{\pi n(x-a)}{b-a}\right) \sin\left(\frac{\pi n(y-a)}{b-a}\right) \exp\left(-\frac{T}{2} \frac{n^2 \pi^2 \nu^2}{(b-a)^2}\right) \right\}.$$

(b) $\lim_{a \rightarrow x} \mathbb{P}_0[\tau_{a,b} \geq T | X_0 = x] = 0$ and $\lim_{x \rightarrow b} \mathbb{P}_0[\tau_{a,b} \geq T | X_0 = x] = 0$.

(c) *For any stopping time $\tau_{a,b}$ with $a < x < b$, the CDF of X_T conditional on $\tau_{a,b} \geq T$ equals*

$$\mathbb{P}_0[X_T \leq z | X_0 = x, \tau_{a,b} \geq T] = \frac{\int_a^z \exp\left(\frac{\mu(y-x)}{\nu^2}\right) \sum_{n=1}^{\infty} \left\{ \sin\left(\frac{\pi n(x-a)}{b-a}\right) \sin\left(\frac{\pi n(y-a)}{b-a}\right) \exp\left(-\frac{T}{2} \frac{n^2 \pi^2 \nu^2}{(b-a)^2}\right) \right\} dy}{\int_a^b \exp\left(\frac{\mu(y-x)}{\nu^2}\right) \sum_{n=1}^{\infty} \left\{ \sin\left(\frac{\pi n(x-a)}{b-a}\right) \sin\left(\frac{\pi n(y-a)}{b-a}\right) \exp\left(-\frac{T}{2} \frac{n^2 \pi^2 \nu^2}{(b-a)^2}\right) \right\} dy}.$$

Suppose that $a = x - \epsilon$ and $b = x + \epsilon'$ for some $\epsilon' > \epsilon > 0$.

(d) If $\mu < 0$, then $\frac{\partial}{\partial \mu} \mathbb{P}_0[X_T \leq z | X_0 = x, \tau_{a,b} \geq T] < 0$ for any $z \in [a, b)$.

(e) If $\mu = 0$, then $\lim_{\epsilon \rightarrow 0} \mathbb{P}_0[X_T \leq x | X_0 = x, \tau_{a,b} \geq T] = 0$.

Suppose that $a = x - \epsilon - \epsilon'$ and $b = x + \epsilon$ for some $\epsilon > 0$ and $\epsilon' \geq 0$. In addition, let $\alpha \in (0, \epsilon)$.

(f) If $\mu \leq 0$, then $\mathbb{P}_0[X_T \leq x - \alpha | X_0 = x, \tau_{a,b} \geq T] \geq \mathbb{P}_0[X_T > x + \alpha | X_0 = x, \tau_{a,b} \geq T]$, holding with a strict inequality whenever $\mu < 0$.

Proof. PART (a). Example 5.1 in Cox and Miller (1977).

PART (b). We prove only the first part here, as the proof of the second part is analogous. To establish the first part, it is sufficient to show that

$$\lim_{a \rightarrow x} \sum_{n=1}^{\infty} \left\{ \sin \left(\frac{\pi n(x-a)}{b-a} \right) \sin \left(\frac{\pi n(y-a)}{b-a} \right) \exp \left(-\frac{T}{2} \frac{n^2 \pi^2 \nu^2}{(b-a)^2} \right) \right\} = 0. \quad (7)$$

As $|\sin \left(\frac{\pi n(x-a)}{b-a} \right) \sin \left(\frac{\pi n(y-a)}{b-a} \right)| \leq 1$ and as $\exp \left(-\frac{T}{2} \frac{n^2 \pi^2 \nu^2}{(b-a)^2} \right) \leq \exp \left(-\frac{T}{2} \frac{n \pi \nu^2}{(b-a)^2} \right)$ and as

$$\sum_{n=1}^{\infty} \exp \left(-\frac{t}{2} \frac{n \pi \nu^2}{(b-a)^2} \right) = \frac{1}{\left(\exp \left(\frac{t}{2} \frac{\pi \nu^2}{(b-a)^2} \right) - 1 \right)} < \infty,$$

we can take the limit in (7) inside the summation. The claim follows from the fact that $\sin(0) = 0$.

PART (c). Follows immediately from Part (a).

PART (d). Consider a threshold stopping time $\tau_{a,b}$ with $a = x - \epsilon$ and $b = x + \epsilon'$ for some $\epsilon' > \epsilon > 0$. The CDF of the corresponding random variable $X_{\tau_{a,b} \wedge T}$ is given by

$$\mathbb{P}_0[X_{T \wedge \tau_{a,b}} \leq z] = \begin{cases} \mathbb{P}_0[\tau_{a,b} < T] \cdot \mathbb{P}_0[X_{\tau_{a,b}} = a] & \text{if } z = a, \\ \mathbb{P}_0[\tau_{a,b} < T] \cdot \mathbb{P}_0[X_{\tau_{a,b}} = a] + \int_a^z q(y, T | X_0 = x) dy & \text{if } a < z < b, \\ 1 & \text{if } z = b. \end{cases}$$

Taking the partial derivative of the CDF at $z \in [a, b)$ with respect to the drift of the process yields

$$\begin{aligned} \frac{\partial}{\partial \mu} \mathbb{P}_0[X_{T \wedge \tau_{a,b}} \leq z] &= \mathbb{P}_0[\tau_{a,b} < T] \cdot \frac{\partial}{\partial \mu} \mathbb{P}_0[X_{\tau_{a,b}} = a] + \mathbb{P}_0[X_{\tau_{a,b}} = a] \cdot \frac{\partial}{\partial \mu} \mathbb{P}_0[\tau_{a,b} < T] \\ &\quad + \int_a^z \frac{\partial}{\partial \mu} q(y, T | X_0 = x) dy. \end{aligned}$$

The first two terms in $\frac{\partial}{\partial \mu} \mathbb{P}_0[X_{T \wedge \tau_{a,b}} \leq z]$ are negative or, at least, non-positive: First, by Lemma 1, $\frac{\partial}{\partial \mu} \mathbb{P}_0[X_{\tau_{a,b}} = a] < 0$. Second, because (i) the drift of the process is negative and (ii) $x - a < b - x$ by assumption, a marginal increase in the drift shifts the distribution of $\tau_{a,b}$ upward, so also $\frac{\partial}{\partial \mu} \mathbb{P}_0[\tau_{a,b} < T] \leq 0$. Together these two observations imply that $\frac{\partial}{\partial \mu} \mathbb{P}_0[X_{T \wedge \tau_{a,b}} \leq a] < 0$.

The remainder of the proof proceeds in two steps: First, we will show that *there exists some $\hat{z} \in (a, b)$ such that $\frac{\partial}{\partial \mu} \mathbb{P}_0[X_{T \wedge \tau_{a,b}} \leq z] < 0$ if and only if $z < \hat{z}$* . Second, we will show that $\frac{\partial}{\partial \mu} \mathbb{P}_0[X_{T \wedge \tau_{a,b}} = b] > 0$ and, thus, $\frac{\partial}{\partial \mu} \{1 - \mathbb{P}_0[X_{T \wedge \tau_{a,b}} = b]\} < 0$, which in turn implies that $\hat{z} = b$.

1. Step: Using the definition of $q(y, T|X_0 = x)$, as provided in Part (a), we obtain

$$\begin{aligned} \int_a^z \frac{\partial}{\partial \mu} q(y, T|X_0 = x) dy &= \frac{1}{\nu^2} \left[\int_a^z yq(y, T|X_0 = x) dy - (x + T\mu) \int_a^z q(y, T|X_0 = x) dy \right] \\ &= \frac{\mathbb{P}_0[\tau_{a,b} \geq T, X_T \leq z]}{\nu^2} \left[\mathbb{E}_0[X_T | \tau_{a,b} \geq T, X_T \leq z] - \mathbb{E}_0[X_T] \right]. \end{aligned}$$

We distinguish two cases, depending on whether the difference in brackets is positive or not. If $\mathbb{E}_0[X_T | \tau_{a,b} \geq T, X_T \leq z] \leq \mathbb{E}_0[X_T]$, then $\int_a^z \frac{\partial}{\partial \mu} q(y, T|X_0 = x) dy \leq 0$, which together with our observations on the first two terms of $\frac{\partial}{\partial \mu} \mathbb{P}_0[X_{T \wedge \tau_{a,b}} \leq z]$ implies that $\frac{\partial}{\partial \mu} \mathbb{P}_0[X_{T \wedge \tau_{a,b}} \leq z] < 0$. Otherwise, because $\mathbb{E}_0[X_T | \tau_{a,b} \geq T, X_T \leq z]$ and $\mathbb{P}_0[\tau_{a,b} \geq T, X_T \leq z]$ are increasing in z , it follows that $\int_a^z \frac{\partial}{\partial \mu} q(y, T|X_0 = x) dy$ increases in z whenever it is positive. In sum, there exists some $\hat{z} \in (a, b)$ such that $\frac{\partial}{\partial \mu} \mathbb{P}_0[X_{T \wedge \tau_{a,b}} \leq z] < 0$ if and only if $z < \hat{z}$, which was to be proven.

2. Step: Now, it is sufficient to show that $\frac{\partial}{\partial \mu} \mathbb{P}_0[X_{T \wedge \tau_{a,b}} = b] > 0$. For the sake of a contradiction, we assume that $\mathbb{P}_0[X_{T \wedge \tau_{a,b}} = b | \mu = \mu'] \geq \mathbb{P}_0[X_{T \wedge \tau_{a,b}} = b | \mu = \mu'']$ for some $\mu' < \mu'' < 0$.

Denote as $M_t := \max_{0 \leq s \leq t} X_s$ the running maximum of the process. By our assumption toward a contradiction, it has to be true that, for some $t \in (0, T)$,

$$\mathbb{P}_0[M_t \geq b | X_s > a, \forall s \leq t, \mu = \mu'] \geq \mathbb{P}_0[M_t \geq b | X_s > a, \forall s \leq t, \mu = \mu''];$$

otherwise, reaching the upper threshold b would be strictly more likely under the process with a less negative drift. Recall that, because both $\mathbb{P}_0[X_{\tau_{a,b}} = a]$ and $\mathbb{P}_0[\tau_{a,b} < T]$ (weakly) decrease with μ , increasing the drift of the process makes it less likely to reach the lower threshold a . Hence, a necessary condition for our assumption toward a contradiction to hold is

$$\mathbb{P}_0[M_t \geq b | \mu = \mu'] \geq \mathbb{P}_0[M_t \geq b | \mu = \mu''];$$

that is, at any time $t \in (0, T)$, conditioning on not having reached the lower threshold is more restrictive for the process with a more negative drift. By the same argument, we can ignore the fact that the process is absorbing in zero. In this case, by the Reflection Principle and Girsanov's Theorem, for an arbitrary drift $\mu < 0$, the distribution of the running maximum, M_t , satisfies

$$\mathbb{P}_0[M_t \geq b] = \exp\left(\frac{2\epsilon'\mu}{\nu^2}\right) \left[1 - \Phi\left(\frac{\epsilon' + \mu t}{\sqrt{t\nu}}\right) \right] + \left[1 - \Phi\left(\frac{\epsilon' - \mu t}{\sqrt{t\nu}}\right) \right],$$

with $b = x + \epsilon'$ for some $\epsilon' > 0$ and with Φ being the standard normal CDF (see Example 5.1 in Cox and Miller, 1977, or Corollary 7.2.2 in Shreve, 2004). Hence, we obtain

$$\begin{aligned} \frac{\partial}{\partial \mu} \mathbb{P}_0[M_t \geq b] &= \frac{2\epsilon'}{\nu^2} \exp\left(\frac{2\epsilon'\mu}{\nu^2}\right) \left[1 - \Phi\left(\frac{\epsilon' + \mu t}{\sqrt{t\nu}}\right) \right] - \exp\left(\frac{2\epsilon'\mu}{\nu^2}\right) \phi\left(\frac{\epsilon' + \mu t}{\sqrt{t\nu}}\right) + \phi\left(\frac{\epsilon' - \mu t}{\sqrt{t\nu}}\right) \\ &= \frac{2\epsilon'}{\nu^2} \exp\left(\frac{2\epsilon'\mu}{\nu^2}\right) \left[1 - \Phi\left(\frac{\epsilon' + \mu t}{\sqrt{t\nu}}\right) \right] + \frac{\sqrt{t}}{\sqrt{2\pi\nu}} \left[\exp\left(\frac{(\epsilon' - \mu t)^2}{2t\nu^2}\right) - \exp\left(\frac{(\epsilon' - \mu t)^2}{2t\nu^2}\right) \right] \\ &= \frac{2\epsilon'}{\nu^2} \exp\left(\frac{2\epsilon'\mu}{\nu^2}\right) \left[1 - \Phi\left(\frac{\epsilon' + \mu t}{\sqrt{t\nu}}\right) \right] \\ &> 0, \end{aligned}$$

where the second equality follows from plugging in the density ϕ of a standard normal distribution. But this implies that $\mathbb{P}_0[M_t \geq b|\mu = \mu'] \geq \mathbb{P}_0[M_t \geq b|\mu = \mu'']$ cannot hold; a contradiction.

PART (e). By Part (b), we have $\lim_{\epsilon \rightarrow 0} \int_{x-\epsilon}^{x+\epsilon'} q(y, T|X_0 = x) dy = 0$ and, as a consequence, also $\lim_{\epsilon \rightarrow 0} \int_{x-\epsilon}^x q(y, T|X_0 = x) dy = 0$. Now, to determine $\lim_{\epsilon \rightarrow 0} \mathbb{P}_0[X_T \leq x|X_0 = x, \tau_{a,b} \geq T]$, we will apply L'Hospital's rule. For that, we have to make a few preliminary observations.

First, if the partial derivative $\frac{\partial}{\partial \epsilon} q(y, T|X_0 = x)$ exists, then it is given by

$$\begin{aligned} \frac{\partial}{\partial \epsilon} q(y, T|X_0 = x) &= \frac{\epsilon'}{(\epsilon + \epsilon')^2} \sum_{n=1}^{\infty} \left\{ \pi n \cos\left(\frac{\pi n \epsilon}{\epsilon + \epsilon'}\right) \sin\left(\frac{\pi n(y-x+\epsilon)}{\epsilon + \epsilon'}\right) \exp\left(-\frac{T}{2} \frac{n^2 \pi^2 \nu^2}{(\epsilon + \epsilon')^2}\right) \right\} \\ &+ \frac{\epsilon' - (y-x)}{(\epsilon + \epsilon')^2} \sum_{n=1}^{\infty} \left\{ \pi n \sin\left(\frac{\pi n \epsilon}{\epsilon + \epsilon'}\right) \cos\left(\frac{\pi n(y-x+\epsilon)}{\epsilon + \epsilon'}\right) \exp\left(-\frac{T}{2} \frac{n^2 \pi^2 \nu^2}{(\epsilon + \epsilon')^2}\right) \right\} \\ &+ \frac{T \nu^2}{(\epsilon + \epsilon')^3} \sum_{n=1}^{\infty} \left\{ \pi^2 n^2 \sin\left(\frac{\pi n \epsilon}{\epsilon + \epsilon'}\right) \sin\left(\frac{\pi n(y-x+\epsilon)}{\epsilon + \epsilon'}\right) \exp\left(-\frac{T}{2} \frac{n^2 \pi^2 \nu^2}{(\epsilon + \epsilon')^2}\right) \right\}, \end{aligned}$$

and since we have

$$\begin{aligned} \left| \frac{\partial}{\partial \epsilon} q(y, T|X_0 = x) \right| &\leq \left[\frac{1}{\epsilon'} + \frac{\epsilon' - 2(y-x)}{\epsilon'^2} \right] \frac{\exp\left(\frac{T}{2} \frac{\pi \nu^2}{\epsilon'^2}\right) \pi}{\left(\exp\left(\frac{T}{2} \frac{\pi \nu^2}{4\epsilon'^2}\right) - 1\right)^2} \\ &+ \frac{T \nu^2 \exp\left(\frac{T}{2} \frac{\pi \nu^2}{\epsilon'^2}\right) \left(\exp\left(\frac{T}{2} \frac{\pi \nu^2}{\epsilon'^2}\right) + 1\right) \pi^2}{\epsilon'^3 \left(\exp\left(\frac{T}{2} \frac{\pi \nu^2}{4\epsilon'^2}\right) - 1\right)^3} < \infty, \end{aligned} \quad (8)$$

it indeed exists. To apply L'Hospital's rule, we need to compute the limit for ϵ approaching zero:

$$\begin{aligned} \lim_{\epsilon \rightarrow 0} \frac{\partial}{\partial \epsilon} q(y, T|X_0 = x) &= \frac{1}{\epsilon'} \sum_{n=1}^{\infty} \lim_{\epsilon \rightarrow 0} \left\{ \pi n \cos\left(\frac{\pi n \epsilon}{\epsilon + \epsilon'}\right) \sin\left(\frac{\pi n(y-x+\epsilon)}{\epsilon + \epsilon'}\right) \exp\left(-\frac{T}{2} \frac{n^2 \pi^2 \nu^2}{(\epsilon + \epsilon')^2}\right) \right\} \\ &+ \frac{\epsilon' - (y-x)}{\epsilon'^2} \sum_{n=1}^{\infty} \lim_{\epsilon \rightarrow 0} \left\{ \pi n \sin\left(\frac{\pi n \epsilon}{\epsilon + \epsilon'}\right) \cos\left(\frac{\pi n(y-x+\epsilon)}{\epsilon + \epsilon'}\right) \exp\left(-\frac{T}{2} \frac{n^2 \pi^2 \nu^2}{(\epsilon + \epsilon')^2}\right) \right\} \\ &+ \frac{T \nu^2}{\epsilon'^3} \sum_{n=1}^{\infty} \lim_{\epsilon \rightarrow 0} \left\{ \pi^2 n^2 \sin\left(\frac{\pi n \epsilon}{\epsilon + \epsilon'}\right) \sin\left(\frac{\pi n(y-x+\epsilon)}{\epsilon + \epsilon'}\right) \exp\left(-\frac{T}{2} \frac{n^2 \pi^2 \nu^2}{(\epsilon + \epsilon')^2}\right) \right\} \\ &= \frac{1}{\epsilon'} \sum_{n=1}^{\infty} \left\{ \pi n \sin\left(\frac{\pi n(y-x)}{\epsilon'}\right) \exp\left(-\frac{T}{2} \frac{n^2 \pi^2 \nu^2}{\epsilon'^2}\right) \right\} \geq 0, \end{aligned} \quad (9)$$

where the first equality follows from the fact that, by the considerations in Eq. (8), we can take the limits into the summations; the second equality holds as $\sin(0) = 0$ and $\cos(0) = 1$; and the inequality follows from the fact that $q(y, T|X_0 = x) \geq 0$ and $\lim_{\epsilon \rightarrow 0} q(y, T|X_0 = x) = 0$, because otherwise $q(y, T|X_0 = x)$ would be negative for ϵ sufficiently close to zero.

Second, we observe that

$$\begin{aligned} \left| \int_{x-\epsilon}^x \frac{\partial}{\partial \epsilon} q(y, T|X_0 = x) dy \right| &\leq \int_{x-\epsilon}^x \left| \frac{\partial}{\partial \epsilon} q(y, T|X_0 = x) \right| dy \\ &\leq \epsilon \left[\frac{1}{\epsilon'} + \frac{\epsilon' - 2(y-x)}{\epsilon'^2} \right] \frac{\exp\left(\frac{T}{2} \frac{\pi \nu^2}{\epsilon'^2}\right) \pi}{\left(\exp\left(\frac{T}{2} \frac{\pi \nu^2}{4\epsilon'^2}\right) - 1\right)^2} + \epsilon \frac{T \nu^2 \exp\left(\frac{T}{2} \frac{\pi \nu^2}{\epsilon'^2}\right) \left(\exp\left(\frac{T}{2} \frac{\pi \nu^2}{\epsilon'^2}\right) + 1\right) \pi^2}{\epsilon'^3 \left(\exp\left(\frac{T}{2} \frac{\pi \nu^2}{4\epsilon'^2}\right) - 1\right)^3} \xrightarrow{\epsilon \rightarrow 0} 0, \end{aligned} \quad (10)$$

where the first inequality follows by the triangle inequality, and where the second inequality follows from Eq. (8). Taking the limit of the final expression is straightforward.

Third, we conclude that

$$\begin{aligned}
\lim_{\epsilon \rightarrow 0} \int_x^{x+\epsilon'} \frac{\partial}{\partial \epsilon} q(y, T | X_0 = x) dy &= \int_x^{x+\epsilon'} \lim_{\epsilon \rightarrow 0} \frac{\partial}{\partial \epsilon} q(y, T | X_0 = x) dy \\
&= \frac{1}{\epsilon'} \int_x^{x+\epsilon'} \sum_{n=1}^{\infty} \left\{ \pi n \sin \left(\frac{\pi n (y-x)}{\epsilon'} \right) \exp \left(-\frac{T n^2 \pi^2 \nu^2}{2 \epsilon'^2} \right) \right\} dy \\
&= \frac{1}{\epsilon'} \int_0^1 \sum_{n=1}^{\infty} \left\{ \sin(\pi n z) \pi n \exp \left(-\frac{T n^2 \pi^2 \nu^2}{2 \epsilon'^2} \right) \right\} dz,
\end{aligned}$$

where the first equality holds by the Theorem of Dominated Convergence, the second one holds by the second to last line in (9), and the third equality follows by substitution. Recall that

$$\sum_{n=1}^{\infty} \left\{ \sin(\pi n z) \pi n \exp \left(-\frac{T n^2 \pi^2 \nu^2}{2 \epsilon'^2} \right) \right\} \geq 0 \quad (11)$$

for any $z \in (0, 1)$, and notice that this inequality is strict for any $z = \frac{1}{k}$ with $k \in \mathbb{N}_{\geq 2}$. The latter follows from the fact that $\sin(\pi n \frac{i}{k}) = -\sin(\pi n \frac{k+i}{k})$ for any $i \leq k$, and $\sin(\pi \frac{i}{k}) \geq 0$ for any $i \leq k$ with a strict inequality for any $i \notin \{0, k\}$, and $\pi n \exp \left(-\frac{T n^2 \pi^2 \nu^2}{2 \epsilon'^2} \right)$ being strictly decreasing in n . Since (11) is continuous in z , we conclude that it is strictly positive on a dense interval around any $z = \frac{1}{k}$ with $k \in \mathbb{N}_{\geq 2}$. This, in turn, implies that $\lim_{\epsilon \rightarrow 0} \int_x^{x+\epsilon'} \frac{\partial}{\partial \epsilon} q(y, T | X_0 = x) dy > 0$.

Combining all the considerations above, we finally conclude that

$$\begin{aligned}
\lim_{\epsilon \rightarrow 0} \mathbb{P}_0[X_T \leq x | X_0 = x, \tau_{a,b} \geq T] &= \lim_{\epsilon \rightarrow 0} \frac{\frac{\partial}{\partial \epsilon} \int_{x-\epsilon}^x q(y, T | X_0 = x) dy}{\frac{\partial}{\partial \epsilon} \int_{x-\epsilon}^{x+\epsilon'} q(y, T | X_0 = x) dy} \\
&= \lim_{\epsilon \rightarrow 0} \frac{\int_{x-\epsilon}^x \frac{\partial}{\partial \epsilon} q(y, T | X_0 = x) dy}{\int_{x-\epsilon}^{x+\epsilon'} \frac{\partial}{\partial \epsilon} q(y, T | X_0 = x) dy} \\
&= \frac{\lim_{\epsilon \rightarrow 0} \int_{x-\epsilon}^x \frac{\partial}{\partial \epsilon} q(y, T | X_0 = x) dy}{\lim_{\epsilon \rightarrow 0} \int_{x-\epsilon}^{x+\epsilon'} \frac{\partial}{\partial \epsilon} q(y, T | X_0 = x) dy} \\
&= 0,
\end{aligned}$$

where the first equality follows by L'Hospital's rule (given that the limit on the right-hand side exists), the second equality follows by the Theorem of Dominated Convergence and by Leibniz's integral rule, the third equality follows from the fact that the limit of the numerator and the limit of the denominator exist, and the last equality holds by (10) and by the fact that $\lim_{\epsilon \rightarrow 0} \int_x^{x+\epsilon'} \frac{\partial}{\partial \epsilon} q(y, T | X_0 = x) dy > 0$ and therefore, by $\lim_{\epsilon \rightarrow 0} \frac{\partial}{\partial \epsilon} q(y, T | X_0 = x) \geq 0$, also $\lim_{\epsilon \rightarrow 0} \int_{x-\epsilon}^{x+\epsilon'} \frac{\partial}{\partial \epsilon} q(y, T | X_0 = x) dy > 0$. This completes the proof.

PART (f). To begin with, let $\epsilon' = 0$. By Part (c), we have to show that

$$\begin{aligned}
&\int_a^{x-\alpha} \exp \left(\frac{\mu(y-x)}{\nu^2} \right) \sum_{n=1}^{\infty} \left\{ \sin \left(\frac{\pi n (x-a)}{b-a} \right) \sin \left(\frac{\pi n (y-a)}{b-a} \right) \exp \left(-\frac{T n^2 \pi^2 \nu^2}{2 (b-a)^2} \right) \right\} dy \\
&\geq \int_{x+\alpha}^b \exp \left(\frac{\mu(y-x)}{\nu^2} \right) \sum_{n=1}^{\infty} \left\{ \sin \left(\frac{\pi n (x-a)}{b-a} \right) \sin \left(\frac{\pi n (y-a)}{b-a} \right) \exp \left(-\frac{T n^2 \pi^2 \nu^2}{2 (b-a)^2} \right) \right\} dy
\end{aligned}$$

for any $\alpha \in (0, \epsilon)$, with a strict inequality if $\mu < 0$. For any $\mu \leq 0$, we have $\exp\left(\frac{\mu(y-x)}{\nu^2}\right) \geq 1$ if and only if $y \leq x$, holding with a strict inequality whenever $y < x$ and $\mu < 0$. This implies that

$$\begin{aligned} & \int_a^{x-\alpha} \exp\left(\frac{\mu(y-x)}{\nu^2}\right) \sum_{n=1}^{\infty} \left\{ \sin\left(\frac{\pi n(x-a)}{b-a}\right) \sin\left(\frac{\pi n(y-a)}{b-a}\right) \exp\left(-\frac{T}{2} \frac{n^2 \pi^2 \nu^2}{(b-a)^2}\right) \right\} dy \\ & \geq \int_{\frac{\pi n}{2}}^{\frac{\pi n}{2} + \frac{\alpha}{\epsilon}} \sum_{n \in \mathbb{N}, n \text{ odd}} \sin\left(\frac{\pi n}{2}\right) \sin\left(\frac{\pi n}{2} - z\right) \exp\left(-\frac{T}{2} \frac{n^2 \pi^2 \nu^2}{4\epsilon^2}\right) dz \\ & = \int_{\frac{\pi n}{2} - \frac{\alpha}{\epsilon}}^{\frac{\pi n}{2}} \sum_{n \in \mathbb{N}, n \text{ odd}} \sin\left(\frac{\pi n}{2}\right) \sin\left(\frac{\pi n}{2} + z\right) \exp\left(-\frac{T}{2} \frac{n^2 \pi^2 \nu^2}{4\epsilon^2}\right) dz \\ & \geq \int_{x+\alpha}^b \exp\left(\frac{\mu(y-x)}{\nu^2}\right) \sum_{n=1}^{\infty} \left\{ \sin\left(\frac{\pi n(x-a)}{b-a}\right) \sin\left(\frac{\pi n(y-a)}{b-a}\right) \exp\left(-\frac{T}{2} \frac{n^2 \pi^2 \nu^2}{(b-a)^2}\right) \right\} dy, \end{aligned}$$

where the two inequalities follow from the fact that $\frac{x-a}{b-a} = \frac{1}{2}$ and $\sin\left(\frac{\pi n}{2}\right) = 0$ for any even $n \in \mathbb{N}$, while the equality holds since $\sin\left(\frac{\pi n}{2} - z\right) = \sin\left(\frac{\pi n}{2} + z\right)$ for any odd $n \in \mathbb{N}$ and any $z \in (0, \frac{\pi n}{2})$. The claim follows from the fact that the inequalities are strict whenever $\mu < 0$.

Fix some $\epsilon > 0$ and $\mu \leq 0$. Now, if $\epsilon' > 0$, the probability that X_T is weakly below x , conditional on reaching the expiration date when playing according to the stopping time $\tau_{a,b}$, $\mathbb{P}_0[X_T \leq x | X_0 = x, \tau_{a,b} \geq T]$, increases compared to the case with $\epsilon' = 0$. This follows from the fact that due to $\epsilon' > 0$ there is now more room below x than above x to reach the expiration date T and from the continuity of the sample paths. \square

A.2: Motivating Example

Proof of Proposition 1. We have to show that, when the value function is linear, then for any point in time $t < T$ with $X_t = x_t$ there exists a stopping time $\tau_{a,b}$ such that $U^s(X_{T \wedge \tau_{a,b}} | \mathcal{C}) > x_t$ or, equivalently,

$$\begin{aligned} & \mathbb{P}_0[\tau_{a,b} < T - t] \cdot \underbrace{[p(a - x_t)\sigma(a, x_t) + (1 - p)(b - x_t)\sigma(b, x_t)]}_{(\star)} \\ & \quad + \mathbb{P}_0[\tau_{a,b} \geq T - t] \cdot \underbrace{\int_{(a,b)} (z - x_t)\sigma(z, x_t) d\Phi_\mu(z)}_{(\star\star)} > 0, \end{aligned}$$

where the probability $p = p(a, b, \mu)$ is defined in Eq. (6) and where $\mathbb{P}_0[\tau_{a,b} < T - t]$ as well as the conditional CDF $\Phi_\mu(z) := \mathbb{P}_0[X_T \leq z | X_0 = x_0, \tau_{a,b} \geq T - t]$ are described in Lemma 2.

Consider a threshold stopping time $\tau_{a,b}$ with $a = x_t - \epsilon$ and $b = x_t + \epsilon'$ for some $\epsilon' > \epsilon > 0$. First, we show that there exists some threshold $\hat{\epsilon} > 0$ such that for any $\epsilon < \hat{\epsilon}$, it holds that $(\star) > 0$. Since $p = \frac{b-x_t}{b-a} = \frac{\epsilon'}{\epsilon + \epsilon'}$, it follows that $(\star) > 0$ holds if and only if

$$\frac{\epsilon \epsilon'}{\epsilon + \epsilon'} [\sigma(x_t + \epsilon', x_t) - \sigma(x_t - \epsilon, x_t)] > 0.$$

The claim then follows from the fact that — due to ordering — the salience weight $\sigma(x_t - \epsilon, x_t)$ monotonically increases in ϵ , and $\sigma(x_t + \epsilon', x_t) > \sigma(x_t, x_t)$ holds.

Second, we show that *there exists some $\check{\epsilon} > 0$, such that for any $\epsilon < \check{\epsilon}$, $(\star\star) > 0$* . We have

$$\begin{aligned} \int_{(a,b)} (z - x_t) \sigma(z, x_t) d\Phi_\mu(z) &\geq \int_{(x-\epsilon, x)} (z - x_t) \sigma(z, x_t) d\Phi_\mu(z) + \int_{(x+\epsilon, x+\epsilon')} (z - x_t) \sigma(z, x_t) d\Phi_\mu(z) \\ &> -\epsilon \bar{\sigma} \int_{(x-\epsilon, x)} d\Phi_\mu(z) + \epsilon \underline{\sigma} \int_{(x+\epsilon, x+\epsilon')} d\Phi_\mu(z) \\ &= \epsilon [(1 - \Phi_\mu(x + \epsilon)) \underline{\sigma} - \Phi_\mu(x) \bar{\sigma}], \end{aligned}$$

where the weak inequality holds as $\epsilon > 0$ and the strict inequality follows by the definition of $\bar{\sigma} := \sup_{(x,y) \in \mathbb{R}_{\geq 0}^2} \sigma(x, y)$ and $\underline{\sigma} := \inf_{(x,y) \in \mathbb{R}_{\geq 0}^2} \sigma(x, y)$. Now recall that $\bar{\sigma} < \infty$ and $\underline{\sigma} > 0$ by assumption. By Lemma 2 (e), we have $\lim_{\epsilon \rightarrow 0} \Phi_\mu(x) = 0$, which yields the claim. \square

A.3: Main Theoretical Result

The proof of our main theoretical result builds on the following lemma that we prove first.

Lemma 3. *The auxiliary utility function $\tilde{u} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ is of exponential growth at $z = x$.*

Proof of Lemma 3. We have to find some $\alpha, \beta \in \mathbb{R}_{> 0}$, so that $[\tilde{u}(z) + \beta] \leq [\tilde{u}(x) + \beta] \exp(\alpha(z - x))$ holds for any $z \geq 0$. Set $\beta = \frac{\sigma(v(x), v(x))v'(x)}{\alpha}$ for some $\alpha > 0$. We need to find some $\alpha > 0$ so that

$$(v(z) - v(x))\sigma(v(z), v(x)) + \beta \leq [(v(x) - v(x))\sigma(v(x), v(x)) + \beta] \exp(\alpha(z - x))$$

holds for all $z \geq 0$. This condition is indeed satisfied if and only if

$$\alpha \frac{v(z) - v(x)}{v'(x)} \frac{\sigma(v(z), v(x))}{\sigma(v(x), v(x))} + 1 \leq \exp(\alpha(z - x))$$

or, equivalently,

$$\frac{v(z) - v(x)}{v'(x)} \frac{\sigma(v(z), v(x))}{\sigma(v(x), v(x))} \leq \frac{\exp(\alpha(z - x)) - 1}{\alpha} \quad (12)$$

holds for all $z \geq 0$. By construction, (12) holds at $z = x$. We distinguish two cases:

1. CASE: Let $z < x$. Divide both sides of (12) by $z - x < 0$, which gives

$$\frac{\frac{v(x) - v(z)}{x - z}}{v'(x)} \frac{\sigma(v(z), v(x))}{\sigma(v(x), v(x))} \geq \frac{\exp(\alpha(x - x)) - \exp(\alpha(z - x))}{x - z} \frac{1}{\alpha}. \quad (13)$$

Since the exponential function is strictly convex, such that, for any $z < x$, we have

$$\frac{\exp(\alpha(x - x)) - \exp(\alpha(z - x))}{x - z} < \alpha \exp(\alpha(x - x)) = \alpha,$$

the right-hand side of (13) is strictly less than 1. Since the value function is (weakly) concave, which implies that, for any $z < x$, we have

$$\frac{v(x) - v(z)}{x - z} \geq v'(x),$$

and since $\sigma(v(z), v(x)) > \sigma(v(x), v(x))$ holds by ordering, the left-hand side of (13) is strictly larger than 1. In sum, we conclude that, for any $z < x$, Condition (12) is satisfied for any $\alpha > 0$.

2. CASE: Let $z > x$. Since both sides of (12) are zero at $z = x$, we can re-write (12) as follows

$$\int_x^z \frac{\partial}{\partial w} \left[\frac{\exp(\alpha(w-x)) - 1}{\alpha} \right] - \frac{\partial}{\partial w} \left[\frac{v(w) - v(x)}{v'(x)} \frac{\sigma(v(w), v(x))}{\sigma(v(x), v(x))} \right] dw \geq 0,$$

which holds if and only if

$$\int_x^z \exp(\alpha(w-x)) - \left[\frac{v'(w)}{v'(x)} \frac{\sigma(v(w), v(x))}{\sigma(v(x), v(x))} + [v(w) - v(x)] \frac{v'(w)}{v'(x)} \frac{\frac{\partial}{\partial v(w)} \sigma(v(w), v(x))}{\sigma(v(x), v(x))} \right] dw \geq 0. \quad (14)$$

A sufficient condition for (14) to hold is that

$$\exp(\alpha(w-x)) \geq \frac{v'(w)}{v'(x)} \frac{\sigma(v(w), v(x))}{\sigma(v(x), v(x))} + [v(w) - v(x)] \frac{v'(w)}{v'(x)} \frac{\frac{\partial}{\partial v(w)} \sigma(v(w), v(x))}{\sigma(v(x), v(x))}$$

for any $w \geq x$. When evaluated at $w = x$, this inequality is tight, since the salience function is differentiable and thus $\lim_{w \rightarrow x} [v(w) - v(x)] \frac{\partial}{\partial v(w)} \sigma(v(w), v(x)) = 0$. Also, if the right-hand side of this inequality is non-positive, the condition is certainly met. So, from now on, consider only $w > x$ for which the right-hand side is positive. Then, we can re-state the condition as follows

$$\alpha \geq \frac{\ln \left(\frac{v'(w)}{v'(x)} \frac{\sigma(v(w), v(x))}{\sigma(v(x), v(x))} + [v(w) - v(x)] \frac{v'(w)}{v'(x)} \frac{\frac{\partial}{\partial v(w)} \sigma(v(w), v(x))}{\sigma(v(x), v(x))} \right)}{w - x}, \quad (15)$$

which has to hold for all relevant $w > x$. The right-hand side of (15) is bounded from above by

$$h(w) := \frac{\ln \left(\frac{\sigma(v(w), v(x)) + [v(w) - v(x)] \frac{\partial}{\partial v(w)} \sigma(v(w), v(x))}{\sigma(v(x), v(x))} \right)}{w - x},$$

since v is (weakly) concave. Hence, a sufficient condition for (15) to hold is given by $\alpha \geq \max_{w \in (x, \infty)} h(w)$. Since we are free to choose any $\alpha > 0$, it is thus sufficient to show that $\max_{w \in (x, \infty)} h(w) < \infty$. First, since $h(w) \geq 0$ for any $w \geq x$, we know that, if the limit $\lim_{w \rightarrow \infty} h(w)$ does not exist, then it has to be positive infinity. Then, by L'Hospital's rule, we conclude that

$$\begin{aligned} 0 \leq \lim_{w \rightarrow x} h(w) &= \lim_{w \rightarrow x} \frac{\frac{2v'(w)}{\partial v(w)} \frac{\partial}{\partial v(w)} \sigma(v(w), v(x)) + v'(w) [v(w) - v(x)] \frac{\partial^2}{\partial v(w)^2} \sigma(v(w), v(x))}{\sigma(v(x), v(x))} \\ &= \frac{\lim_{w \rightarrow x} 2v'(w) \frac{\partial}{\partial v(w)} \sigma(v(w), v(x)) + v'(w) [v(w) - v(x)] \frac{\partial^2}{\partial v(w)^2} \sigma(v(w), v(x))}{\lim_{w \rightarrow x} \sigma(v(w), v(x)) + [v(w) - v(x)] \frac{\partial}{\partial v(w)} \sigma(v(w), v(x))} \\ &= \frac{2v'(x)}{\sigma(v(x), v(x))} \frac{\partial}{\partial v(w)} \sigma(v(w), v(x)) \Big|_{w=x} < \infty. \end{aligned}$$

where the third equality as well as the inequality follow from the fact that the salience function is twice differentiable, which implies, in particular, $\lim_{w \rightarrow x} [v(w) - v(x)] \frac{\partial}{\partial v(w)} \sigma(v(w), v(x)) = 0$ as well as $\lim_{w \rightarrow x} [v(w) - v(x)] \frac{\partial^2}{\partial v(w)^2} \sigma(v(w), v(x)) = 0$. The fact that the limit exists further justifies the application of L'Hospital's rule. Second, again by L'Hospital's rule, we have

$$\begin{aligned}
0 \leq \lim_{w \rightarrow \infty} h(w) &= \lim_{w \rightarrow \infty} \frac{2v'(w) \frac{\partial}{\partial v(w)} \sigma(v(w), v(x)) + v'(w) [v(w) - v(x)] \frac{\partial^2}{\partial v(w)^2} \sigma(v(w), v(x))}{\sigma(v(w), v(x)) + [v(w) - v(x)] \frac{\partial}{\partial v(w)} \sigma(v(w), v(x))} \\
&\leq \lim_{w \rightarrow \infty} \frac{2v'(w) \frac{\partial}{\partial v(w)} \sigma(v(w), v(x)) + v'(w) [v(w) - v(x)] \frac{\partial^2}{\partial v(w)^2} \sigma(v(w), v(x))}{[v(w) - v(x)] \frac{\partial}{\partial v(w)} \sigma(v(w), v(x))} \\
&= \lim_{w \rightarrow \infty} \left\{ \frac{2v'(w)}{v(w) - v(x)} + v'(w) \frac{\frac{\partial^2}{\partial v(w)^2} \sigma(v(w), v(x))}{\frac{\partial}{\partial v(w)} \sigma(v(w), v(x))} \right\} \\
&\leq \lim_{w \rightarrow \infty} \frac{2v'(w)}{v(w) - v(x)} = 0,
\end{aligned}$$

where the second inequality follows from the fact that $\lim_{\Delta \rightarrow \infty} \frac{\partial^2}{\partial \Delta^2} \sigma(x + \Delta, x) \leq 0$, as otherwise $\sigma(x + \Delta, x)$ could not be strictly increasing in Δ on $(0, \infty)$ and bounded from above.

Since $h(z)$ is continuous on (x, ∞) , it follows from $\lim_{w \rightarrow x} h(w) < \infty$ and $\lim_{w \rightarrow \infty} h(w) < \infty$ that $\max_{w \in (x, \infty)} h(w)$ exists. This, in turn, implies that there exists a constant $\alpha > 0$, such that Condition (14) is satisfied for any $z > x$, which was to be proven. \square

Proof of Theorem 1. The statement follows from Proposition 1 in Ebert and Strack (2015). We restate their argument here in terms of our notation. By Lemma 3, the auxiliary utility function is of exponential growth at $z = x$, so that we can find $\alpha, \beta \in \mathbb{R}_{>0}$ such that, for any $z \geq 0$, we have

$$[\tilde{u}(z) + \beta] \leq [\tilde{u}(x) + \beta] \exp(\alpha(z - x)). \quad (16)$$

Recall that the preferences of an EUT agent are invariant under positive affine transformations, which implies that the utility function $\hat{u}(z) := \tilde{u}(z) + \beta$ represents the exact same preferences. We should also keep in mind that $\hat{u}(x) = \beta > 0$.

Consider an EUT agent with a utility function \hat{u} , and an ABM $X_t = x + \mu t + \nu W_t$ with an initial value x and a drift $\mu < -\frac{1}{2}\alpha\nu^2 =: \tilde{\mu}$. For any stopping time τ with $\mathbb{P}_0[\tau > 0] > 0$, we have

$$\begin{aligned}
\mathbb{E}[\hat{u}(X_\tau)] &\leq \hat{u}(x) \mathbb{E}[\exp(\alpha(X_\tau - x))] \\
&= \hat{u}(x) \mathbb{E}[\exp(\alpha\mu\tau + \alpha\nu W_\tau)] \\
&= \hat{u}(x) \mathbb{E} \left[1 + \int_0^\tau \left(\alpha\mu + \frac{1}{2}\alpha^2\nu^2 \right) \exp(\alpha\mu s + \alpha\nu W_s) ds + \int_0^\tau \alpha\nu \exp(\alpha\mu s + \alpha\nu W_s) dW_s \right] \\
&= \hat{u}(x) \mathbb{E} \left[1 + \int_0^\tau \left(\alpha\mu + \frac{1}{2}\alpha^2\nu^2 \right) \exp(\alpha\mu s + \alpha\nu W_s) ds \right] \\
&< \hat{u}(x),
\end{aligned}$$

where the first inequality holds by (16), the second equality holds by Itô's Lemma, the third

equality holds by Doob's Optional Sampling Theorem, and the second inequality holds by $\hat{u}(x) > 0$ and the assumption that $\mu < \tilde{\mu}$ (so that the expectation is less than one). Hence, an EUT agent with a utility function \hat{u} and, thus, the naïve salient thinker does not start to gamble. \square

Proof of Corollary 1. Consider an Arithmetic Brownian Motion $X_t = x + \mu t + \nu W_t$ with a negative drift, which a naïve salient thinker with a linear value function would not start. We transform the process using the strictly increasing *scale function* (Revuz and Yor, 1999, pp. 302)

$$\Psi : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}, z \mapsto \int_0^z \exp\left(-2\frac{\mu}{\nu^2}y\right) dy,$$

which yields a scaled process $(\Psi(X_t))_{t \in \mathbb{R}_{\geq 0}}$ with zero drift and an initial value $\Psi(x)$.

Now consider a salient thinker with the exact same salience function, but a value function $v(z) = \frac{\nu}{2|\mu|} \ln\left(1 + \frac{2|\mu|}{\nu}z\right)$, which is strictly increasing and concave. Since $v(z) = \Psi^{-1}(z)$, we conclude that, for any stopping time τ , it has to hold that

$$\mathbb{E}[(v(\Psi(X_\tau)) - v(\Psi(x)))\sigma(v(\Psi(X_\tau)), v(\Psi(x)))] = \mathbb{E}[(X_\tau - x)\sigma(X_\tau, x)] \leq 0,$$

where the inequality follows from the assumption that the naïve salient thinker with a linear value function does not start the process $(X_t)_{t \in \mathbb{R}_{\geq 0}}$. Consequently, the naïve salient thinker with a value function $v(\cdot)$ does not start the scaled process $(\Psi(X_t))_{t \in \mathbb{R}_{\geq 0}}$ with zero drift. \square

A.4: Additional Results on Stop-Loss and Take-Profit Strategies

Proof of Proposition 2. Consider a stop-loss and take-profit strategy that corresponds to the threshold stopping time $\tau_{a,b}$ with $a := x_t - \epsilon - \epsilon'$ and $b := x_t + \epsilon$ for some $\epsilon > 0$, $\epsilon' \geq 0$, and $\epsilon + \epsilon' \leq x_t$, and that is therefore not a loss-exit strategy. Again, we denote as

$$\Phi_\mu(z) := \mathbb{P}_t[X_T \leq z | X_t = x_t, \tau_{a,b} \geq T]$$

the CDF of X_T conditional on reaching the expiration date. Then, it follows that

$$\begin{aligned}
U^s(X_{T \wedge \tau_{a,b}} | \mathcal{C}) - v(x_t) &\propto \mathbb{P}_t[\tau_{a,b} < T] \times \left[p(v(x_t - \epsilon - \epsilon') - v(x_t))\sigma(v(x_t - \epsilon - \epsilon'), v(x_t)) \right. \\
&\quad \left. + (1 - p)(v(x_t + \epsilon) - v(x_t))\sigma(v(x_t + \epsilon), v(x_t)) \right] \\
&\quad + \mathbb{P}_t[\tau_{a,b} \geq T] \times \int_{(a,b)} (v(z) - v(x_t))\sigma(v(z), v(x_t)) d\Phi_\mu(z) \\
&< \mathbb{P}_t[\tau_{a,b} < T] \times \sigma(v(x_t + \epsilon), v(x_t)) \times [pv(x_t - \epsilon - \epsilon') + (1 - p)v(x_t + \epsilon) - v(x_t)] \\
&\quad + \mathbb{P}_t[\tau_{a,b} \geq T] \times \int_{(-\epsilon, \epsilon)} (v(x_t + z) - v(x_t))\sigma(v(x_t + z), v(x_t)) d\tilde{\Phi}_\mu(z) \\
&< \mathbb{P}_t[\tau_{a,b} \geq T] \times \int_{(-\epsilon, \epsilon)} (v(x_t + z) - v(x_t))\sigma(v(x_t + |z|), v(x_t)) d\tilde{\Phi}_\mu(z) \\
&\leq \mathbb{P}_t[\tau_{a,b} \geq T] \times \int_{(0, \epsilon)} \left([v(x_t + z) - v(x_t)] - [v(x_t) - v(x_t - z)] \right) \\
&\quad \times \sigma(v(x_t + z), v(x_t)) d\tilde{\Phi}_\mu(z) \leq 0,
\end{aligned}$$

with $p = p(a, b, \mu)$ defined in Eq. (6). The first inequality follows from ordering, diminishing sensitivity, and the fact that $v(x_t + \epsilon) - v(x_t) \leq v(x_t) - v(x_t - \epsilon - \epsilon')$ due to the concavity of the value function as well as the construction of $\tilde{\Phi}_\mu$, which is defined as $\tilde{\Phi}_\mu(z) := \Phi_\mu(x_t + z)$ for any $z \geq -\epsilon$ and $\tilde{\Phi}_\mu(z) := 0$ for any $z < -\epsilon$. The second one follows from the drift being non-positive in combination with Jensen's Inequality, and diminishing sensitivity of the salience function. The weak inequality holds by Lemma 2 (f), and the last inequality holds by concavity of the value function, which implies $v(x_t + z) - v(x_t) \leq v(x_t) - v(x_t - z)$ for any $z > 0$. \square

Proof of Proposition 3. Let $\mu' < 0$. If the naïve salient thinker does start, there exists some threshold stopping time $\tau_{a,b}$ such that $U^s(X_{T \wedge \tau_{a,b}} | \mathcal{C}) > v(x)$. By Proposition 2, the stopping time $\tau_{a,b}$ represents a loss-exit strategy; that is, the stopping thresholds satisfy $x - a < b - x$. By Lemma 2 (d), as the drift increases to $\mu'' > \mu'$, the distribution of $X_{T \wedge \tau_{a,b}}$ improves in terms of first-order stochastic dominance. Hence, by Proposition 1 in Dertwinkel-Kalt and Köster (2020), also $U^s(X_{T \wedge \tau_{a,b}} | \mathcal{C})$ increases as we move from a drift μ' to a drift μ'' .

In sum, if the naïve salient thinker does start a process with drift $\mu' < 0$, he does start any process with drift $\mu'' > \mu'$. Likewise, if the naïve salient thinker does not start a process with $\mu' \leq 0$, then he does not start for any process with $\mu'' < \mu'$. \square

A.5: Salience and the Disposition Effect

Salience theory predicts the disposition effect if

$$\frac{\sigma(v(x_t - \epsilon), v(x_t))}{\sigma(v(x_t + \epsilon'), v(x_t))} \times \frac{v(x_t) - v(x_t - \epsilon)}{v(x_t + \epsilon') - v(x_t)} \quad (17)$$

is increasing in x_t . With a linear value function, this term simplifies to

$$\frac{\sigma(x_t - \epsilon, x_t)}{\sigma(x_t + \epsilon', x_t)} \times \frac{\epsilon}{\epsilon'},$$

and substituting Bordalo *et al.* (2012)'s salience function into this term gives

$$\frac{\frac{|x_t - (x_t - \epsilon)|}{|x_t| + |x_t - \epsilon| + \theta}}{\frac{|x_t - (x_t + \epsilon')|}{|x_t| + |x_t + \epsilon'| + \theta}} \times \frac{\epsilon}{\epsilon'} = \frac{|x_t| + |x_t + \epsilon'| + \theta}{|x_t| + |x_t - \epsilon| + \theta} \times \frac{|\epsilon|}{|-\epsilon'|} \times \frac{\epsilon}{\epsilon'}.$$

Since we exclude negative outcomes, x_t , $x_t + \epsilon'$, and $x_t - \epsilon$ must be positive, and we can remove the absolutes. Similarly, because ϵ and ϵ' are positive by definition, we have $|\epsilon| = \epsilon$ and $|-\epsilon'| = \epsilon'$, which gives

$$\frac{\epsilon^2(2x_t + \epsilon' + \theta)}{(\epsilon')^2(2x_t - \epsilon + \theta)}.$$

Taking the first derivative with respect to x_t yields

$$-\frac{2\epsilon^2(\epsilon' + \epsilon)}{(\epsilon')^2(2x_t + \theta - \epsilon)^2}.$$

Since ϵ and ϵ' are positive, the fraction's numerator and denominator are always positive. Hence, the whole term is always negative, and the original term (17) with the standard salience function and a linear value function is decreasing in x_t . Hence, salience theory's standard specification cannot explain the disposition effect.

A.6: Stopping a Process Prematurely

In this section, we present simulations that suggest that (1) loss-exit strategies induce right-skewed return distributions, and (2) salience theory could predict, after having started, stopping before the expiration date.

To do this, we simulate the development of an ABM with starting value $x_t = 100$, drift $\mu = 0$, standard deviation $\sigma = 5$ and expiration date $T = 100$, for a given loss-exit strategy $\tau_{60,180}$ with lower threshold 60 and upper threshold 180 exactly 500 million times at each $t \in \{0, \dots, 100\}$. If the ABM hits the bounds at any time, we terminate the process prematurely and register the boundary value as the outcome. This approach provides us with an approximation of the outcome distribution for the truncated ABM, enabling us to compute its skewness.

Figure 7 shows that loss-exit strategies indeed induce right-skewed return distributions, but close to the expiration date the return distribution is approximately symmetric as the selected thresholds are unlikely to be met before the expiration date.

For (2), we show that gambling becomes less attractive as time passes (i.e., t increases and the remaining time $T - t$ decreases). Specifically, we show that the salience-weighted utility of gambling (Definition 2) is decreasing in t for a range of points in time. For that, we need to assume a specific functional form of the salient thinker's value and salience functions. We use a linear value function and as a salience function we use an adaptation of the salience function




Figure 7: *Skewness of the outcome distribution induced by $\tau_{60,180}$*

$|x - y|/(|x| + |y| + 0.1)$ that Bordalo *et al.* (2012) use in their rank-based salience model (which is equal to the salience function (4) that we presented in the main text with $\theta = 0.1$ that Bordalo *et al.* use in their Appendix), namely

$$\sigma(x, y) = \delta \frac{|x - y|}{|x| + |y| + 0.1}.$$

Here, $\delta > 1$ controls the strengths of the salience distortion, and as there are no previous calibrations of the continuous salience model that could guide our choice of δ , we set $\delta = 2$.²⁵ With these functional-form assumptions and the simulated outcome distribution, we can calculate the salience-weighted utility of gambling at each $t \in \{0, \dots, 100\}$. However, to make this computationally feasible, we need to reduce the number of states by assigning the outcomes of our simulated distributions into bins of size 0.1.²⁶

Figure 8 shows that, like the skewness of the outcome distribution, the salience-weighted utility of gambling is decreasing over time for a considerable range of points in time. Initially, the salience-weighted utility exceeds 100, so that gambling is attractive, but falls below this value as the expiration date approaches. This demonstrates that a salient thinker might start to gamble and stop before the expiration date. Notably, the fact that the salience-weighted utility is increasing close to the expiration date follows from decreasing variance of the return distribution together with diminishing sensitivity.²⁷

²⁵Because we use the continuous salience model, we cannot rely on the calibration of δ of Bordalo *et al.* (2012) and others that relied on the rank-based salience model. Assuming $\delta = 2$ gives us only quite mild salience distortion, which however suffices for our purpose of demonstrating in-between stopping.

²⁶This is necessary because each simulation could lead to a different final value of the ABM, leaving us with up to 500 million states.

²⁷The intuition behind this result is the following: The return distribution close to the expiration date is quite




Figure 8: Saliency-Weighted Utility of Gambling over Time

Appendix B: Sophisticated Stopping Behavior Without Commitment

B.1: Statement of the Result

To solve for a sophisticate's stopping behavior, we adopt the equilibrium concept of Ebert and Strack (2018), which says that a stopping time τ constitutes an equilibrium if and only if at any point in time it is optimal to follow the strategy, taking as given that all future selves will do so.

Definition 6 (Equilibrium). *A stopping time τ constitutes an equilibrium if and only if at every point in time t it is optimal to take the prescribed decision, given that all future selves will follow this strategy.*

We find that a sophisticated salient thinker, who is restricted to choose stop-loss and take-profit strategies, does not start any process with a non-positive drift, which implies that naïvete is a necessary assumption to explain (unfair) casino gambling within the salience framework.

Proposition 4. *Suppose that the agent can only choose stop-loss and take-profit strategies. Fix an initial wealth level $x \in \mathbb{R}_{>0}$, and consider only processes with a non-positive drift. Then, in any equilibrium, the sophisticated salient thinker does not start.*

To fix ideas, let us assume that $T = \infty$. For any threshold stopping time $\tau_{a,b}$, there exists some wealth level $y' \in (a, b)$ such that the downside of the binary lottery $X_{\tau_{a,b}}$ is salient when evaluated in the choice set $\mathcal{C} = \{X_{\tau_{a,b}}, y'\}$. Moreover, if the process has a non-positive drift, then, at any wealth level y , we have $\mathbb{E}[X_{\tau_{a,b}}] \leq y$. Since a salient thinker, with a weakly concave value function, values a binary lottery with a salient downside strictly less than its expected value, the sophisticated agent anticipates to stop no later than at wealth level y' . Thus, by Definition

symmetric as the bounds are unlikely to be hit, and by diminishing sensitivity the lower payoffs are overweighted relative to the higher payoffs. This effect, however, decreases when the variance of the process decreases.

6, the threshold stopping time $\tau_{a,b}$ cannot constitute an equilibrium. In contrast, at any initial wealth level $x \in \mathbb{R}_{>0}$, not starting can be supported as an equilibrium outcome: given that all future selves will not start, the current self is indifferent between not starting and starting the process, so that it is indeed optimal to stop at every point in time. As we prove in the following, the argument extends to processes with a finite expiration date.

B.2: Proof of Proposition 4 for a Finite Expiration Date

Fix an initial wealth level x and a non-positive drift $\mu \leq 0$. It remains to be shown that the arguments presented above extend to processes with a finite expiration date $T \in \mathbb{R}_{>0}$.

Consider a stop-loss and take-profit strategy, which can be represented by a threshold stopping time $\tau_{a,b}$. We now argue that it cannot be an equilibrium to play according to stopping time $\tau_{a,b}$ with $b \in (x, \infty)$. At any time t with a wealth level $X_t = y \in (a, b)$, a salient thinker follows the stop-loss and take-profit strategy that is represented by $\tau_{a,b}$ if and only if

$$\begin{aligned} \mathbb{P}_t[\tau_{a,b} < T] \times \left[p(v(a) - v(y))\sigma(v(a), v(y)) + (1 - p)(v(b) - v(y))\sigma(v(b), v(y)) \right] \\ + \mathbb{P}_t[\tau_{a,b} \geq T] \times \int_{(a,b)} (v(z) - v(y))\sigma(v(z), v(y)) d\Phi_\mu(z) \geq 0 \end{aligned}$$

holds, where the probability $p = (a, b, \mu)$ is defined as in Eq. (6) and where the conditional CDF $\Phi_\mu(z) := \mathbb{P}_0[X_T \leq z | X_0 = x_0, \tau_{a,b} \geq T - t]$ is described in Lemma 2.

Notice that $\sigma(v(a), v(y)) > \sigma(v(b), v(y))$ holds for any wealth level y sufficiently close to b . Also, we have $\mathbb{E}_t[X_{\tau_{a,b}} | X_t = y] \leq y$ due to the non-positive drift. This implies, together with the concave value function, that, for any wealth level y sufficiently close to b , it holds that

$$p(v(a) - v(y))\sigma(v(a), v(y)) + (1 - p)(v(b) - v(y))\sigma(v(b), v(y)) < 0.$$

Since, for any fixed t , we have $\lim_{y \rightarrow b} \mathbb{P}_t[\tau_{a,b} < T] = 1$ by Lemma 2 (b), we thus conclude that for any $\tau_{a,b}$ there is some $y' \in (a, b)$ such that

$$\begin{aligned} \mathbb{P}_t[\tau_{a,b} < T] \times \left[p(v(a) - v(y'))\sigma(v(a), v(y')) + (1 - p)(v(b) - v(y'))\sigma(v(b), v(y')) \right] \\ + \mathbb{P}_t[\tau_{a,b} \geq T] \times \int_{(a,b)} (v(z) - v(y'))\sigma(v(z), v(y')) d\Phi_\mu(z) < 0. \end{aligned}$$

Hence, if the agent is restricted to choose from the set of all stop-loss and take-profit strategies, there exists no equilibrium in which a sophisticated salient thinker does start. By the same arguments as for the case of $T = \infty$, not starting can be supported as an equilibrium outcome, which proves the claim. \square

Appendix C: Additional Details on the Main Experiment

This appendix contains supplementary material to the experiment that we conducted.

C.1: Parameters and Layout of the Static Choices

After making the six selling decisions, subjects had to choose twelve times between a binary lottery and the safe option paying its expected value. The parameters of the lotteries as well as the classification of skewness-seeking choices are depicted in Table 2. Figure 9 further illustrates the layout that we used for these static choices in the experiment.

Lottery	Safe Option	Skewness	Skewness-Seeking Choice
(37.5, 80%; 0, 20%)	30	-1.5	Safe
(41.25, 64%; 10, 36%)	30	-0.6	Safe
(45, 50%; 15, 50%)	30	0	Safe
(60, 20%; 22.5, 80%)	30	1.5	Lottery
(75, 10%; 25, 90%)	30	2.7	Lottery
(135, 2%; 27.85, 98%)	30	6.9	Lottery
(57.5, 80%; 20, 20%)	50	-1.5	Safe
(61.25, 64%; 30, 36%)	50	-0.6	Safe
(65, 50%; 35, 50%)	50	0	Safe
(80, 20%; 42.5, 80%)	50	1.5	Lottery
(95, 10%; 45, 90%)	50	2.7	Lottery
(155, 2%; 47.85, 98%)	50	6.9	Lottery

Table 2: Lotteries used to elicit skewness seeking in static choices. The safe option is equal to the lottery's expected value. In addition, all lotteries have the same variance, so that the first and the second set of lotteries, respectively, differ only in terms of skewness.




Figure 9: The figure illustrates the layout of the static choices in the experiment (in German).

C.2: Experimental Instructions

Screen 1 — Instructions: Overview of the Experiment

Please note that you are not allowed to use your mobile phone or talk to other participants during the experiment. After you have finished the experiment and your payment appears on the screen, please stay seated and wait for the other participants to finish. At this point you are allowed to use your phone again. If you have a question, please raise your hand and a lab assistant will come to you.

In this experiment you will make investment decisions. More precisely, you will have to decide at what time and price you want to sell an asset. The price at which you can sell the asset will change over time.

In total you will make 6 such investment decisions. At the end of the experiment, we will choose one of your decisions at random and pay you the price at which you sold this asset. Irrespective of this, you will receive a show up fee of 4 Euro. During the experiment, we will denote all monetary values in the currency Taler, which will be converted to Euro at an exchange rate of 1 Euro = 10 Taler.

The only thing that changes between the different decisions is the long-term profitability of the asset. The maximum time for which you can hold the asset will be 10 seconds in all decisions. If you do not sell the asset in the first 10 seconds, it will automatically be sold at its price after 10 seconds. The initial value of the asset will always be 100 Taler.

In the following, we will explain to you the development of the asset step by step. In particular, we will show you how the long-term profitability varies across the different assets. Moreover, we will explain in detail which selling strategies you will be able to choose.

Screen 2 — Instructions: Development of the Asset Price

Below you can see a graph, which depicts the development of the price of an asset. As soon as you press "Start", a line which represents the value of the asset will appear.

Please press "Start" now.

[Subjects are shown a graph of an exemplary price path with a final price of 100 Taler.]

As mentioned previously, you cannot hold the asset for longer than 10 seconds. The final asset price is 100 Taler.

Screen 3 — Instructions: Different Drifts

In this experiment, you will see assets of varying profitability. How profitable an asset is in the long run is described by the drift of the asset. The drift denotes the average change in the value of the process per second.

A positive drift implies that the asset will increase in value in the long run, while a negative

drift implies that the asset will decrease in value in the long run. Notice that the value of the asset varies over time. Hence, even an asset with a negative drift sometimes increases in value.

In order for you to get a feeling for how the value of an asset changes with the drift, we will show you a few examples of different drifts on the next screens.

Screens 4-6 — Instructions: The Drift of an Asset

The drift of this asset is 0 [or 2 or -2]. Please press "Start" and watch the development of the asset's price.

[Subjects are shown a graph of an exemplary price path with a final price of 100 Taler for the process with drift 0, 120 Taler for the process with drift 2, and 80 for the process with drift -2.]

Screen 7 — Instructions: No Negative Prices

The asset does never take a negative value. Once the asset's value reaches zero, it does not rise again, but will stay at zero permanently. Please press "Start" and watch the development of the asset's price.

[Subjects are shown a graph illustrating that the process is absorbing in zero.]

Screen 8 — Instructions: The Process is not Bounded from Above

Independent of the drift, the value of the asset can, in principle, become arbitrarily large. The probability that the asset's value indeed becomes very large is the smaller the more negative the drift is. But even an asset with a very negative drift can attain a very large value.

Screen 9 — Instructions: Strategies with an Upper and a Lower Bound

In each decision, you will set an upper and a lower bound at which you are willing to sell the asset. If the price reaches the upper bound, the process will stop and you will be able to sell the asset. If you sell the asset, you will receive the price that you have set as the upper bound. If the price reaches the lower bound, you can also sell the asset. In this case you will receive the price that you have set as the lower bound.

The upper bound must always be above the current value of the process. The lower bound must always be below the current value of the process. You can adjust the bounds by clicking on the red lines and moving them around. Important: throughout the experiment, you will have to move the upper and the lower bound at least once, before you can start the process. Please move the bounds now and then click "Start".

[Subjects are shown a graph similar to the ones depicted in Figure 1. After moving the bounds and starting the process, subjects cannot pause the process or adjust the bounds anymore.]

Screen 10 — Instructions: Pausing the Process

After you have started the process, you can pause it at any time. While the process is paused, you can move the upper and the lower bound. While the process is moving, you cannot move the bounds.

Now you have to complete the following steps in the order listed below:

1. Move the upper and the lower bound.
2. Start the process.
3. Pause the process.
4. Move the upper and the lower bound again.
5. Start the process again.

[Subjects are shown a graph similar to the ones depicted in Figure 1, but without the opportunity to sell the asset immediately.]

Screen 11 — Instructions: Sell Immediately

Before you start the process, you can instead sell the asset immediately by clicking on "Sell Immediately". You can only do this before you start the process for the first time. After you have started the process, the "Sell Immediately" button will disappear. Afterwards the process will only stop prematurely if it either hits the upper or the lower bound. You can now either "Sell Immediately" or — after moving each bound at least once — start the process.

[Subjects are shown a graph similar to the ones depicted in Figure 1.]

Screen 12 — Instructions: Change Bounds Before Starting the Process

In the first 10 seconds on each decision screen, you will only be able to move the bounds. After that you can "Sell Immediately" or start the process.

If the process reaches either bound, it stops and you can sell the asset at the price at which this bound is set. Alternatively, you can move the bounds and restart the process. Corresponding buttons for both options will appear once the process reached a bound.

[Subjects are shown a graph similar to the ones depicted in Figure 1.]

Screen 13 — The Task is About to Start

You will now participate in three practice rounds. Afterwards you will play the decision round. The drift in the practice rounds will be identical to the one in the decision round. The practice rounds will give you the opportunity to get an intuition for how the process will develop during the decision round.

The drift in the practice rounds and the subsequent decision round will be 0 [or -1 or -3 or -5 or -10 or -20].

Screen 14 — Practice Round

The drift in the practice rounds and the subsequent decision round will be 0 [or -1 or -3 or -5 or -10 or -20].

[Subjects are shown a graph as depicted in the left panel of Figure 2.]

Screen 15 — Instructions: The drift of an asset

On this page you see 10 exemplary paths of an asset with a drift of 0 [or -1 or -3 or -5 or -10 or -20].

[Subjects are shown a graph as depicted in the right panel of Figure 2.]

Screen 16 — Decision

The practice rounds are over — now it is getting serious! Please make your selling decision. The drift in this round is 0 [or -1 or -3 or -5 or -10 or -20].

[Subjects are shown a graph as depicted in Figure 1.]

Screen 17 — Additional Questions I: Instructions

On the next pages you will make 12 choices between a lottery and a safe payoff. From now on all outcomes will be displayed in Euro.

At the end of the experiment, we will select one participant of this session at random. For this participant, we will randomly select one of the 12 decisions and determine the outcome of the chosen lottery. This participant will receive the corresponding payoff from the chosen lottery.

Example

If you select Lottery 1 in the example below, you will receive either 135 Euro or 27.85 Euro. The probability that you receive 135 Euro is 2% and the probability that you receive 27.85 Euro is 98%. Alternatively, if you select Lottery 2, you will receive 30 Euro for sure.

[Subjects are shown the graph depicted in Figure 9.]

Screen 18: Additional Questions I - Decision 1

Please choose a lottery. As soon as you have chosen a lottery, a button labelled "Next Page" will appear.

[Subjects are shown the graph as depicted in Figure 9.]

Screen 19 — Additional Questions II

Please answer the following questions. For every correct answer, you will receive one Taler.

If 10 machines take 10 minutes to make 10 nails, how many minutes do 100 machines need to make 100 nails?

A part of a pond is covered with water lilies. Every day the area covered with water lilies doubles. If it takes 24 days until the whole pond is covered with water lilies, how many days does it take until half of the pond is covered with water lilies?

If three elves can wrap three presents in one hour, how many elves does it take to wrap six presents in two hours?

Jerry has both the 15th best and the 15th worst grade in his class. How many students are in the class?

In a sports team tall members are three times as likely to win medals as short members. This year the team won 60 medals in total. How many medals were won by short team members?

Screen 20 — Additional Questions III

Please answer the following questions. For every correct answer, you will receive one Taler.

Suppose you had 100 Euro in a savings account and the interest rate was 2% per year. After 5 years, how much do you think you would have in the account if you left the money to grow?

[Options: "More than 102 Euro", "Exactly 102 Euro", "Less than 102 Euro".]

Suppose you had 100 Euro in a savings account and the interest rate was 20% per year and you never withdraw money or interest payments. After 5 years, how much would you have on this account in total?

[Options: "More than 200 Euro", "Exactly 200 Euro", "Less than 200 Euro".]

Imagine that the interest rate on your savings account was 1% per year and the inflation was 2% per year. After 1 year, how much would you be able to buy with the money in this account?

[Options: "More than today", "As much as today", "Less than today".]

Assume a friend inherits 10.000 Euro today and his brother inherits 10.000 Euro three years from now. There is a positive interest rate. Who is richer because of the inheritance?

[Options: "My friend", "Her brother", "Both are equally rich".]

Suppose that your income and all prices double in the next year. How much will you be able to buy with your income?

[Options: "More than today", "As much as today", "Less than today".]

Screen 21 — General Information About You

Please enter your age:

Please choose your gender:

Screen 22 — Enter Station Number

Please enter your station number:

Screen 23 — Payment

Your decision from round 2 will be paid.

You sold the asset for 100.00 Taler.

You received 1 Taler from answering the additional questions.

You are the participant whose lottery choice is paid. You receive an additional 80.00 Euro from the lottery.

Your payment including the show up fee of 4 Euro is 94.10 Euro.

Appendix D: Salience Predictions on Static Skewness Seeking

In this section, we extend a result from Dertwinkel-Kalt and Köster (2020) on a salient thinker's skewness seeking in static settings from the case of a linear value function to the case of a weakly concave value function. Assuming a linear value function, Dertwinkel-Kalt and Köster (2020) study, in particular, a salient thinker's choice between a binary lottery with an expected value E , a variance V , and a skewness S , which we denote by $L(E, V, S)$, and the safe option paying the lottery's expected value E with certainty, and they show that:

Proposition 5 (Dertwinkel-Kalt and Köster, 2020). *There exists some $\hat{S} = \hat{S}(E, V) \in \mathbb{R}$, such that a salient thinker with a linear value function chooses $L(E, V, S)$ over E if and only if $S > \hat{S}$.*

The proposition says that a salient thinker with a linear value function chooses a binary lottery over its expected value if and only if this lottery is sufficiently skewed. In the following, we will show that the same comparative static holds when assuming a weakly concave value function. This provides a theoretical foundation for why we look at the empirical relationship between a subject's share of skewness-seeking choices in the static choices and the share of loss-exit strategies this subject has chosen in the stopping problems (see Result 44 and Figure 5). A positive correlation between the two measures indicates that a subject seeking skewness consistent with salience theory in static choices does so also in dynamic choices.

To begin with, recall that the parameters of the binary lottery $L(E, V, S)$ — i.e. the outcomes x_1 and x_2 as well as the probability p that x_1 is realized — are uniquely defined by (Ebert, 2015):

$$x_1 = E - \sqrt{\frac{V(1-p)}{p}}, \quad x_2 = E + \sqrt{\frac{Vp}{1-p}}, \quad \text{and } p = \frac{1}{2} + \frac{S}{2\sqrt{4+S^2}}.$$

Now consider a salient thinker with a weakly concave value function $v(\cdot)$, who faces the choice between the lottery $L(E, V, S)$ and the safe option paying its expected value E with certainty. The salient thinker chooses lottery $L(E, V, S)$ over its expected value E if and only if

$$\begin{aligned} p \left[v \left(E - \sqrt{\frac{V(1-p)}{p}} \right) - v(E) \right] \sigma \left(v \left(E - \sqrt{\frac{V(1-p)}{p}} \right), v(E) \right) \\ + (1-p) \left[v \left(E + \sqrt{\frac{Vp}{1-p}} \right) - v(E) \right] \sigma \left(v \left(E + \sqrt{\frac{Vp}{1-p}} \right), v(E) \right) > 0, \end{aligned}$$

or, equivalently,

$$\pi \frac{v \left(E + \sqrt{\frac{V}{\pi}} \right) - v(E)}{v(E) - v \left(E - \sqrt{V\pi} \right)} > \frac{\sigma \left(v \left(E - \sqrt{V\pi} \right), v(E) \right)}{\sigma \left(v \left(E + \sqrt{\frac{V}{\pi}} \right), v(E) \right)}, \quad (18)$$

where $\pi := \frac{1-p}{p}$ denotes the relative likelihood of the lottery's upside. To establish that a salient thinker chooses the lottery if and only if it is sufficiently skewed, we will show that both the left-hand side and the right-hand side of (18) are monotonic in the likelihood ratio π ; namely,

that the left-hand side decreases in π , while the right-hand side increases in π .

First, by the ordering property, the right-hand side of (18) monotonically increases in π . Second, we observe that the left-hand side monotonically decreases in the likelihood ratio π :

$$\begin{aligned}
\frac{\partial}{\partial \pi} \left(\pi \frac{v\left(E + \sqrt{\frac{V}{\pi}}\right) - v(E)}{v(E) - v\left(E - \sqrt{V\pi}\right)} \right) &= \frac{v\left(E + \sqrt{\frac{V}{\pi}}\right) - v(E)}{v(E) - v\left(E - \sqrt{V\pi}\right)} \\
&\quad - \frac{\sqrt{\frac{V}{\pi}} v'\left(E + \sqrt{\frac{V}{\pi}}\right) [v(E) - v\left(E - \sqrt{V\pi}\right)] + \sqrt{V\pi} v'\left(E - \sqrt{V\pi}\right) [v\left(E + \sqrt{\frac{V}{\pi}}\right) - v(E)]}{[v(E) - v\left(E - \sqrt{V\pi}\right)]^2} \\
&\propto \left[v\left(E + \sqrt{\frac{V}{\pi}}\right) - v(E) - \sqrt{\frac{V}{\pi}} v'\left(E + \sqrt{\frac{V}{\pi}}\right) \right] [v(E) - v\left(E - \sqrt{V\pi}\right)] \\
&\quad - \sqrt{V\pi} v'\left(E - \sqrt{V\pi}\right) \left[v\left(E + \sqrt{\frac{V}{\pi}}\right) - v(E) \right] \\
&\propto 1 - \frac{v'\left(E + \sqrt{\frac{V}{\pi}}\right)}{\underbrace{\frac{v\left(E + \sqrt{\frac{V}{\pi}}\right) - v(E)}{\sqrt{\frac{V}{\pi}}}}_{\geq 0 \text{ since } v(\cdot) \text{ is increasing}}} - \frac{v'\left(E - \sqrt{V\pi}\right)}{\underbrace{\frac{v(E) - v\left(E - \sqrt{V\pi}\right)}{\sqrt{V\pi}}}_{\geq 1 \text{ since } v(\cdot) \text{ is concave}}} \\
&\leq 0,
\end{aligned}$$

where, after taking the derivative, we first multiply by $[v(E) - v(E - \sqrt{V\pi})]^2$ and rearrange, and then divide by $v(E) - v(E - \sqrt{V\pi})$ and $v(E + \sqrt{\frac{V}{\pi}}) - v(E)$ to arrive at the final expression. Combining these two observations, we conclude that there exists some $\hat{\pi} \geq 0$, such that (18) is satisfied if and only if $\pi < \hat{\pi}$. Since π monotonically decreases in the probability p , and since the probability p monotonically increases in the skewness S , we arrive at the following result:

Proposition 6. *There exists some $\hat{S} = \hat{S}(E, V) \in \mathbb{R} \cup \{\infty\}$, such that a salient thinker with a weakly concave value function chooses $L(E, V, S)$ over E if and only if $S > \hat{S}$.*

This proposition confirms that the comparative static on the lottery's skewness derived in Dertwinkel-Kalt and Köster (2020), under the assumption of a linear value function, is robust to allowing for a weakly concave value function. The only difference compared to the result in Dertwinkel-Kalt and Köster (2020) is that a salient thinker with a sufficiently concave value function will not choose the binary lottery, irrespective of how skewed it is. Formally, it follows that the threshold value \hat{S} in Proposition 6 satisfies $\hat{S} < \infty$ if and only if

$$\lim_{\pi \rightarrow 0} \frac{\partial}{\partial \pi} \left(\pi \frac{v\left(E + \sqrt{\frac{V}{\pi}}\right) - v(E)}{v(E) - v\left(E - \sqrt{V\pi}\right)} - \frac{\sigma\left(v\left(E - \sqrt{V\pi}\right), v(E)\right)}{\sigma\left(v\left(E + \sqrt{\frac{V}{\pi}}\right), v(E)\right)} \right) < 0,$$

which depends both on the curvature of the value and on the curvature of the salience function. But, as illustrated in Proposition 5, the above inequality is certainly satisfied for a linear value function and, by continuity, it will hold for mildly concave value functions as well.

Appendix E: Additional Figures




Figure 10: The left panel depicts the share of subjects holding the asset until the expiration date, separately for the different drifts. The right panel depicts the share of subjects selling the asset immediately.




Figure 11: The left panel depicts the share of subjects holding the asset until the expiration date, separately for the different drifts and below- and above-median subjects in terms of cognitive skills. The right panel depicts the share of below- and above-median subjects selling the asset immediately.




Figure 12: The figure depicts the share of initial loss-exit strategies chosen for the different drifts.

		After		After		
		Loss-Exit	Gain-Exit	Loss-Exit	Gain-Exit	
Before	Loss-Exit	63.31%	10.31%	Loss-Exit	72.65%	0%
	Gain-Exit	12.01%	14.37%	Gain-Exit	3.48%	23.87%

Figure 13: The left (right) table gives a categorization of all strategy adjustments that we observe throughout the experiment when a bound (no bound) is hit. "Before" indicates, in the left table, which type of strategy the subject has chosen last, and, in the right table, the type of strategy that is played in the moment in which the process is paused.

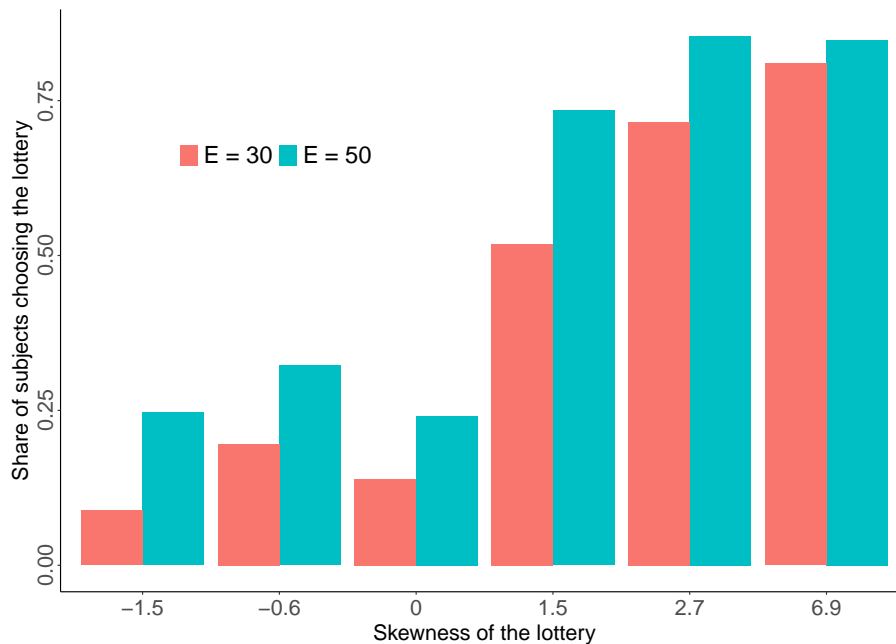


Figure 14: The figure depicts the share of subjects choosing each of the lotteries depicted in Table 2 over its expected value.

Appendix F: Stopping Behavior under Cumulative Prospect Theory

In this section, we analyze the stopping behavior of a naïve CPT-agent under the assumption of a finite expiration date. Ebert and Strack (2015) study the case without an expiration date and show that, under mild regularity assumptions on the probability weighting function,²⁸ a naïve CPT-agent never stops an ABM irrespective of its drift. In what follows, we will show numerically that this strong result still holds for a finite expiration date.

CPT preferences. Let X be a real-valued random variable. A CPT-agent evaluates each outcome of this random variable relative to a *reference point* $r \in \mathbb{R}$ via a strictly increasing *value function* $U : \mathbb{R} \rightarrow \mathbb{R}$. All outcomes larger than the reference point are classified as *gains*, while outcomes smaller than the reference point are classified as *losses*. Throughout this section, we assume a (weakly) *S-shaped* value function (Ebert and Strack, 2015, Online Appendix W.2),

$$U(x) = \begin{cases} (x - r)^\alpha & \text{if } x \geq r, \\ -\lambda \cdot (r - x)^\alpha & \text{if } x < r, \end{cases} \quad (19)$$

with parameters $\alpha \in (0, 1]$ and $\lambda > 1$.²⁹ According to Tversky and Kahneman (1992), cumulative probabilities are distorted by a *weighting function*. More precisely, there are (potentially different) non-decreasing weighting functions $w^-, w^+ : [0, 1] \rightarrow [0, 1]$ for gains and losses with $w^-(0) = w^+(0) = 0$ and $w^-(1) = w^+(1) = 1$. Throughout this section, we use the following weighting functions, which have been proposed by Tversky and Kahneman (1992):³⁰

$$w^-(p) = \frac{p^\delta}{(p^\delta + (1-p)^\delta)^{1/\delta}} \quad \text{and} \quad w^+(p) = \frac{p^\gamma}{(p^\gamma + (1-p)^\gamma)^{1/\gamma}} \quad \text{for } 0.279 < \delta, \gamma < 1.$$

The CPT value of the random variable X can be defined as (see Kothiyal *et al.*, 2011)

$$\begin{aligned} CPT(X) &:= \int_{\mathbb{R}_+} w^+(\mathbb{P}[U(X) > y]) dy - \int_{\mathbb{R}_-} w^-(\mathbb{P}[U(X) < y]) dy \\ &= \int_{\mathbb{R}_+} w^+(\mathbb{P}[X > r + y^{1/\alpha}]) dy - \int_{\mathbb{R}_-} w^-(\mathbb{P}[X < r - (-y)^{1/\alpha}/\lambda]) dy, \end{aligned} \quad (20)$$

where the second equality holds due to a (weakly) S-shaped value function in Eq. (19).

Stopping strategies. Consider a threshold stopping time $\tau_{a,b}$ and therefore induces a random wealth level $X_{T \wedge \tau_{a,b}}$. If the reference point $r \in \mathbb{R}$ satisfies $r \in [a, b]$, then the CPT value associated

²⁸In their Online Appendix W.1, Ebert and Strack (2015) verify that common CPT specifications satisfy the sufficient conditions that they impose on the probability weighting function to derive their main result.

²⁹As argued in Wakker (2010, p. 270), the model is ill-specified when taking different α for gains and losses.

³⁰The bounds on the parameters are taken from Dhami (2016, p. 122).

with this random variable is given by

$$\begin{aligned} CPT(X_{T \wedge \tau_{a,b}}) &= \int_{(0, (b-r)^\alpha)} w^+(\mathbb{P}_t[X_{T \wedge \tau_{a,b}} > r + y^{1/\alpha}]) dy \\ &\quad - \int_{(-\lambda(r-a)^\alpha, 0)} w^-(\mathbb{P}_t[X_{T \wedge \tau_{a,b}} < r - (-y)^{1/\alpha}/\lambda]) dy. \end{aligned} \quad (21)$$

For $a \geq r$, in contrast, the CPT value of the random variable $X_{T \wedge \tau_{a,b}}$ equals

$$CPT(X_{T \wedge \tau_{a,b}}) = \int_{((a-r)^\alpha, (b-r)^\alpha)} w^+(\mathbb{P}_t[X > r + y^{1/\alpha}/\lambda]) dy + (a-r)^\alpha, \quad (22)$$

while for $b \leq r$ it is given by

$$CPT(X_{T \wedge \tau_{a,b}}) = - \int_{(-\lambda(r-a)^\alpha, -\lambda(r-b)^\alpha)} w^-(\mathbb{P}_t[X < r - (-y)^{1/\alpha}/\lambda]) dy - \lambda(r-b)^\alpha. \quad (23)$$

At time $t < T$ with a current wealth level $x_t \in \mathbb{R}_{>0}$, we consider the following class of threshold stopping times:³¹ for $k \in \mathbb{R}_{>0}$ and $p \in (0, \frac{1}{2})$, define $a_{t,k} = x_t - k \cdot p$ and $b_{t,k} = x_t - k \cdot (1-p)$. Notice that, for any drift $\mu \leq 0$, these threshold stopping times are not only loss-exit strategies, but also induce a right-skewed distribution of returns.

Numerical analysis of stopping behavior. To ease the illustration of the results, we assume that the reference point constantly adjusts to the current wealth level (i.e., $r_t = x_t$ for any t). This implies, in particular, that the wealth level itself does not matter for a CPT-agent's stopping behavior, which makes the numerical analysis much more convenient. Based on the estimates in Tversky and Kahneman (1992), we set $\alpha = 0.88$ and $\lambda = 2.25$ as well as $\delta = 0.69$ and $\gamma = 0.61$.

Assuming a drift of $\mu = -2$ and a volatility of $\nu = 5$, Figure 15 depicts, for a given point in time t , the CPT value of the random variable $X_{T \wedge \tau_{a_{t,k}, b_{t,k}}}$ as a function of the remaining time, $T - t$, until the expiration date for the strategies with $k \in \{2, 4, 6, 8, 10\}$ and $p = 0.01$. Since we have $r_t = x_t$ by assumption, a CPT-agent does not stop at time t as long as there exists a stopping strategy that yields a strictly positive CPT value. We observe from Figure 15 that for any remaining time until the expiration date, there indeed exists a threshold stopping time that yields a strictly positive CPT value. When shifting the stopping thresholds closer to the current wealth level (by shifting the parameter k closer to zero), we obtain a similar picture for any arbitrarily negative drift. Hence, at least for the chosen parameter values, a naïve CPT-agent does not stop before the expiration date or, in other words, the stark never-stopping result derived by Ebert and Strack (2015) still holds for a finite expiration date.

Figure 15 highlights a couple of numerical regularities that are suggestive for the result not to hinge on the exact parameters chosen here: First, the CPT value derived from the depicted stopping strategies becomes flat in the remaining time until the expiration date relatively quickly and the earlier so the closer the stopping thresholds are to the current wealth level (i.e., the closer is k to zero). This suggests that the result derived by Ebert and Strack (2015) — which is

³¹These strategies are similar to those used in the proof of Theorem 1 in Ebert and Strack (2015).

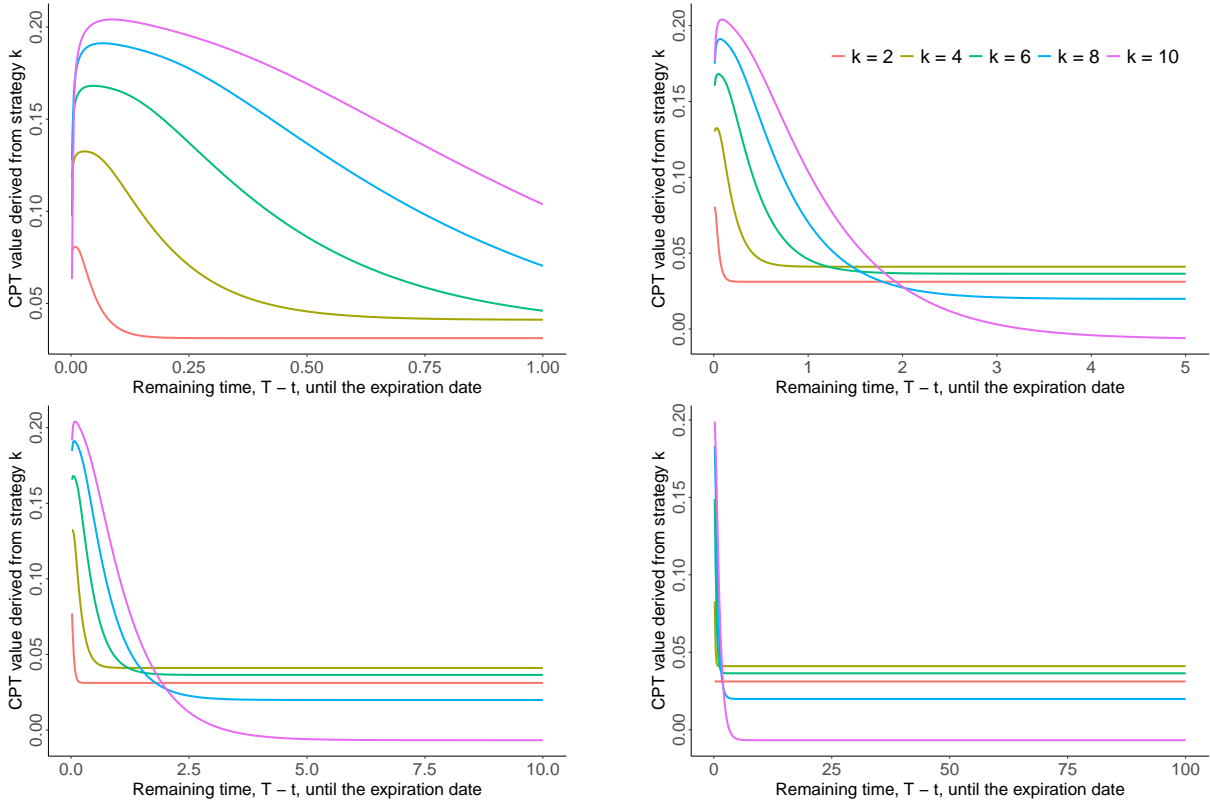


Figure 15: The figure depicts $CPT(X_{T \wedge \tau_{\alpha_t, k, b_t, k}})$ as a function of the remaining time, $T - t$, until the expiration date for time invariant strategies with $k \in \{2, 4, 6, 8, 10\}$ and $p = 0.01$ as described above. We assume a drift parameter of $\mu = -2$ and a volatility parameter of $\nu = 5$. The preference parameters are set to $\alpha = 0.88$ and $\lambda = 2.25$, and the parameters of the weighting function are $\delta = 0.69$ and $\gamma = 0.61$.

proven by the explicit use of strategies with thresholds close to the current wealth level — does not rely on T being infinity, but should hold already for rather short expiration dates. Second, as the remaining time until the expiration date gets smaller, the CPT value of the depicted loss-exit strategies increases (before it eventually falls to zero). Both patterns are robust to different specifications (e.g., a piece-wise linear value function with $\alpha = 1$ or a reference point of $r_t = 0$). This suggests that the never-stopping result derived by Ebert and Strack (2015) is indeed robust to allowing for a finite expiration date. All numerical results are available upon request.

Appendix G: Saliency-Driven Decoy Effects

G.1: Overview of the Experimental Design and Results

The most striking difference compared to the alternative models discussed above is that saliency theory is first and foremost a model of context-dependent behavior: the evaluation of an option depends on the alternatives at hand. To conclusively rule out models of context-independent behavior, we next study the role of (dominated) *decoys* in stopping problems. Such decoy effects are not only inconsistent with the models discussed in the preceding section, but also with other prominent models such as expectation-based loss aversion (Kőszegi and Rabin, 2007).

The modified stopping problem. Suppose there is no longer just one asset, but two assets. We refer to these assets as $G(\text{green})$ and $B(\text{blue})$, and assume that their prices follow the ABMs

$$dX_t = \mu_G dt + \nu dW_t \quad \text{and} \quad dY_t = \mu_B dt + \nu dU_t, \quad \text{respectively.}$$

Both assets share the same initial value $X_0 = z = Y_0$ and the same volatility $\nu \in \mathbb{R}_{>0}$, but Asset G has a larger drift than Asset B , $\mu_G > \mu_B$. We further assume that the standard Brownian Motions $(W_t)_{t \in \mathbb{R}_{\geq 0}}$ and $(U_t)_{t \in \mathbb{R}_{\geq 0}}$ are independent of each other. There is no expiration date.

	$p_G p_B$	$p_G(1 - p_B)$	$(1 - p_G)p_B$	$(1 - p_G)(1 - p_B)$
Invest in Asset G	a	a	b	b
Invest in Asset B	a	b	a	b
No Investment	z	z	z	z

Table 3: Joint distribution of the different options.

If the agent invests in either asset, he is restricted to choose a loss-exit strategy, represented by the threshold stopping time $\tau_{a,b}$ with $a < z < b$, which he *cannot* revise over time. Since there is no expiration date, the asset is sold at one of these thresholds, and since the agent cannot revise his strategy, this happens in finite time with probability one. Denote as $p_G := p(a, b, \mu_G)$ the probability of Asset G being sold at the lower price $a < z$, and as $p_B := p(a, b, \mu_B)$ the corresponding probability for Asset B . Since $\mu_G > \mu_B$ and, therefore, $p_G < p_B$, investing in Asset B is dominated (in the sense of first-order stochastic dominance) by investing in Asset G . We, therefore, refer to Asset B as a (*dominated*) *decoy* (e.g., Huber *et al.*, 1982). We compare two scenarios: the choice set is either $\{X_{\tau_{a,b}}, z\}$ (*no decoy*) or $\{X_{\tau_{a,b}}, Y_{\tau_{a,b}}, z\}$ (*decoy*).

Salience-driven decoy effects. According to salience theory, investing in Asset G can become more attractive in the presence of the dominated Asset B . More precisely, a salient thinker compares the selling price of Asset G state-by-state to some *reference point* given by a convex combination of the other options' outcomes (see Table 3 for the joint distribution of the different options).³² As a consequence, when adding the dominated decoy to the choice set, the salience of Asset G 's upside — i.e. selling it at a price b — changes: by the contrast effect, it becomes more salient whenever the dominated asset is sold at a price $a < z$ and less salient when it is sold at $b > z$. Likewise, the downside of investing in Asset G — i.e. selling it at a price a — becomes more salient when the dominated asset is sold at a price $b > z$ and less salient when it is sold at a price $a < z$ instead. Since the dominated asset is the more likely to be sold at the lower price the more negative its drift μ_B is, the upside of Asset G becomes relatively more salient as μ_B decreases. In other words, according to salience theory, (i) the presence of the dominated asset can boost demand for Asset G , and (ii) this decoy effect is (weakly) stronger for dominated assets with a more negative drift (see Appendix G.3 for a formal derivation).

³²While Bordalo *et al.* (2012) assume that the reference point is given by the state-wise average over the alternative options, we allow for a much more flexible functional form (see Appendix G.2 for details).

Experimental design and implementation. We conducted an online experiment to test for such decoy effects on stopping behavior.³³ As in the main experiment, we set the initial price to $z = 100$ Taler (this time converted to £ at a ratio of 60:1) and the volatility to $\nu = 5$. Throughout, Asset *G* has zero drift, while the dominated asset's drift is either $\mu_B = -10$ or $\mu_B = -20$. Subjects could always choose the outside option of *No Investment*. Figure 16 illustrates the decision screens with (right panel) and without (left panel) a dominated decoy.



Figure 16: Screenshots of the decision screens with and without a decoy.

If a subject decides to invest in an asset, he can sell it only at pre-specified prices 90 and 190. More precisely, if a subject invests, the price of the asset will change until it reaches either 90 or 190. Recall that such a loss-exit strategy is potentially attractive to a salient thinker (see Proposition 2). If a subjects does not invest, he receives an asset's initial price with certainty.

Each subject made three investment decisions: one decision with a binary choice set (*Asset Green* vs. *No Investment*) and two with a larger choice set (*Asset Green* vs. *Asset Blue* vs. *No Investment*) for $\mu_B \in \{-10, -20\}$. The order of decisions was randomized at the subject level.

As in the main experiment, to explain the drift of an ABM to the subjects, they had to successively draw three sample paths from the underlying process, and then saw an overview of five additional sample paths of this process before making a decision (see Figure 17 for examples of the latter with and without a decoy). The sample paths were randomly drawn at the subject level; that is, different subjects saw different sample paths of the same underlying process.

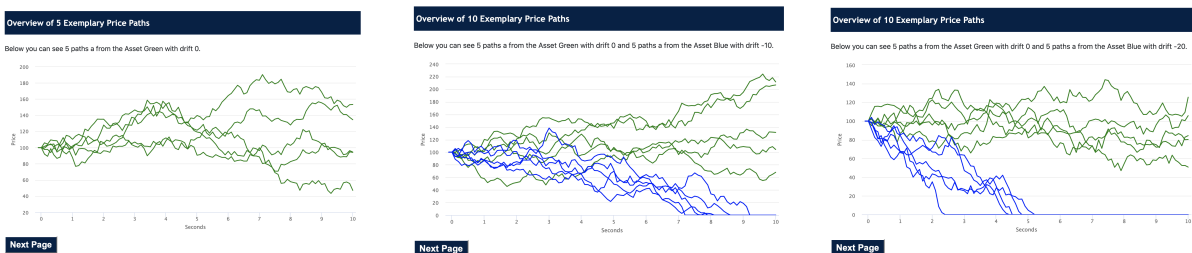


Figure 17: Screenshots of the sampling screens with and without a decoy.

At the end of the experiment, for each subject, one of the three investment decisions was

³³The experimental design, including the fully specified salience model and its predictions, was pre-registered in the AEA RCT registry as trial AEARCTR-0007398. Moreover, the experiment is available via the following link: <https://os-experiment-archive.herokuapp.com/demo>.

randomly drawn by the computer to be payoff-relevant. All subjects received an additional £3 for their participation in the experiment.

We conducted 7 sessions in March 2021 via the Oxford Nuffield CESS lab, using the software oTree (Chen *et al.*, 2016). A total number of $n = 247$ subjects completed the experiment.³⁴ The experiment lasted for around 11 minutes on average. Subjects earned £4.68 on average, with earnings ranging from £4.50 to £6.17.

Experimental results. The share of subjects investing in the viable Asset *Green* increases by roughly 10 p.p. when a dominated decoy is available (see Figure 18). As pre-registered, when a decoy was available, we drop the subjects who chose the dominated option — roughly, 4% of the subjects for $\mu_B = -10$ and 2% of the subjects for $\mu_B = -20$ — before calculating the share of subjects investing in the viable asset. This entails the implicit assumption that choosing the decoy is a mistake that is independent of whether a subject actually wants to invest or not. It is important to note, however, that the observed decoy effect is robust (but slightly smaller) to dealing with subjects who chose the decoy in a different way (e.g. assuming that all of them actually prefer not to invest).³⁵ This finding is consistent with salience theory, but it is clearly inconsistent with models of context-independent behavior such as EUT or CPT.

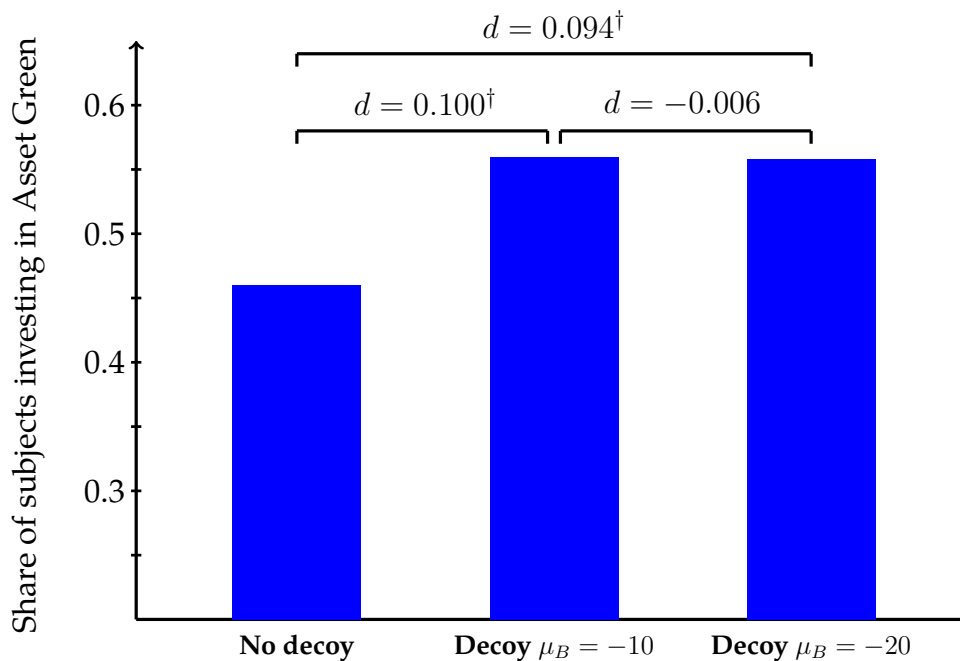


Figure 18: The figure depicts the share of subjects that invested in Asset *Green* with and without a decoy. As pre-registered, when a decoy was available, we drop the subjects who chose the dominated option (roughly, 4% for $\mu_B = -10$ and 2% for $\mu_B = -20$) before calculating these shares. We further present the results of *t*-tests with standard errors being clustered at the subject level. †: *p*-value = 0.013.

³⁴All of these subjects passed an attention check that tested for comprehension of the experimental instructions.

³⁵The most conservative way to test for decoy effects is to assume that all subjects who chose the decoy actually prefer not to invest. When making this assumption, the decoy effect ranges from 7.7 p.p. (if $\mu_B = -10$) to 8.5 p.p. (if $\mu_B = -20$). Using *t*-tests with standard errors being clustered at the subject level, we find that the latter effect is statistically significant at the 5%-level (*p*-value = 0.027), while the former effect almost is (*p*-value = 0.053).

We do not find a significant effect of the dominated decoy’s drift on how attractive Asset *Green* appears to be. At the same time, we cannot reject the null-hypothesis that the share of subjects choosing Asset *Green* is (weakly) larger for $\mu_B = -20$ than for $\mu_B = -10$. Hence, the observed behavior is indeed consistent with salience theory.

G.2: A More General Salience Model

Consider a choice set $\mathcal{C} = \{X_i\}_{i=1}^n$. The random variables (or *lotteries*) X_1 to X_n are non-negative with a joint cumulative distribution function $F : \mathbb{R}_{\geq 0}^n \rightarrow [0, 1]$. A state of the world refers to a tuple of outcomes, $(x_1, \dots, x_n) \in \mathbb{R}_{\geq 0}^n$. A salient thinker compares the value of a given lottery, $v(X_i)$, to a *reference point* $R_i = \phi(v(X_1), \dots, v(X_{i-1}), v(X_{i+1}), \dots, v(X_n))$. Bordalo *et al.* (2012) assume that the reference point is given by the state-wise average over all alternative options: $R_i = \frac{1}{n-1} \sum_{j \neq i} v(X_j)$. We, in contrast, allow for a more general reference point $\phi : \mathbb{R}^{n-1} \rightarrow \mathbb{R}$ that (i) strictly increases in each of its arguments and (ii) satisfies $\phi(z, \dots, z) = z$.

Definition 7. *The salience-weighted utility of lottery X_i evaluated in $\mathcal{C} = \{X_j\}_{j=1}^n$ equals*

$$U^s(X|\mathcal{C}) = \frac{1}{\int_{\mathbb{R}_{\geq 0}^n \sigma(v(x_i), r_i) dF(x_1, \dots, x_n)} \int_{\mathbb{R}_{\geq 0}^n v(x_i) \cdot \sigma(v(x_i), r_i) dF(x_1, \dots, x_n),$$

where $r_i = \phi(v(x_1), \dots, v(x_{i-1}), v(x_{i+1}), \dots, v(x_n))$, and where $\sigma : \mathbb{R}_{\geq 0}^2 \rightarrow \mathbb{R}_{>0}$ is a salience function that is bounded away from zero.

G.3: Salience Predictions on Decoy Effects

Consider the choice set $\mathcal{C} = \{X_{\tau_{a,b}}, Y_{\tau_{a,b}}, z\}$. Here, the reference point R_k relative to which Asset $k \in \{G, B\}$ is evaluated has the following distribution:

	$p_G p_B$	$p_G(1 - p_B)$	$(1 - p_G)p_B$	$(1 - p_G)(1 - p_B)$
R_G	$\phi(v(z), v(a))$	$\phi(v(z), v(b))$	$\phi(v(z), v(a))$	$\phi(v(z), v(b))$
R_B	$\phi(v(z), v(a))$	$\phi(v(z), v(a))$	$\phi(v(z), v(b))$	$\phi(v(z), v(b))$

Table 4: *Distribution of the reference points in the larger choice set.*

If the choice set includes also the dominated Asset *B* and if this dominated asset has a sufficiently negative drift, then — compared to the case with a binary choice set — Asset *G* becomes more attractive to a salient thinker. Hence, a sufficiently “bad” Asset *B* serves as a decoy that boosts demand for Asset *G*. More formally, we obtain the following proposition:

Proposition 7 (Salience-Driven Decoy Effect).

- (a) *The salient thinker will never invest in the dominated Asset B.*
- (b) *The salience-weighted utility derived from investing in Asset G monotonically increases in p_B .*
- (c) *There is some $\hat{\mu} \in \mathbb{R} \cup \{-\infty\}$, so that a salient thinker invests in Asset G if and only if $\mu_B < \hat{\mu}$.*

(d) If the salient thinker invests in Asset G when facing the binary choice set $\{X_{\tau_{a,b}}, z\}$, then $\hat{\mu} \in \mathbb{R}$.

Proof. Part (b). The salience-weighted utility from investing in Asset G is

$$U^s(X_{\tau_{a,b}}|\mathcal{C}) = v(a)\pi(p_G, p_B) + v(b)[1 - \pi(p_G, p_B)],$$

where

$$\pi(p_G, p_B) := \frac{p_G s_a(p_B)}{p_G s_a(p_B) + (1 - p_G) s_b(p_B)},$$

with $s_a(p_B) := p_B \sigma(v(a), \phi(v(z), v(a))) + (1 - p_B) \sigma(v(a), \phi(v(z), v(b)))$ being the average salience of a and $s_b(p_B) := p_B \sigma(v(b), \phi(v(z), v(a))) + (1 - p_B) \sigma(v(b), \phi(v(z), v(b)))$ being that of b .

Since $\phi(v(z), v(b)) > \phi(v(z), v(a))$ and, thus, $\sigma(v(a), \phi(v(z), v(b))) > \sigma(v(a), \phi(v(z), v(a)))$ by the ordering property, $s_a(p_B)$ is strictly decreasing in p_B . Analogously, ordering implies that $\sigma(v(b), \phi(v(z), v(a))) > \sigma(v(b), \phi(v(z), v(b)))$, so that $s_b(p_B)$ is strictly increasing in p_B . It follows that $\pi(p_G, p_B)$ is strictly decreasing and, thus, $U^s(X_{\tau_{a,b}}|\mathcal{C})$ is strictly increasing in p_B .

Part (a). Next, we observe that

$$\begin{aligned} \frac{\partial}{\partial p_G} \pi(p_G, p_B) &= \frac{s_a(p_B)[p_G s_a(p_B) + (1 - p_G) s_b(p_B)] - p_G s_a(p_B)[s_a(p_B) - s_b(p_B)]}{(p_G s_a(p_B) + (1 - p_G) s_b(p_B))^2} \\ &= \frac{s_a(p_B) s_b(p_B)}{(p_G s_a(p_B) + (1 - p_G) s_b(p_B))^2} > 0, \end{aligned}$$

which, in turn, implies that $U^s(X_{\tau_{a,b}}|\mathcal{C})$ is strictly decreasing in p_G .

Combining this with Part (b) and the fact that $p_B > p_G$, we conclude:

$$\begin{aligned} U^s(X_{\tau_{a,b}}|\mathcal{C}) &= v(a)\pi(p_G, p_B) + v(b)[1 - \pi(p_G, p_B)] \\ &> v(a)\pi(p_G, p_G) + v(b)[1 - \pi(p_G, p_G)] \\ &> v(a)\pi(p_B, p_G) + v(b)[1 - \pi(p_B, p_G)] = U^s(Y_{\tau_{a,b}}|\mathcal{C}). \end{aligned}$$

Part (c). Follows immediately from the fact that $U^s(X_{\tau_{a,b}}|\mathcal{C})$ is increasing in p_B .

Part (d). Given the binary choice set $\mathcal{C}' = \{X_{\tau_{a,b}}, z\}$, the salience-weighted utility from investing in Asset G is given by $U^s(X_{\tau_{a,b}}|\mathcal{C}') = v(a)\tilde{\pi} + v(b)[1 - \tilde{\pi}]$, where

$$\tilde{\pi} := \frac{p_G \sigma(v(a), v(z))}{p_G \sigma(v(a), v(z)) + (1 - p_G) \sigma(v(b), v(z))}.$$

Notice that

$$\lim_{p_B \rightarrow 1} \pi(p_G, p_B) = \frac{p_G \sigma(v(a), \phi(v(z), v(a)))}{p_G \sigma(v(a), \phi(v(z), v(a))) + (1 - p_G) \sigma(v(b), \phi(v(z), v(a)))} < \tilde{\pi},$$

since, by ordering, $\sigma(v(a), \phi(v(z), v(a))) < \sigma(v(a), v(z))$ and $\sigma(v(b), \phi(v(z), v(a))) > \sigma(v(b), v(z))$. Hence, if $U^s(X_{\tau_{a,b}}|\mathcal{C}') > 0$, then also $\lim_{p_B \rightarrow 1} U^s(X_{\tau_{a,b}}|\mathcal{C}) > 0$. \square

G.4: Experimental Instructions

Screen 1 — Instructions: Overview of the Experiment

In this experiment you will make investment decisions. More precisely, in each decision, you will have to decide whether to invest in an asset and, if so, in which one. In total you will make 3 investment decisions.

At the end of the experiment, we will choose one of your 3 decisions at random to be payoff-relevant. Your payment depends on whether you invested in an asset and, if so, at which price the asset is sold. During the experiment, we will denote all monetary values in the currency Taler which will be converted to £ at an exchange rate of £1 = 60 Taler.

You will receive a show up fee of £3 for your participation in the experiment.

Screen 2 — Instructions: Development of the Asset Price

Below you can see a graph, which depicts the development of the price of an asset, for a period of 10 seconds. As soon as you press "Start", a line which represents the value of the asset will appear. Please press "Start" now.

[Subjects are shown a graph of an exemplary price path with a final price of 100 Taler]

Screen 3 — Instructions: Different Drifts

The assets that you can choose in this experiment differ only in their long-term profitability. How profitable an asset is in the long run is described by the drift of the asset. The drift denotes the average change in the asset's value per second.

A positive drift implies that the asset will increase in value in the long run, while a negative drift implies that the asset will decrease in value in the long run. A drift of zero implies that the asset's value will neither increase nor decrease in the long run.

Notice, though, that even an asset with a negative drift sometimes increases in value, and that – independent of its drift – the value of an asset can, in principle, become arbitrarily large. The probability that an asset's value indeed becomes very large is the smaller the more negative the drift is.

To give you a feeling of how the long-run value of an asset changes with its drift we will show you now a few examples of assets with different drifts.

Screens 4-6 — Instructions: The Drift of an Asset

The drift of this asset is 0 [or 2 or -2]. Please press "Start" and watch the development of the asset's price.

[Subjects are shown a graph of an exemplary price path with a final price of 100 Taler for the process with drift 0, 120 Taler for the process with drift 2, and 80 for the process with drift -2.]

Screen 7 — Instructions: Upper and Lower Bounds on the Asset's Price

If you decide to invest in an asset, it will be sold at pre-specified prices: either at 90 (lower bound) or at 190 (upper bound). All available assets will be automatically sold as soon as their prices hit one of these bounds. If the price of an asset reaches the upper bound of 190, the asset will be sold and you will receive 190 Taler. If the price reaches the lower bound of 90, the asset will be sold as well and you will receive 90 Taler.

Importantly, there is no expiration time. The price of an asset will change until it hits one of the two bounds.

[Subjects are shown a graph with a button labelled "Asset Green". When they click the button, a price path appears]

Screen 8 — Instructions: Upper and Lower Bounds on the Asset's Price

You have to decide repeatedly whether to invest in one of at most two different assets. You can always opt for the alternative of "no investment" in which case you will receive 100 Taler (i.e. the starting value of each asset) with certainty.

In the example below you can choose among the following three options:

- No investment
- Invest in Asset Green (which has drift 0)
- Invest in Asset Blue (which has drift -2)

Please make a decision now.

[Subjects are shown a graph with three buttons labelled "Asset Green", "Asset Blue" and "Sell Immediately"]

Screen 9 — Questions about the Instructions

[Subjects see a graph with an upper bound of 190, a lower bound of 90 and a starting value of 100]

[Subjects need to answer the following questions]

- Suppose you choose "No Investment". How many Taler do you get?
- What is the starting value of the asset?
- Suppose you have invested in the asset and it hits the lower bound. How many Taler do you get?
- If an asset has a drift of 2 it will:

[In the last question they choose between the following options]

- increase by an amount of 2 every second
- increase by an amount of 2 per second on average over a long time period
- vary by an amount of 2 every second

Screen 10 — Decision round x of 3

In the upcoming decision you will be able to choose between the following options:

- No investment. (100 Taler with certainty)
- An asset with a drift of 0. (Green Asset)
- An asset with a drift of [-10 / -20] (Blue Asset)

Before you make your decision we will show you a few exemplary price paths of the asset.

Screen 11 — Sampling page x of 3

Below you can see exemplary price paths of Asset Green (with drift 0), for a period of ten seconds.

[Subjects see one individual price path from Asset Green. If there is also an Asset Blue, they see a price path from Asset Blue as well. This screen is shown 3 times in a row]

Screen 12 — Overview of 5 Exemplary Price Paths

Below you can see 5 paths from the Asset Green with drift 0.

[Subjects see 5 price paths from Asset Green. If there is also an Asset Blue, they see 5 price paths from Asset Blue as well.]

Screen 13 — Decision x of 3

Please choose between Asset Green (with drift 0) [and Asset Green (with drift -10/-20)] and "No Investment".

[Subjects are shown a graph with three buttons labelled "Asset Green", "Asset Blue" and "Sell Immediately"]

Screen 14 — Payoff

Your decision from round [a] will be paid.

You sold the asset for [b] Taler.

You earned a bonus payment of £[c] on top of your show up fee of £3. Your total payment is £[d].