Optimal Stopping in a Dynamic Salience Model

Markus Dertwinkel-Kalt, Jonas Frey, Mats Köster

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

Optimal Stopping in a Dynamic Saliency Model

Abstract

While many puzzles in static choices under risk can be explained by a preference for positive and an aversion toward negative skewness, little is known about the implications of such skewness preferences for decision making in dynamic problems. Indeed, skewness preferences might play an even bigger role in dynamic environments because, even if the underlying stochastic process is symmetric, the agent can endogenously create a skewed distribution of returns through the choice of her stopping strategy. Guided by saliency theory, we theoretically and experimentally analyze the implications of skewness preferences for optimal stopping problems. We find strong support for all saliency-based predictions in a laboratory experiment, and we verify that saliency theory coherently explains skewness preferences revealed in static and dynamic decisions. Based on these findings we conclude that the static saliency model—unlike (static) cumulative prospect theory—can be reasonably applied to dynamic decision problems. Our results have important implications for common optimal stopping problems such as when to sell an asset, when to stop gambling, when to enter the job market or to retire, and when to stop searching for a house or a spouse.

JEL-Codes: D010, D810, D900.

Keywords: saliency theory, prospect theory, skewness preferences, behavioural stopping.

Markus Dertwinkel-Kalt
Frankfurt School of Finance & Management
Frankfurt / Germany
m.dertwinkel-kalt@fs.de

Jonas Frey
Saïd Business School
University of Oxford / United Kingdom
Jonas.Frey.DPHIL2017@saïd.oxford.edu

Mats Köster
Düsseldorf Institute for Competition
Economics (DICE)
HHU Düsseldorf / Germany
koester@dice.hhu.de

August 2020

The experimental design was pre-registered in the AEA RCT registry as trial AEARCTR-0005359. We thank Sebastian Ebert, Katrin Gödker, Paul Heidhues, Alex Imas, Joshua Miller, Philipp Strack, and Peter Wakker for helpful comments and suggestions. We gratefully acknowledge financial support by the German Science Foundation (DFG projects 404416232 for Markus Dertwinkel-Kalt and 235577387/GRK1974 for Mats Köster). Part of the research was carried out while Mats Köster was visiting UC Berkeley and he is grateful for their hospitality.

1 Introduction

1.1 Overview and Contribution

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. While most puzzles regarding static choices under risk (e.g., the Allais paradoxes or the simultaneous demand for insurance and gambling) can be explained by a preference for positive and an aversion toward negative skewness, little is known about the implications of such *skewness preferences* for dynamic decision making. Indeed, skewness preferences might play an even bigger role in dynamic environments because, even if the underlying stochastic process is symmetric, the decision maker can create a skewed distribution of returns through the choice of her stopping strategy (see, e.g., Barberis, 2012; Ebert and Strack, 2015). By adapting a static model of skewness preferences to an otherwise standard optimal stopping problem, we study theoretically and experimentally the implications of *endogenously* created skewness for stopping behavior.

Skewness preferences revealed in static decisions are typically rationalized via models of non-linear probability weighting such as cumulative prospect theory (CPT; Tversky and Kahneman, 1992) or salience theory of choice under risk (ST; Bordalo *et al.*, 2012). Ebert and Strack (2015, 2018) show, however, that skewness preferences as modelled by CPT result in unrealistic, as too extreme, predictions for optimal stopping problems: as long as the available strategies enable the agent to create a sufficiently skewed return distribution, a naïve¹ CPT agent will never stop and instead gamble *until the bitter end*, independent of how unprofitable the gamble is in expectation. Salience theory, in contrast to CPT, imposes natural limits on the strength of skewness preferences in static settings (see Dertwinkel-Kalt and Köster, forthcoming), and, as we will show in this paper, therefore predicts also more sensible stopping behavior.

We study the problem of when to stop an Arithmetic Brownian Motion (ABM) with a non-positive drift and a finite expiration date. By allowing for a finite expiration date, we modify the setup in Ebert and Strack (2015, 2018) in a way that makes our predictions testable in an incentivized lab experiment. Here, a naïve CPT agent will, independent of the drift of the process, never stop gambling before bankruptcy, since at every point in time, he can find a *stop-loss and take-profit strategy*² that induces a sufficiently right-skewed distribution of returns to be more attractive than stopping immediately. Our main theoretical result shows that a naïve *salient thinker* does not exhibit the same extreme stopping behavior as a CPT agent, namely, a naïve salient thinker immediately stops any ABM with a drift below some threshold. As a corollary

---

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 CPT, 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, which is also supported by our experimental results.

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.*, 2020) 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.*, 2020). Moreover, these strategies allow agents to create skewness in the return distribution and they therefore allow us to study the effect of endogenously created skewness on stopping behavior.

to our main result, we further show that salience theory is consistent with (excessive) gambling for “fair” processes with a drift of zero, but does not predict that people will necessarily hold such a fair process until the expiration date. Our main theoretical result, thus, distinguishes salience theory not only from CPT-specifications adapted in the literature (as, e.g., by Barberis, 2012; Ebert and Strack, 2015; Heimer *et al.*, 2020), which postulate that subjects will never stop before the expiration date, but also from expected utility theory (EUT) with a concave utility function, according to which the agent stops any process with a non-positive drift immediately.

In a next step, we restrict the choice set to all stop-loss and take-profit strategies to learn more about how a naïve salient thinker plans to stop and how he will revise this plan over time. These additional predictions allow for a more powerful experimental test of salience theory against alternative models of skewness preferences. 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.*, 2020): 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. Second, a naïve salient thinker does not necessarily follow his initial plan, but may instead revise his strategy over time. Salience theory can rationalize, in particular, stopping behavior that is consistent with the well-known disposition effect (e.g., Shefrin and Statman, 1985; Odean, 1998; Barberis, 2012; Imas, 2016).

We test the predictions of salience theory in a lab experiment in which subjects have to decide when to stop ABMs with different non-positive drifts. Using a within-subjects design, we validate our approach of adapting the static salience model to a dynamic setting by further eliciting skewness preferences in static choices. For that, building on Dertwinkel-Kalt and Köster (forthcoming), we first 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. To test for this non-parametric salience prediction, we let subjects choose repeatedly between a binary lottery with a fixed expected value and a fixed variance, but varying skewness, and a safe payoff equal to the lottery’s expected value. If salience is the psychological mechanism driving skewness preferences in general, it should coherently explain revealed attitudes toward skewness in both the static and the dynamic choices.

Our experimental results strongly support all salience-based predictions on stopping behavior. First, most subjects start gambling the fair process with a drift of zero, but then stop before the expiration date. Second, subjects stop the earlier, the more negative is the drift of the process. Third, for the median subject, more than 70% of all chosen strategies are loss-exit strategies. Fourth, 93% of the subjects revise their initial strategy at least once, and actual behavior is reminiscent of the disposition effect. Finally, we find a positive correlation between static and dynamic gambling behavior, which is both, statistically and economically, significant: subjects that reveal a stronger preference for positive skewness in static choices also have a larger propensity to choose loss-exit strategies in the dynamic ones. In sum, not only do our results suggest that endogenously created skewness plays an important role for stopping behavior, but also that the static salience model can be reasonably applied to a dynamic context. We conclude that it is not the non-linear probability weighting per se, but the specificities of common CPT-specifications that produce implausible dynamic predictions (see also Duraj, 2020).

1.2 Related Literature

We contribute to four strands of the literature: to (i) the literature on skewness preferences (which has, so far, mostly focused on static choices), to (ii) the theoretical and (iii) the experimental literature on behavioral stopping as well as to (iv) the literature on the testable implications of salience theory of choice under risk.

The fact that people have a preference for right-skewed (or positively skewed) risks and an aversion toward left-skewed (or negatively skewed) risks is one of the central findings by Kahneman and Tversky (1979). Such *skewness preferences* allow us to explain, for instance, simultaneous purchases of insurance and lottery tickets or the Allais paradoxes (Dertwinkel-Kalt and Köster, forthcoming). In a laboratory experiment, Dertwinkel-Kalt and Köster (forthcoming) have established a causal link between the attractiveness of a binary lottery and its skewness: fixing both, the lottery’s expected value and its variance, they find that the share of subjects choosing the lottery over its expected value monotonically increases in the lottery’s skewness. Besides replicating their findings, we establish in this paper a positive relationship between static and dynamic skewness preferences.

Second, we contribute to the theoretical literature on behavioral stopping, which has so far mainly analyzed the implications of non-linear probability weighting for dynamic gambling, with a focus on the behavior predicted by CPT (e.g. 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-linearities in the probability weights imply time-inconsistent preferences (e.g., Machina, 1989), which makes the analysis of stopping behavior both interesting and challenging. Predicted behavior depends, in particular, on whether the agent is *naïve* about his time-inconsistency or not. 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 given strategy becomes important. The literature has studied the stopping behavior of naïve agents without commitment (e.g., Barberis, 2012; Ebert and Strack, 2015)³ as well as with partial or full commitment (e.g., Xu and Zhou, 2013; Henderson *et al.*, 2017; He *et al.*, 2019). Duraj (2020) discusses further models (such as disappointment aversion) that produce less extreme predictions than CPT, all of which we discuss in Section 6, when reviewing alternative models. We add to the theoretical literature on behavioral stopping by proposing a dynamic version of salience theory of choice under risk (Bordalo *et al.*, 2012), and we describe stopping behavior of naïve (and briefly that of sophisticated) salient thinkers without commitment.

Our main theoretical result contrasts the non-parametric predictions of salience theory on the stopping behavior of a naïve agent with those of CPT. For an infinite process, Ebert and Strack (2015) show that, independent of the drift of the process, a naïve CPT agent will never stop gambling before bankruptcy, since at every point in time, he can find a sufficiently right-

---

³Barberis (2012) investigates numerically how, in a series of symmetric gambles, suitable stopping strategies allow to create skewness in the return distribution, thereby inducing time-inconsistent gambling behavior of a naïve CPT-agent. As Barberis (2012) builds on a discrete-time setup with a short expiration date, close to which the agent runs out of skewness, he does not find the never-stopping result that Ebert and Strack (2015) derive.

skewed stop-loss and take-profit strategy that is more attractive than stopping immediately. In contrast, our salience model predicts that gambling behavior is sensitive to the drift of the process: a naïve salient thinker gambles if and only if the drift of the process is not too negative. We further show that these predictions carry over to the case of a finite expiration date, which makes our theory testable in the context of an incentivized lab experiment.

The reason why salience theory predicts more plausible stopping behavior than CPT lies in the different mechanism that induces a preference for right-skewed and an aversion toward left-skewed risks. Salience theory builds on the idea of the *contrast effect* (e.g., Schkade and Kahneman, 1998), which implies that states of the world in which the attainable outcomes differ a lot, attract a great deal of attention and are perceived as more likely than they actually are. Thus, in contrast to CPT, it is not the probability of a given state of the world that mechanically determines the size of the “probability distortion,” but the contrast in attainable payoffs in this state. As a consequence, an attractive stop-loss and take-profit strategy under salience theory is very different from an attractive stop-loss and take-profit strategy under CPT. While a CPT-agent is attracted by a stopping strategy with thresholds arbitrarily close to the current wealth level, according to salience theory the take-profit threshold has to be much larger than the current wealth level in order to attract a salient thinker’s attention. But, since salience distortions are bounded, a stopping strategy with a large take-profit threshold can only be attractive as long as the returns to gambling are not too negative. Thus, salience theory imposes natural bounds on the strength of skewness preferences, both in static and in dynamic setups.

Third, by testing our salience predictions in a lab experiment, we also contribute to the small, but growing experimental literature on behavioral stopping (e.g., Imas, 2016; Imas *et al.*, 2017; Strack and Viefers, forthcoming; Heimer *et al.*, 2020). Existing papers have focused on the question in how far stopping decisions are path-dependent, and in particular in how far the realization of previous gains and losses affects behavior. Closest related to our paper is Heimer *et al.* (2020), who study optimal stopping behavior with an underlying process consisting of repeated (fair) coin tosses. Similarly to us, they focus on stop-loss and take-profit strategies, 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, which is also reflected in our data. We regard the study by Heimer *et al.* (2020) as complementary to ours, as their study is purely experimental and, therefore, less suited to precisely test different models of skewness preferences against each other, and as it does not investigate drift-effects—the key feature of our paper—, which turn out to be of first-order importance.

More generally, we contribute to the theoretical and experimental literature on salience theory of choice under risk (Bordalo *et al.*, 2012), which has been axiomatized by Lanzani (2020). Köster (2020) extends the axioms introduced in Lanzani (2020) to capture the choice between multi-dimensional risks, thereby unifying the applications of salience theory to choice under risk and consumer choice in one general framework. The predictions of salience theory for static decisions under risk have been confirmed in several lab experiments. Frydman and Mormann (2018) show, for instance, that the salience-based explanation of the Allais paradoxes fits subjects’ behavior much better than the CPT-based explanation. We add to the theoretical literature on salience theory by extending the model from static to dynamic decision problems and

to the experimental literature by testing the model’s predictions in a controlled lab experiment.

Finally, basic stopping problems also play a crucial role in the classical literature, including the optimal stopping literature following Rust (1987) and the search literature in industrial organization (for reviews of the search literature see Armstrong, 2017; Anderson and Renault, 2018; Gonzales, 2018). These papers completely abstract from any non-linear probability weighting and skewness effects despite the fact that they have been shown to be important in static settings. Integrating models of non-linear probability weighting into these setups thus offers a novel behavioral perspective on the classical problems of optimal stopping.

2 A Dynamic Version of Saliency Theory of Choice under Risk

2.1 Static Model

Consider an agent who has to choose from a choice 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 $S \subseteq \mathbb{R}_{\geq 0}^2$. If a random variable is degenerate, we call it a *safe* option.

According to saliency 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 $s \in S$ that depends on the state’s objective probability and on its saliency. The saliency of a state is assessed by a so-called *saliency function*, which is defined as follows:

Definition 1 (Saliency Function). *We say that a symmetric, bounded, and twice differentiable function $\sigma : \mathbb{R}_{\geq 0}^2 \rightarrow \mathbb{R}_{> 0}$ is a saliency 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 $\epsilon > 0$, we have*

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

We say that a given state of the world $(x, y) \in S$ is the more salient the larger its saliency value is. The ordering property implies that a state 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 saliency 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 higher outcome levels, thereby qualifying the contrast effect.

---

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 $\sigma(-x, -y) > \sigma(-w, -z)$.

Intrinsically, a salient thinker is (weakly) risk-averse, as he evaluates monetary outcomes via some strictly increasing and (weakly) concave value function $v : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ that is twice differentiable. But, depending on the salience of outcomes, which, in turn, depends on the composition of the choice set $\mathcal{C} = \{X, Y\}$, a salient thinker might behave as if being risk-seeking. More specifically, a salient thinker maximizes a *salience-weighted utility* 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.⁵

Notice that the salience-weighted probabilities are normalized so that they sum to one (see, e.g., Bordalo *et al.*, 2012; Dertwinkel-Kalt and Köster, forthcoming). This implies, in particular, that 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.

2.2 Dynamic Model

Stochastic process. We build on the setup by Ebert and Strack (2015, 2018), who assume that the wealth of an agent, who steadily participates in a gamble, evolves according to a Markov diffusion. Specifically, we consider an *Arithmetic Brownian Motion* (ABM),

$$dX_t = \mu dt + \nu dW_t, \text{ with } X_0 = x,$$

with a constant drift $\mu \in \mathbb{R}$ and volatility $\nu \in \mathbb{R}_{> 0}$, and 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) in two ways: First, we assume that the process is non-negative, $X_t \geq 0$, and absorbing in zero. Second, we allow for a finite *expiration date* $T \in \mathbb{R}_{> 0} \cup \{\infty\}$. The latter is necessary,⁶ but substantially complicates the theoretical analysis of stopping behavior.

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 potentially random 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 $\tau \leq T$ then implements a random wealth level X_τ with a cumulative distribution function denoted by F_τ .

⁵Alternatively, we could assume that the arguments of the salience function are not utils, but the numerical outcomes of the random variables (see Dertwinkel-Kalt and Wenner, 2020, for a discussion). This would not change any of the results derived in this paper, but would simply make the exposition of the arguments a bit more clumsy.

⁶As we discuss in more detail in Section 4, the only alternative to such a deterministic termination rule would be to implement a random termination rule, according to which, at time t , the process is automatically stopped with some probability $\omega_t \in [0, 1]$. The problem with a random termination rule is, however, that the theoretical analysis of stopping behavior, under the assumption of non-linear probability weighting, becomes completely intractable.

For our main result, we do not impose any restrictions on the set of stopping times that the agent can choose from. In particular, we also allow for randomized strategies as analyzed in Henderson *et al.* (2017) or Duraj (2020). 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* $\tau_{a,b}$, which is defined as the first leaving time of the interval (a, b) and which implements a random wealth level $X_{T \wedge \tau_{a,b}}$. Notably, 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.*, 2020) 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.*, 2020).

Solution concept under salience theory. Non-linear probability weighting, whether it is endogenous as in salience theory or not, 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 thus 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 naïveté as sensible. But, in Section 4 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 immediately. 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 (Naïve Decision 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 he chooses the random variable X_τ from the set $\{X_\tau, x_t\}$. 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, meaning that he compares it only to the alternative of stopping immediately. Since salience theory is about context-dependent behavior, it is necessary to impose some assumption on the *consideration set*—that is, the set of strategies that the agent takes into account when making the stopping decision—to derive testable predictions. We regard this specification as plausible, especially in the context of our experiment, where subjects can observe only a single strategy at a time (see Section 4 for details). To tie our hands, we pre-registered our assumptions on the consideration set before running the experiment.

3 Stopping Behavior of a Naïve Salient Thinker

3.1 Motivating Example

As an illustration of the salience mechanism, consider an agent with a linear value function, $v(x) = x$, who decides when to stop a “fair” process with a drift of zero. Suppose that the agent

only considers stop-loss and take-profit strategies, which can be represented by a threshold stopping time $\tau_{a,b}$. For the sake of illustration, we abstract from a finite expiration date and assume $T = \infty$. Then, any threshold stopping time $\tau_{a,b}$ induces a binary lottery $X_{\tau_{a,b}} = (a, p; b, 1 - p)$, which has—as the process has a drift of zero—an expected value of $\mathbb{E}[X_{\tau_{a,b}}] = x_t$ at time t .

When does this salient thinker prefer to gamble according to the stopping time $\tau_{a,b}$ over stopping at a current wealth level of x_t ? By construction, if the upside and the downside of the binary lottery $X_{\tau_{a,b}}$ are equally salient, then a salient thinker with a linear value function behaves as if he was risk-neutral; that is, he is indifferent between the lottery $X_{\tau_{a,b}}$ and the safe option x_t . Precisely, the salient thinker strictly prefers the binary lottery $X_{\tau_{a,b}}$ over the current wealth level x_t if and only if the lottery's upside, b , is more salient than its downside, a ; that is, if and only if $\sigma(b, x_t) > \sigma(a, x_t)$ holds. Since $\sigma(b, x_t) > \sigma(x_t, x_t)$ due to ordering, and since the salience function is continuous, we can find, for any x_t , stopping thresholds a and b , such that a salient thinker with a linear value function strictly prefers the binary lottery $X_{\tau_{a,b}}$ over the safe option x_t . Hence, at any positive level of wealth, the naïve salient thinker finds some stopping time $\tau_{a,b}$ that he strictly prefers to stopping immediately and he, therefore, never stops a process with zero drift. It is easily verified that the result still holds for a finite expiration date. This and all the other 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.*

Outlook. We have seen that a naïve salient thinker with a linear value function never stops a fair process, since at any point in time he finds a threshold stopping time that induces a binary lottery with a salient upside, which is more attractive than stopping immediately. But how general is this result? And how sensitive is a salient thinker's behavior to the drift of the process?

As we will work out in the next subsection, there are two reasons for why a salient thinker would stop the process before the expiration date: either the drift of the process is sufficiently negative or the salient thinker is intrinsically risk-averse (i.e., he has a sufficiently concave value function). More precisely, while a naïve salient thinker with a linear value function holds also processes with a slightly negative drift until the expiration date, a salient thinker with a sufficiently concave value function stops even a fair process. But, irrespective of his intrinsic risk-aversion, any salient thinker immediately stops processes with a sufficiently negative drift, which distinguishes salience theory from models like CPT and constitutes our main result.

3.2 Main Theoretical Result

Motivated by the stark never-stopping result, which Ebert and Strack (2015) have derived for naïve agents with CPT-preferences, we ask whether the stopping behavior of a naïve salient thinker is sensitive to the drift of the process. The prediction that a naïve CPT-agent never stops, irrespective of the drift of the process, relies on the fact that CPT induces a strong preference for right-skewed lotteries in static settings (Ebert and Strack, 2015, Theorem 1). Salience theory, in contrast, imposes limits on the strength of skewness preferences in static choices (Dertwinkel-Kalt and Köster, forthcoming, Corollary 2): the premium that a salient thinker is willing to pay

for taking up a right-skewed binary risk approaches zero, as the lottery's skewness becomes arbitrarily large. In other words, more skewness is not necessarily attractive to a salient thinker. As a consequence, salience theory implies less extreme behavior in optimal stopping problems: precisely, as we will show in the following, *any* naïve salient thinker immediately stops any process with a sufficiently negative drift.

Consider a general stopping time $\tau \leq T$, which induces a random wealth level X_τ with a cumulative distribution function denoted by F_τ . A naïve salient thinker then stops immediately at time $t = 0$ if and only if, for any such stopping time τ , we have

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

For an initial value $X_0 = x$, we define an auxiliary function $\tilde{u}(z) := (v(z) - v(x)) \sigma(v(z), v(x))$, which constitutes a classical utility function that is strictly increasing and continuous 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 a fixed initial value $X_0 = x$, a naïve salient thinker stops the process immediately if and only if an EUT-agent with a utility function $\tilde{u}(\cdot)$ stops it immediately.

The main step, in proving that a salient thinker immediately stops any process with a sufficiently negative drift, is to derive a bound on how risk-seeking a salient thinker can ever be, which is in a sense reminiscent of the limit on skewness preferences derived in Dertwinkel-Kalt and Köster (forthcoming). Our first result thus approximates a salient thinker's willingness to take risk by that of an EUT-agent with an exponential utility. Roughly, an EUT-agent is the more risk-seeking at a given wealth level the more convex his utility function is at this point. Since the salience function is differentiable and bounded, the corresponding auxiliary utility function $\tilde{u}(\cdot)$ is of *exponential growth* at $z = x$, which means that, in this point, it is not more convex than *all* exponential functions. Precisely, there exists an EUT-agent with an exponential utility who, because of his more convex utility function, takes up some risks that a salient thinker would certainly avoid, thereby imposing a bound on the salient thinker's willingness to take risk.

Lemma 1. *The auxiliary utility function $\tilde{u} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ is of exponential growth at $z = x$.*

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 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 immediately stops any process with a drift $\mu < \tilde{\mu}$.*

Theorem 1, which constitutes our main theoretical result, provides a clear-cut and testable distinction between our salience model and those models of exogenous probability-weighting, such as CPT, which yield the stark never-stopping result. Any naïve salient thinker, irrespective of his intrinsic risk-aversion, immediately stops any process with a sufficiently negative drift.

As a corollary to our main result, we observe that an intrinsically risk-averse salient thinker does not necessarily hold even a fair process until the expiration date. Depending on how concave a salient thinker's value function is, he stops a process with a drift of zero earlier or later. While a naïve salient thinker with a linear value function never stops a fair process, a naïve salient thinker with a sufficiently concave value function stops it immediately. A naïve salient thinker with a mildly concave value function, in contrast, starts to gamble fair processes with a drift of zero, but might then stop before reaching the expiration date.

Corollary 1. *A naïve salient thinker might stop a process with zero drift before the expiration date.*

3.3 Additional Results on Stop-Loss and Take-Profit Strategies

Next, in order to learn more about the stopping behavior of a naïve salient thinker, we restrict the agent's choice set to all stop-loss and take-profit strategies. First, we derive a stronger result on the limits of naïve gambling, which states that stopping behavior is monotonic in the drift of the process. Second, we characterize the type of stop-loss and take-profit strategies that are attractive to a salient thinker. Finally, 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 (see, e.g., Shefrin and Statman, 1985; Odean, 1998; Weber and Camerer, 1998; Imas, 2016; Heimer *et al.*, 2020, for empirical and experimental evidence).

A stronger result on the limits of naïve gambling. If we restrict the agent's choice set to all stop-loss and take-profit strategies and assume that the expiration date is sufficiently large, our model predicts that the agent's stopping behavior depends on the drift in a monotonic way.

To fix ideas, let us get back to the example of no expiration date. In this case, any stop-loss and take-profit strategy is associated with a threshold stopping time $\tau_{a,b}$ and therefore induces a binary lottery $X_{\tau_{a,b}} = (a, p; b, 1 - p)$ over wealth. For any such stopping time, 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 over wealth induced by the stopping time $\tau_{a,b}$ in terms of first-order stochastic dominance. By Proposition 1 in Dertwinkel-Kalt and Köster (forthcoming), 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 stopping immediately for any drift $\mu > \mu'$. In sum, a naïve salient thinker stops immediately if and *only if* the drift falls below some threshold.

What happens if we allow for a finite expiration date instead? Still, for any threshold stopping time $\tau_{a,b}$, an increase in the drift of the process makes it less likely to stop a and more likely to stop at b . But it is not clear, in general, that the distribution of the corresponding random variable, $X_{T \wedge \tau_{a,b}}$, improves in terms of first-order stochastic dominance. Thus, we cannot simply invoke Proposition 1 in Dertwinkel-Kalt and Köster (forthcoming) to establish a monotonic relationship between the drift of the process and a salient thinker's stopping behavior. The following proposition provides a *sufficient* condition for such a monotonic relationship to exist:

Proposition 2. *Suppose that the agent can only choose stop-loss and take-profit strategies, and fix an initial wealth level $x \in \mathbb{R}_{>0}$ as well as a volatility $\nu \in \mathbb{R}_{>0}$. For any two drifts μ and μ' with $0 \geq \mu' > \mu$, there exists some threshold value $\hat{T} \in \mathbb{R}_{\geq 0}$ such that, for any expiration date $T > \hat{T}$, a naïve salient thinker immediately stops the process with drift μ if he does so for the process with drift μ' .*

While a large expiration date is sufficient for a salient thinker to react to changes in the drift in a monotonic way, it is by no means necessary for this result to hold. For a rather short expiration date, it is simply the case that, for *some* threshold stopping times, $\tau_{a,b}$, the CDF of the corresponding random variable, $X_{T \wedge \tau_{a,b}}$, might not shift downwards in *all*, but only in some, points as the drift increases. The fact that there might exist stopping times, for which an increase in the drift does not improve the corresponding distribution in terms of first-order stochastic dominance, makes it hard to derive general results. But, using the functional forms typically assumed in the literature, it is easily verified that the sufficient condition derived in Proposition 2 is very weak in the sense that $\hat{T}$ is (close to) zero in most cases.

The role of skewness in naïve gambling. By the contrast effect, a salient thinker may choose a positively skewed risk over its expected value, but certainly avoids any negatively skewed risk (see Appendix D for a formal analysis of static skewness preferences). A preference for positive skewness is also what drives a salient thinker's stopping behavior. To make this point precisely, we introduce the notion of a *loss-exit* strategy, which not only induces a right-skewed distribution of returns for processes with a non-positive drift, but also looks right-skewed.

Definition 4 (Loss-Exit Strategy). *A stop-loss and take-profit strategy is a loss-exit strategy at time t if and only if the thresholds of the corresponding stopping time $\tau_{a,b}$ satisfy $b - x_t > x_t - a$.⁷*

Ordering and diminishing sensitivity together imply that a salient thinker gambles according to a stop-loss and take-profit strategy only if it is a loss-exit strategy. Consider again the case without an expiration date and assume a zero drift. If $b - x_t \leq x_t - a$, then not only is the payoff level lower in the state (a, x_t) compared to the state (b, x_t) , but also the contrast in outcomes is weakly larger in the former state. Hence, as the value function is weakly concave, the downside a is more salient than the upside b , which makes this stopping strategy unattractive to a salient thinker. Conversely, due to diminishing sensitivity of the salience function and concavity of the value function, $b - x_t > x_t - a$ does not imply that the upside b is more salient than the downside a , so that a salient thinker does not find every loss-exit strategy attractive. The argument extends to processes with a negative drift as well as to a setup with a finite expiration date.

Proposition 3. *A naïve salient thinker chooses a stop-loss and take-profit strategy only if it is a loss-exit strategy.*

---

⁷Not any stopping strategy that induces a right-skewed distribution over wealth is also a loss-exit strategy. Let $T = \infty$ and $\mu < 0$. Suppose that the current wealth level is x , and consider the stopping time $\tau_{a,b}$ with $a = x - \epsilon - \epsilon'$ and $b = x + \epsilon$ for $\epsilon, \epsilon' > 0$, which gives rise to a binary lottery $X_{\tau_{a,b}}$. As shown by Ebert (2015), a binary lottery is (unambiguously) right-skewed if and only if its lower payoff is strictly more likely than its higher payoff. Thus, for ϵ' sufficiently close to zero, the binary lottery $X_{\tau_{a,b}}$ is right-skewed, as $\mathbb{P}[X_{\tau_{a,b}} = a] > \frac{1}{2}$ due to the negative drift. But obviously the corresponding stopping strategy is not a loss-exit strategy. Only in the special case of zero drift, a stop-loss and take-profit strategy is a loss-exit strategy if and only if the corresponding binary lottery is right-skewed.

Can salience theory explain the disposition effect? Even if we restrict the agent to choose stop-loss and take-profit strategies, so that *planned* behavior is path-independent, salience theory can explain *actual* behavior consistent with the so-called disposition effect; i.e., the tendency to rather stop processes that have increased in value than those that have decreased in value (e.g., Odean, 1998; Heimer *et al.*, 2020). Our salience-based explanation of the disposition effect is similar in spirit to the one proposed by Barberis (2012) based on CPT: 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 why salience theory can explain disposition-effect-like behavior, it is again sufficient to think through the extreme case without an expiration date. At any point in time t , with a current wealth level x_t , the set of stop-loss and take-profit strategies can be represented by the set of threshold stopping times $\{\tau_{x_t-\epsilon, x_t+\epsilon'} : \epsilon, \epsilon' \geq 0\}$. A naïve salient thinker thus stops at time t if and only if, for any $\epsilon, \epsilon' > 0$, it holds that

$$p(v(x_t - \epsilon) - v(x_t))\sigma(v(x_t - \epsilon), v(x_t)) + (1 - p)(v(x_t + \epsilon') - v(x_t))\sigma(v(x_t + \epsilon'), v(x_t)) \leq 0,$$

where $p = p(\epsilon, \epsilon', \mu)$ denotes the probability of stopping at the stop-loss threshold. We can re-write the preceding inequality as follows,

$$\frac{\sigma(v(x_t - \epsilon), v(x_t))}{\sigma(v(x_t + \epsilon'), v(x_t))} \times \frac{v(x_t) - v(x_t - \epsilon)}{v(x_t + \epsilon') - v(x_t)} \geq \frac{1 - p}{p}. \quad (2)$$

Studying this inequality shows us that, depending on the exact functional forms of salience and value function, salience theory can indeed rationalize disposition-effect-like behavior. The right-hand side of (2) is independent of the current wealth level x_t due to the assumption of a constant drift (see Lemma 2 in Appendix A). If the left-hand side of (2) is increasing in x_t , the salience model thus predicts a *disposition effect*, because in this case stopping becomes more likely after the process has increased in value and less likely after a decrease in value. If the left-hand side of (2) 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.

4 An Experiment on Dynamic Gambling Behavior

In this section we present and discuss our experimental design. You can also click through the translated experiment via the following link: <https://stopping.herokuapp.com/demo/>.

4.1 Experimental Design

We conducted a pre-registered lab experiment in which subjects had to repeatedly decide at which price to sell an asset. 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. The selling price followed an Arithmetic Brownian Motion (ABM) with a drift parameter $\mu \in \{0, -0.1, -0.3, -0.5, -1, -2\}$ and a volatility $\nu = 5$. The process was updated every tenth of

a second, which implies $T = 100$.⁸ The asset's initial price was $x = 100$ Taler, an experimental currency that was converted into Euro at a ratio of 10:1 at the end of the experiment.

Figure 1: Screenshots of the decision screen for the process with zero drift (in German). 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.

We restricted the choice set at any given point in time to all stop-loss and take-profit strategies, as it is illustrated in Figure 1: 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. In addition, 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. Before starting the process, subjects could further decide to sell the asset immediately (see the upper left panel). At the beginning, the upper and lower threshold were centered symmetrically around the initial price (see also the upper left panel). In order to start the process, subjects had to move each threshold at least once.

Overall, subjects made six selling decisions, one decision for each of the drift parameters. The order of drifts was randomized at the subject level. To explain the drift of an ABM to the subjects, they were subsequently shown three sample paths from the underlying process and they observed an overview of ten additional sample paths of this process before making

⁸Notice 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.

a decision (see Figure 2). The sample paths were randomly drawn at the subject level; that is, different subjects saw different sample paths of the same underlying process.

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

After completing the six stopping problems, subjects faced a series of static choices under risk. Precisely, subjects made twelve choices between a binary lottery and a safe option paying the lottery’s expected value. For an overview of the lotteries, which are exactly those of Experiment 1 in Dertwinkel-Kalt and Köster (forthcoming), see Table 1 in Appendix C. Specifically, we used two sets of lotteries with the same expected value and the same variance, but different levels of skewness. Half of the lotteries have an expected value of $E = 30$ Euro and the other half have an expected value of $E = 50$ Euro. 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, one of the six selling decisions was randomly drawn by the computer to be payoff-relevant. We further randomly selected one subject in each session for whom, in addition, one of the twelve binary choices was randomly chosen to be payoff-relevant. Subjects were further rewarded for correctly answered CRT and financial literacy questions (1 Taler per question), and they received an additional 4 Euros for their participation.

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

4.2 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 for 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.⁹ Subjects were told that the processes they would see in the experiment differ *only* in their drift. In particular, we told subjects that all processes have in common that they are non-negative:

“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.”

This statement was accompanied by a graph illustrating that the process is non-negative and absorbing in zero (see Figure 3).

Figure 3: The figure illustrates that the process cannot take negative values (in German).

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.”

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.

⁹The sample paths 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.

Features of the process. To make our theory testable, we deviate from the setup proposed by Ebert and Strack (2015) in two ways: First, we introduce a finite expiration date. Second, we restrict ourselves to non-negative processes. In addition, we are only interested in processes with a non-positive drift, as processes with a positive drift do not allow us to separate between the predictions of different models such as EUT, CPT, and salience theory.

Since it is impossible to implement in practice a process that can run forever with probability one, we implemented—similar as Heimer *et al.* (2020)—a finite expiration date. Alternatively, we could have implemented a random termination rule, according to which, at any point in 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.¹⁰ In particular, when using a random termination rule, the stop-loss and take-profit strategies lose their intuitive properties regarding the induced skewness, because the thresholds themselves affect the shape of the probability distribution much less than in the case of a finite expiration date. Hence, a random termination rule is less suited to study the role of skewness for dynamic gambling behavior. Finally, a finite expiration date is easier to explain to subjects, which we regard—given the complexity of the experiment—as a major advantage.

To ensure incentive-compatibility, we restrict the processes to be non-negative, as we cannot credibly implement losses in a laboratory experiment. In 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. Moreover, given the framing of the experiment (as selling an asset at a certain price), it feels more natural to restrict the processes to take only non-negative values (or prices).

Restriction of the choice set. We restricted the subjects' choice set at any point in time to all stop-loss and take-profit strategies, which can be represented by a stop-loss and a take-profit threshold. This design choice was made based on both practical and theoretical considerations.

First, since we are interested in the role of endogenously created skewness for stopping behavior, we need an experimental design that allows us to learn something about the actual strategies that subjects choose. So, while we could have allowed subjects to adopt completely general strategies by just providing them with a STOP-button, this way we would not have learnt anything about their stopping strategies and, consequently, about the role of endogenous skew-

¹⁰Suppose, for instance, that we impose a random termination rule with $\omega_t = \omega \in (0, 1)$ for any $t \in \mathbb{N}$ and $\omega_t = 0$ otherwise. Then, the CDF induced by the threshold stopping time $\tau_{a,b}$ is given by

$$\mathbb{P}_0^\omega[X_{\tau_{a,b}} \leq z] = \begin{cases} \sum_{k=1}^{\infty} (1 - \omega)^{k-1} \mathbb{P}_0[\tau_{a,b} \leq k] \mathbb{P}_0[X_{\tau_{a,b}} = a] + (1 - \omega)^{k-1} \omega \mathbb{P}_0[\tau_{a,b} > k] \mathbb{P}_0[X_k \leq z | \tau_{a,b} > k] & \text{if } z < b, \\ 1 & \text{if } z = b, \end{cases}$$

where $\mathbb{P}_0[X_{\tau_{a,b}} = a]$ is defined in Lemma 2 and where $\mathbb{P}_0[\tau_{a,b} \leq k]$ as well as $\mathbb{P}_0[X_k \leq z | \tau_{a,b} > k]$ are defined in Lemma 3, both of which can be found in Appendix A. To put this into perspective, notice that, under the assumption of a finite expiration date $T < \infty$, the CDF induced by the same stopping time $\tau_{a,b}$ equals

$$\mathbb{P}_0[X_{T \wedge \tau_{a,b}} \leq z] = \begin{cases} \mathbb{P}_0[\tau_{a,b} \leq T] \cdot \mathbb{P}_0[X_{\tau_{a,b}} = a] + \mathbb{P}_0[\tau_{a,b} > T] \cdot \mathbb{P}_0[X_T \leq z | \tau_{a,b} > T] & \text{if } z < b, \\ 1 & \text{if } z = b. \end{cases}$$

ness. Obviously, we could have restricted the choice set in a different way, for instance, by asking subjects to simply state for how long they would like to hold the asset. While these deterministic stopping strategies could be easily elicited as well, they do not allow subjects to systematically add skewness to the return distribution, which makes them ill-suited to study the role of endogenous skewness for stopping behavior.¹¹ Stop-loss and take-profit strategies, in contrast, are not only easy to elicit, but also allow subjects to create a substantial amount of skewness in returns via the choice of the thresholds, thereby meeting all necessary criteria to address our research question. Second, as demonstrated in contemporaneous work by Heimer *et al.* (2020), stop-loss and take-profit strategies are highly relevant in practice, which is further 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).

A potential concern that remains to be addressed is whether we succeeded in restricting the choice set to all stop-loss and take-profit strategies. In other words, what we assume is that at any point in time, a subject sets a single stop-loss threshold and a single take-profit threshold, and believes to only stop at either of these two thresholds or at the expiration date. We 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. While, in principle, a subject could, already at time $t = 0$, come up with the plan to adjust the initial strategy later on, the red lines indicating the upper and lower threshold stick out on the screen, making this stop-loss and take-profit strategy particularly salient. Most importantly, even if subjects adopted strategies other than stop-loss and take-profit 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.

Non-binding strategies and costless adjustments. We allowed the subjects to costlessly adjust the stop-loss and take-profit thresholds at any point in time during a selling decision. Precisely, the 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 again and to continue the process. Again we made both design choices for practical and theoretical reasons.

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 could introduce partial commitment when subjects anticipate that they will not be able to adjust their strategy fast enough, namely, before the process hits either threshold. In any case, the ability to (partially) commit to a strategy would make the problem less interesting from a theoretical perspective. As argued in Henderson *et al.* (2017) and He *et al.* (2019), with partial commitment, the never-stopping result derived by Ebert and Strack (2015) for subjects with CPT-preferences breaks down. Thus, to experimentally differentiate between salience theory and other models of skewness preferences, such as CPT, we need to rule out

---

¹¹It is easily verified that, due to the lack skewness in the return distribution, *any* naïve salient thinker who is restricted to choose from the set of deterministic stopping times stops the process before the expiration date.

any commitment power (see Section 6 for a thorough discussion of alternative models), which is why we made the strategies non-binding and adjustments of the thresholds costless. Second, 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.* (2020) using brokerage data. Third, since subjects have a non-zero reaction time, non-binding strategies reduce the noise in measuring the *intended* stopping time. We regard this last argument as particularly important, as this noise would be asymmetric (making stopping disproportionately more likely than non-stopping) and hard to model.

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.

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.¹² 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 behaves as if she is 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, never gambles 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.

4.3 Experimental Predictions

We now state the precise predictions of salience theory that guided our experimental design. We thereby 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, the first prediction differs from what was pre-registered, while the remaining predictions are identical to the pre-registered ones. Notably, the last prediction on the relationship between static and dynamic skewness

---

¹²In contrast, under the random termination rule described in Footnote 10, any adjustment of the initial stop-loss and take-profit strategy could have been interpreted as time-inconsistent behavior. Beside the fact that the theoretical analysis under a random termination rule is intractable, we decided to implement a finite expiration date, as we are interested in the implications of skewness preferences for dynamic gambling behavior and not so much in time-inconsistency per-se.

preferences is motivated by our presumption that salience is the psychological mechanism that drives skewness preferences in general, that is, in both static and dynamic choices.

Our first prediction builds on the observation that, by Corollary 1, for the process with a zero drift our salience model is consistent with stopping immediately as well as with stopping before or at the expiration date. Investigating this prediction is still interesting, as it shows that salience theory is consistent with behavior ruled out by other models such as EUT or CPT.

Prediction 1. *If $\mu = 0$, subjects might start to gamble and they might stop before the expiration date.*

Our second prediction is derived from our main theoretical result, which allows us to clearly distinguish between salience theory and models, such as CPT, that yield the stark never-stopping result. More precisely, by Proposition 2, we expect a monotonic relationship between the drift of the process and a subject's stopping behavior.

Prediction 2. *The share of subjects selling the asset immediately monotonically decreases in the drift.*

Based on Proposition 3, we further predict that subjects choose loss-exit strategies.

Prediction 3. *Conditional on not selling the asset, subjects choose a loss-exit strategy.*

Since we extend a theory of static choice under risk to a dynamic setup, we are also interested in the empirical relationship between a subject's attitude toward static and dynamic risks. If salience is indeed the psychological mechanism that drives skewness preferences in general, it should coherently explain skewness preferences revealed in static and dynamic choices. To test this hypothesis, we classify both static and dynamic choices into being *skewness-seeking* or not. We say that one of the twelve static choices is skewness-seeking if the subject chooses a right-skewed lottery over a safe option or a safe option over a left-skewed or symmetric lottery (see Table 1 in Appendix C for the exact lotteries and the classification). 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. In particular, the stronger is the contrast effect (relative to the level effect) the closer this threshold is to zero, such that a larger share of skewness-seeking choices indicates that a subject is more susceptible to salience effects. In the dynamic decisions, we classify a chosen stopping strategy as being skewness-seeking if it is a loss-exit strategy and, therefore, induces a right-skewed return distribution. Based on Prediction 3, we then expect a positive correlation between the share of skewness-seeking choices in the 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.*

5 Experimental Results

We first present our main results that allow us to test Predictions 1-4. Subsequently, we present further results on strategy adjustments, and we discuss in how far they speak to the salience mechanism that drives our stopping predictions. Finally, we present additional, exploratory results on disposition-effect-like behavior, and on how cognitive skills affect stopping decisions.

5.1 Main Test of the Saliency Predictions

Our central experimental result—as depicted in Figure 4—addresses our main theoretical prediction that there should be a monotonic relationship between the drift of the process and a subject’s stopping behavior. At the hand of this figure we will successively discuss the results corresponding to Predictions 1 and 2.

Figure 4: The figure depicts the smoothed empirical cumulative distribution functions of stopping times.

To address Prediction 1, we look into stopping behavior for the fair process with a 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 (see Figure 4).

Result 1. *For the process with $\mu = 0$, the median subject holds the asset for 7.65 out of 10 seconds, and around 19% of the subjects hold the asset until the expiration date.*

While saliency theory does not make a precise prediction on when subjects stop a fair process (see Prediction 1), Result 1 is inconsistent both with EUT with a concave utility function—which predicts that subjects sell the asset immediately—as well as with CPT—which predicts that subjects will hold the asset until the expiration date.

Next, we investigate the effect of the drift on a subject’s decision whether to start a process and—if it is started—when to stop it. The share of subjects stopping immediately monotonically decreases in the drift of the process (see Figure 4 and the right panel of Figure 9 in Appendix E). This revealed sensitivity to the drift of the process is consistent with Prediction 2: regressing a binary indicator of whether a subject stops immediately or not on the drift of the process yields a point estimate of -0.171 (p -value < 0.001 , standard errors clustered at the subject level).

Result 2 (a). *The share of subjects selling immediately monotonically decreases in the drift of the process.*

Also this result is 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 stop 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.¹³

Result 2 (b). *Subjects stop earlier for processes with more negative drifts.*

We now turn to the subjects' strategy choices. Consistent with Prediction 3, a majority of subjects chooses loss-exit strategies initially, and this result on initial strategies holds across all the different drifts that we considered (see Figure 11 in Appendix E).

Result 3 (a). *Conditional on not selling immediately, 65% of initial strategies are loss-exit strategies.*

When we aggregate all the strategies that 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 5). This gives the second part of our third result:

Result 3 (b). *For the median subject, 73% of all strategies chosen throughout the experiment are loss-exit strategies.*

Figure 5: This figure gives the empirical distribution of the share of loss-exit strategies across subjects.

Finally, as depicted in Figure 6, subjects behave quite consistently in the static and the dynamic decision problems. To test for the link between static and dynamic skewness preferences,

¹³This 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.

we regress the share of loss-exit strategies amongst all strategies chosen throughout the six asset-selling tasks on the share of skewness-seeking choices in the 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 6: The figure depicts the relationship between static and dynamic gambling behavior. We provide the estimated slope-coefficient of the depicted linear regression, which is significantly larger than zero.

One might be concerned that Result 4 conflicts with the “discrepancy” between static and dynamic risk taking documented in Heimer *et al.* (2020). Heimer *et al.* (2020), 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 pretty right-skewed distributions of returns. In other words, while we study the relationship between skewness preferences in static and dynamic decisions, Heimer *et al.* (2020) look at the difference between static and dynamic problems that results from the fact that the latter enables subjects to endogenously create a skewed distribution of returns.

5.2 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 for at least one of the six processes, (ii) subjects adjust, on average, their strategy more than once per selling-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 bounds. Moreover, if a subject chooses a loss-exit strategy and the process hits a bound, 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 12 in Appendix E),¹⁴ which is consistent with Prediction 3.

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 indicates that subjects are time-inconsistent in their strategy choices, 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 main salience prediction that subjects are sensitive to the drift in a “reasonable” way.

Digging deeper into the results on stopping later than initially planned shows that 35% of the processes fall below the initial lower bound, but only 12% of the processes rise above the initial upper bound. Taken together these results indicate exactly the type of time-inconsistency 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 in order to create positive skewness in the return distribution, and they then adjust their strategy as soon as the loss-exit threshold is hit, thereby continuing to gamble 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.

5.3 Additional Findings

Subjects reveal a disposition effect. To test for disposition-effect-like behavior, it is insightful to split up the share of processes that are stopped later than initially planned into those that fall below the lower bound (namely, 35%) and those that rise above the upper bound (namely, 12%). Keeping the asset “too long” (compared to the subject’s initial strategy) when the process has decreased in value rather than when it has increased in value is reminiscent of the disposition effect, whereby processes are rather sold in the gain and rather held in the loss domain.

Another test for the disposition effect is to compare the tendency to sell 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. The subjects, who have revised their initial strategy for a respective process at least once, are, on average, more likely to sell a process of value $100 + x$ than one of value $100 - x$ (see Figure 7). To make selling decisions comparable, we have here considered only processes with a drift of zero, for which gains and losses are equally likely.¹⁵

¹⁴Strategy adjustments conditional on not hitting a bound 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 12 in Appendix E).

¹⁵A 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 here on the fair processes.

Figure 7: The figure depicts the smoothed empirical cumulative distribution function 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.

Cognitive skills matter. Below-median subjects in terms of cognitive skills are particularly likely to gamble in our experiment, whereby cognitive skills are measured by the sum of correct answers to the modified CRT and the financial literacy questions. 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 (Figure 10 in Appendix E). Notably, below- and above-median subjects in terms of cognitive skills behave both consistently with our main prediction: both are responsive to a change in the drift of the process.

6 Discussion of Alternative Models

6.1 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 1 in Section 5).

In order to rationalize Result 1 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 13 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 turning point from concave to convex would have to change with the outcome level. Saliency theory, in contrast, predicts skewness-dependent risk attitudes for *any* outcome level (see Appendix D and Dertwinkel-Kalt and Köster, forthcoming), and is thus consistent with the data. Moreover, EUT, in contrast to saliency theory, does in general not explain why subjects prefer loss-exit strategies over gain-exit strategies (Result 3 in Section 5). In sum, EUT cannot coherently explain our findings on static and dynamic risk attitudes.

6.2 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 stopping immediately. As we numerically show in Appendix F, at the example of the representative CPT-agent proposed by Tversky and Kahneman (1992),¹⁶ 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 1) nor that stopping behavior is sensitive to the drift of the process (i.e., Result 2).¹⁷ As a consequence of this never-stopping result, CPT is also inconsistent with the disposition effect in a setting like ours (or the one by Ebert and Strack, 2015, as they argue). CPT can, however, also account for Result 3: as 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.*, 2020).

6.3 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.¹⁸ Moreover, Barberis and Xiong (2012) derive results that are seemingly similar to the drift-sensitivity of a

¹⁶It 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.*, 2020) we focus on the representative CPT-agent estimated by Tversky and Kahneman (1992).

¹⁷CPT 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.

¹⁸Barberis 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.

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. But, 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. Frydman *et al.*, 2014; 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 $r = x_0$.¹⁹ 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

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

or, equivalently,

$$2(1-p)(b - x_0) \leq (1 + \lambda)p(x_0 - a). \quad (3)$$

A sufficient condition for Eq. (3) 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 cannot account for our Result 1.

More generally, the preceding analysis highlights that some form of non-linear probability weighting is necessary to explain our results on skewness preferences, not only in the dynamic selling decisions, but also in the static choices, which we analyze in Appendix E (see Figure 13). The former point is made in an informal way also in Heimer *et al.* (2020). 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.

---

¹⁹Precisely, the case of $\lambda = 1$ refers to Eq. (7) in Barberis and Xiong (2012), while $\lambda > 1$ corresponds to Eq. (18) in their paper.

6.4 Disappointment Aversion

Gul (1991) proposes a theory of disappointment aversion to explain the Allais paradox, in particular, the certainty effect.²⁰ The model can, in principle, rationalize a preference for positive skewness 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 a preference for positive skewness, it also predicts that subjects will not stop a process with zero drift before the expiration date, which is inconsistent with Result 1. Moreover, just like EUT and CPT, disappointment aversion cannot account for context-dependent behavior, thereby conflicting with existing evidence from static settings (e.g., Dertwinkel-Kalt and Köster, forthcoming).

If we abstract from a finite expiration date (i.e., if $T = \infty$ holds), a disappointment-averse agent evaluates 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 1) and static choices (see Figure 13 in Appendix E).

So, let us assume that not only $\beta > 0$, but also that $u(\cdot)$ is convex. As in our experiment, we 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 1). 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 preferences in static and dynamic settings.

---

²⁰Disappointment aversion is a special case of cautious expected utility (Cerreia-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).

6.5 Other Saliency Specifications

Adaptive reference point. While in its original formulation saliency theory does not incorporate a reference point in the value function (see Bordalo *et al.*, 2012), we could, in principle, adjust our saliency model by introducing a reference point, r_t , into the value function that might even evolve over time.²¹ Depending on the evolution of the reference point this model can potentially explain a lot of different stopping behaviors. Making specific assumptions on the evolution of the reference point would allow us to falsify the model. But, without any guidance on how the reference point evolves over time, it is impossible to provide a comprehensive analysis. One natural candidate would be a reference point that instantaneously adapts to the current wealth level. But this model would predict, for instance, that subjects have the same disposition to sell assets that have increased or decreased in value, which is clearly inconsistent with our data on disposition-effect-like behavior.

Adaptive consideration set. So far, we have assumed that the salient thinker compares each stopping strategy to the alternative of stopping immediately. While we regard this as a plausible assumption, in particular, in the context of our experiment, one might argue that previously chosen strategies affect the perception of whatever strategy is considered next. As in the case of an adaptive reference point, this additional degree of freedom will probably allow us to rationalize almost any stopping behavior, except for the never-stopping result of CPT. Since a salient thinker's drift-sensitivity relies on the boundedness of the saliency function and is, therefore, unaffected by consideration set effects, even with an adaptive consideration set a salient thinker would stop immediately for a sufficiently negative drift. Again, making specific assumptions on the evolution of the consideration set would allow us to falsify the model. But also here, without any guidance on how the consideration set changes over time, it is impossible to provide a comprehensive analysis. One natural assumption would be 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 saliency 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. This prediction is inconsistent with our data.

7 Conclusion

We theoretically and experimentally investigate the role of endogenously created skewness for how people make dynamic decisions under risk. We find that, if they can create skewness through their stopping strategy, people take on gambles that are fair or slightly unfair. More

---

²¹Alternatively, one could introduce a reference point, $r_t = r$, that is constant over time. In fact, our model could be interpreted as a special case of this more general model with $r = 0$. Depending on the exact reference point r and on whether we want to allow for a differently curved value function over gains and losses, more or less predictions will change. Assuming, for instance, $r = x_0$ and an s-shaped value function as in Barberis (2012), that is not too convex in the loss domain, gives a much more flexible model, but yields predictions quite similar to those of our model.

specifically, gambling with loss-exit strategies is attractive to subjects. These strategies do not only induce a right-skewed distribution of returns, but they also *look* right-skewed, as the stop-loss threshold is closer to the current value of the process than the take-profit threshold. That subjects are particularly attracted by these loss-exit strategies might indicate that it is rather the differential contrasts on the up- and downside than the differential probabilities of up- and downside payoffs that makes subjects gamble. This speaks to the mechanism behind salience theory, according to which the decision weights assigned to salient outcomes are distorted, whereas in CPT the decision weights assigned to extreme probabilities are distorted.

While Ebert and Strack (2015) conclude that “probability weighting in combination with naïvete leads to unrealistic predictions for a wide range of dynamic setups,” we show that this conclusion relies on the specificities of CPT, but does not hold for models of non-linear probability weighting in general. In fact, salience-based probability weighting yields more plausible predictions for dynamic gambling behavior, which we find support for in a laboratory experiment. Our finding that static and dynamic skewness preferences are positively correlated in a way that is consistent with salience theory also suggests that at least salience-based static models of choice under risk can be reasonably applied to a dynamic context.

To account for reasonable stopping behavior, we point to two properties of our salience model that have not been in the focus of the preceding salience literature: boundedness of the salience function (as established by Bordalo *et al.*, 2012) and concavity of the value function (as also put forward by Nielsen *et al.*, 2020). First, boundedness of the salience function is essential to overcome the never-stopping result of CPT. Second, we point to a qualitative difference between diminishing sensitivity of the salience function and concavity or diminishing sensitivity of the value function. Concavity of the value function is the “stronger” concept, meaning that it allows us to rationalize that subjects do not necessarily gamble even fair processes until the end. Jointly, these two properties limit skewness preferences to a “reasonable” degree.

Our results point out the importance of skewness preferences in explaining dynamic choices under risk when the underlying stochastic process is symmetric, but we would expect them to play an even bigger role when the process is positively skewed itself. 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 preferences, 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 is attractive not only due to the skewness created by the stopping strategy, but also 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.*, 2016; Grove *et al.*, 2018). A similar type of argument applies to the classical problem of searching for a job, one of our introductory examples from the stopping-literature. Here, skewness preferences 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 that results from the chosen stopping strategy is amplified by the skewness of the process itself. In sum, skewness preferences can explain time-inconsistent behavior in trying to reach an elusive goal. While the fundamental role that skewness plays for optimal stopping problems has been overlooked by the classical literature, we show that it can be satisfyingly addressed by incorporating a static model of skewness preferences, namely that of salience.

References

- ANDERSON, S. P. and RENAULT, R. (2018). Firm pricing with consumer search. In L. Corchon and M. Marini (eds.), *Handbook of Game Theory and Industrial Organization, Volume II*, Edward Elgar Publishing.
- ARMSTRONG, M. (2017). Ordered consumer search. *Journal of the European Economic Association*, **15** (5), 989–1024.
- BARBERIS, N. 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.
- BARBERIS, N. C. (2012). A model of casino gambling. *Management Science*, **58** (1), 35–51.
- 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., NICOLA GENNAIOLI and SHLEIFER, A. (2012). Salience 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. (2016). The effect of superstar firms on college major choice. <http://www.pbcsf.tsinghua.edu.cn/research/loudong/paper/W%20The%20Effect%20of%20Superstar%20Firms%20on.pdf>, unpublished.

- COX, D. R. and MILLER, H. D. (1977). *The Theory of Stochastic Processes*. Chapman and Hall Ltd.
- DETWINKEL-KALT, M. and KÖSTER, M. (forthcoming). Salience and skewness preferences. *Journal of the European Economic Association*.
- and WENNER, L. (2020). Testing fundamentals of salience theory. http://sesp.bplaced.net/dertwinkel/wp-content/uploads/2020/06/DW_v17.pdf, unpublished.
- DHAMI, S. (2016). *The Foundations of Behavioral Economic Analysis*. Oxford University Press.
- DURAJ, J. (2020). Optimal stopping with general risk preferences.
- EBERT, S. (2015). On skewed risks in economic models and experiments. *Journal of Economic Behavior & Organization*, **112**, 85–97.
- 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. *The Review of Financial Studies*, **30** (6), 2110–2129.
- FRYDMAN, C., BARBERIS, N., CAMERER, C., BOSSAERTS, P. and RANGEL, A. (2014). Using neural data to test a theory of investor behavior: An application to realization utility. *Journal of Finance*, **69** (2), 907–946.
- and MORMANN, M. M. (2018). The role of salience in choice under risk: An experimental investigation. Unpublished.
- GONZALES, J. L. M. (2018). Consumer search in digital markets. In J.-J. Ganuza and G. Llobet (eds.), *Economic Analysis of the Digital Revolution*, Funcas, pp. 181–208.
- 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. (2018). Career lotto? labor supply in winner-take-all markets. <http://people.bath.ac.uk/klp33/tennis.pdf>, unpublished.
- GUL, F. (1991). A theory of disappointment aversion. *Econometrica*, pp. 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., ILIEWA, Z., IMAS, A. and WEBER, M. (2020). 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.
- IMAS, A. (2016). The realization effect: Risk-taking after realized versus paper losses. *American Economic Review*, **106** (8), 2086–2109.
- , LAMÉ, D. and WILSON, A. J. (2017). Preference reversals between one-shot and repeated decisions: The case of regret. <http://www.aleximas.com/s/PrefReversalsRegret-1.pdf>, unpublished.
- 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.
- KÖSTER, M. (2020). A multivariate salience theory of choice under risk. <https://drive.google.com/file/d/1u4DMfvCulPL4NjLnnkoRYK6dZ2Dx6X-i/view>, mimeo.
- KOTHIYAL, A., SPINU, V. and WAKKER, P. P. (2011). Prospect theory for continuous distributions: A preference foundation. *Journal of Risk and Uncertainty*, **42** (3), 195.
- LANZANI, G. (2020). Correlation made simple: Applications to salience and regret theory. <http://economics.mit.edu/files/18851>, unpublished.
- 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, C. S., SEBALD, A. C. and SØRENSEN, P. N. (2020). Testing for salience effects in choices under risk. *Working Paper, University of Copenhagen*.
- 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.
- 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.
- RUST, J. (1987). Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica*, **55** (5), 999–1033.
- 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.
- STRACK, P. and VIEFERS, P. (forthcoming). Too proud to stop: Regret in dynamic decisions. *Journal of the European Economic Association*.
- 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 2. *If $\mu \neq 0$, then, for any threshold stopping time $\tau_{a,b}$ with $a < x < b$, we have*

$$\mathbb{P}_0[X_{\tau_{a,b}} = a] = \frac{\exp(-(2\mu/\nu^2)b) - \exp(-(2\mu/\nu^2)x)}{\exp(-(2\mu/\nu^2)b) - \exp(-(2\mu/\nu^2)a)}. \quad (4)$$

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

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

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

(e) For any stopping time $\tau_{a,b}$ with $a < x < b$ and $z \in [a, b]$, $\frac{\partial}{\partial \mu} \mathbb{P}_0[X_T \leq z | X_0 = x, \tau_{a,b} \geq T] \leq 0$, holding with a strict inequality for any $z < b$. Hence, an increase in the drift of the process improves the distribution of X_T conditional on $\tau_{a,b} \geq T$ in terms of first-order stochastic dominance.

(f) Let $\mu < 0$. Then, for any $T > -\frac{(x-a)}{\mu}$, we have $\frac{\partial}{\partial \mu} \mathbb{P}_0[\tau_{a,b} \geq T | X_0 = x] > 0$.

(g) Let $\mu = 0$. Then, $\frac{\partial}{\partial \mu} \mathbb{P}_0[\tau_{a,b} \geq T | X_0 = x] = 0$.

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

(h) 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 $\mu < 0$.

Finally, suppose that $a = x - \epsilon$ and $b = x + \epsilon'$ for some $\epsilon' > \epsilon$. In addition, let $\mu = 0$.

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

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

$$\lim_{a \rightarrow x} \sum_{n=1}^{\infty} \left\{ \sin \left(\frac{\pi n(x-a)}{b-a} \right) \sin \left(\frac{\pi n(y-a)}{b-a} \right) \exp \left(-\frac{T}{2} \frac{n^2 \pi^2 \nu^2}{(b-a)^2} \right) \right\} = 0. \quad (5)$$

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 (5) inside the summation. The claim follows from the fact that $\sin(0) = 0$.

PART (c). Fix an initial wealth level $x \in \mathbb{R}_{>0}$ and a stopping time $\tau_{a,b}$. Then, we have

$$\begin{aligned} \mathbb{P}_0[\tau_{a,b} \geq T | X_0 = x] &\propto \int_a^b \frac{\exp \left(\frac{\mu(y-x)}{\nu^2} \right)}{\exp \left(\frac{T}{2} \frac{\mu^2}{\nu^2} \right)} \sum_{n=1}^{\infty} \left\{ \sin \left(\frac{\pi n(x-a)}{b-a} \right) \sin \left(\frac{\pi n(y-a)}{b-a} \right) \exp \left(-\frac{T}{2} \frac{n^2 \pi^2 \nu^2}{(b-a)^2} \right) \right\} dy \\ &\leq \frac{1}{\exp \left(\frac{T}{2} \frac{\mu^2}{\nu^2} \right)} \int_a^b \exp \left(\frac{\mu(y-x)}{\nu^2} \right) \sum_{n=1}^{\infty} \exp \left(-\frac{T}{2} \frac{n^2 \pi^2 \nu^2}{(b-a)^2} \right) dy \\ &\leq \frac{1}{\exp \left(\frac{T}{2} \frac{\mu^2}{\nu^2} \right)} \int_a^b \exp \left(\frac{\mu(y-x)}{\nu^2} \right) \sum_{n=1}^{\infty} \exp \left(-n \cdot \frac{T \pi^2 \nu^2}{2(b-a)^2} \right) dy \\ &= \frac{1}{\exp \left(\frac{T}{2} \frac{\mu^2}{\nu^2} \right) \left(\exp \left(\frac{T \pi^2 \nu^2}{2(b-a)^2} \right) - 1 \right)} \int_a^b \exp \left(\frac{\mu(y-x)}{\nu^2} \right) dy \xrightarrow{T \rightarrow \infty} 0, \end{aligned}$$

where the first inequality follows from the fact that $|\sin \left(\frac{\pi n(x-a)}{b-a} \right) \sin \left(\frac{\pi n(y-a)}{b-a} \right)| \leq 1$ and the second inequality holds as $n \geq 1$ and $\frac{T \pi^2 \nu^2}{2(b-a)^2} > 0$.

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

PART (e). Denote $q(y, T) = q(y, T|X_0 = x)$. It follows from Part (d) that the claimed inequality, $\frac{\partial}{\partial \mu} \mathbb{P}_0[X_T \leq z|X_0 = x, \tau_{a,b} \geq T] \leq 0$, holds if and only if

$$\left(\int_a^z (y-x)q(y, T) dy \right) \cdot \left(\int_a^b q(y, T) dy \right) \leq \left(\int_a^b (y-x)q(y, T) dy \right) \cdot \left(\int_a^z q(y, T) dy \right),$$

or, equivalently, $\mathbb{E}_0[X_T|X_0 = x, \tau_{a,b} \geq T, X_T \leq z] \leq \mathbb{E}_0[X_T|X_0 = x, \tau_{a,b} \geq T]$, which is true by the definition of the conditional expectation.

PART (f). Let $\mu < 0$. Then, we have

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

where, in the third line, we multiply with ν^2 and divide by $\mathbb{P}_0[\tau_{a,b} \geq T|X_0 = x]$, and where the inequality follows from the fact that $\mathbb{E}_0[X_T|X_0 = x, \tau_{a,b} \geq T] > a$ by construction.

PART (g). By the proof of Part (f), for $\mu = 0$, we have

$$\frac{\partial}{\partial \mu} \mathbb{P}_0[\tau_{a,b} \geq T|X_0 = x] \propto \mathbb{E}_0[X_T|X_0 = x, \tau_{a,b} \geq T] - x.$$

By Doob's Optional Sampling Theorem, we know that

$$\begin{aligned} \mathbb{E}_0[X_{T \wedge \tau_{a,b}}] &= \mathbb{P}_0[\tau_{a,b} > T] \cdot \mathbb{E}_0[X_{\tau_{a,b}}|\tau_{a,b} < T] + \mathbb{P}_0[\tau_{a,b} \geq T] \cdot \mathbb{E}_0[X_T|\tau_{a,b} \geq T] \\ &= \mathbb{P}_0[\tau_{a,b} > T] \cdot \mathbb{E}_0[X_{\tau_{a,b}}] + \mathbb{P}_0[\tau_{a,b} \geq T] \cdot \mathbb{E}_0[X_T|\tau_{a,b} \geq T] \\ &= \mathbb{P}_0[\tau_{a,b} > T] \cdot x + \mathbb{P}_0[\tau_{a,b} \geq T] \cdot \mathbb{E}_0[X_T|\tau_{a,b} \geq T] \\ &= x, \end{aligned}$$

has to hold. But this implies that $\mathbb{E}_0[X_T|X_0 = x, \tau_{a,b} \geq T] = x$, which proves the claim.

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

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

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

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

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.

PART (i). 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

$$\begin{aligned}
\left| \frac{\partial}{\partial \epsilon} q(y, T | X_0 = x) \right| &\leq \left[\frac{1}{\epsilon'} + \frac{\epsilon' - 2(y-x)}{\epsilon'^2} \right] \frac{\exp\left(\frac{T}{2} \frac{\pi \nu^2}{\epsilon'^2}\right) \pi}{\left(\exp\left(\frac{T}{2} \frac{\pi \nu^2}{4\epsilon'^2}\right) - 1\right)^2} \\
&+ \frac{T \nu^2 \exp\left(\frac{T}{2} \frac{\pi \nu^2}{\epsilon'^2}\right) \left(\exp\left(\frac{T}{2} \frac{\pi \nu^2}{\epsilon'^2}\right) + 1\right) \pi^2}{\epsilon'^3 \left(\exp\left(\frac{T}{2} \frac{\pi \nu^2}{4\epsilon'^2}\right) - 1\right)^3} < \infty,
\end{aligned} \tag{6}$$

this partial derivative indeed exists. To apply L'Hospital's rule, we need to compute the limit of

$\frac{\partial}{\partial \epsilon} q(y, T|X_0 = x)$ for ϵ approaching zero, which is given by

$$\begin{aligned}
\lim_{\epsilon \rightarrow 0} \frac{\partial}{\partial \epsilon} q(y, T|X_0 = x) &= \frac{1}{\epsilon'} \sum_{n=1}^{\infty} \lim_{\epsilon \rightarrow 0} \left\{ \pi n \cos \left(\frac{\pi n \epsilon}{\epsilon + \epsilon'} \right) \sin \left(\frac{\pi n (y - x + \epsilon)}{\epsilon + \epsilon'} \right) \exp \left(-\frac{T}{2} \frac{n^2 \pi^2 \nu^2}{(\epsilon + \epsilon')^2} \right) \right\} \\
&\quad + \frac{\epsilon' - (y - x)}{\epsilon'^2} \sum_{n=1}^{\infty} \lim_{\epsilon \rightarrow 0} \left\{ \pi n \sin \left(\frac{\pi n \epsilon}{\epsilon + \epsilon'} \right) \cos \left(\frac{\pi n (y - x + \epsilon)}{\epsilon + \epsilon'} \right) \exp \left(-\frac{T}{2} \frac{n^2 \pi^2 \nu^2}{(\epsilon + \epsilon')^2} \right) \right\} \\
&\quad + \frac{T \nu^2}{\epsilon'^3} \sum_{n=1}^{\infty} \lim_{\epsilon \rightarrow 0} \left\{ \pi^2 n^2 \sin \left(\frac{\pi n \epsilon}{\epsilon + \epsilon'} \right) \sin \left(\frac{\pi n (y - x + \epsilon)}{\epsilon + \epsilon'} \right) \exp \left(-\frac{T}{2} \frac{n^2 \pi^2 \nu^2}{(\epsilon + \epsilon')^2} \right) \right\} \\
&= \frac{1}{\epsilon'} \sum_{n=1}^{\infty} \left\{ \pi n \sin \left(\frac{\pi n (y - x)}{\epsilon'} \right) \exp \left(-\frac{T}{2} \frac{n^2 \pi^2 \nu^2}{\epsilon'^2} \right) \right\} \\
&\geq 0,
\end{aligned} \tag{7}$$

where the first equality follows from the fact that, by the considerations in Eq. (6), 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 $\lim_{\epsilon \rightarrow 0} q(y, T|X_0 = x) = 0$ (as otherwise $q(y, T|X_0 = x)$ would be negative for ϵ sufficiently close to zero).

Second, we observe that

$$\begin{aligned}
\left| \int_{x-\epsilon}^x \frac{\partial}{\partial \epsilon} q(y, T|X_0 = x) dy \right| &\leq \int_{x-\epsilon}^x \left| \frac{\partial}{\partial \epsilon} q(y, T|X_0 = x) \right| dy \\
&\leq \epsilon \left[\frac{1}{\epsilon'} + \frac{\epsilon' - 2(y - x)}{\epsilon'^2} \right] \frac{\exp \left(\frac{T}{2} \frac{\pi \nu^2}{\epsilon'^2} \right) \pi}{\left(\exp \left(\frac{T}{2} \frac{\pi \nu^2}{4 \epsilon'^2} \right) - 1 \right)^2} \\
&\quad + \epsilon \frac{T \nu^2 \exp \left(\frac{T}{2} \frac{\pi \nu^2}{\epsilon'^2} \right) \left(\exp \left(\frac{T}{2} \frac{\pi \nu^2}{\epsilon'^2} \right) + 1 \right) \pi^2}{\epsilon'^3 \left(\exp \left(\frac{T}{2} \frac{\pi \nu^2}{4 \epsilon'^2} \right) - 1 \right)^3} \xrightarrow{\epsilon \rightarrow 0} 0,
\end{aligned} \tag{8}$$

where the first inequality follows by the triangle inequality, and where the second inequality follows from Eq. (6). Taking the limit of the final expression is straightforward.

Third, we conclude that

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

where the first equality holds by the Theorem of Dominated Convergence, the second one holds by the second to last line in (7), and the third equality follows by substitution. Recall that

$$\sum_{n=1}^{\infty} \left\{ \sin(\pi n z) \pi n \exp \left(-\frac{T}{2} \frac{n^2 \pi^2 \nu^2}{\epsilon'^2} \right) \right\} \geq 0 \tag{9}$$

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 $\pi n \exp\left(-\frac{T}{2} \frac{n^2 \pi^2 \nu^2}{\epsilon'^2}\right)$ being strictly decreasing in n . Since (9) 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 (8) 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. $\square$

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,

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

where the probability $p = p(a, b, \mu)$ is defined in Eq. (4) 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 3.

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 $(*) > 0$. Since $p = \frac{b - x_t}{b - a} = \frac{\epsilon'}{\epsilon + \epsilon'}$, it follows that $(*) > 0$ 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, $\sigma(x_t - \epsilon, x_t)$ 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 3 (i), we have $\lim_{\epsilon \rightarrow 0} \Phi_\mu(x) = 0$, which yields the claim. $\square$

A.3: Main Theoretical Result

Proof of Lemma 1. 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,

$$\frac{v(z) - v(x)}{v'(x)} \frac{\sigma(v(z), v(x))}{\sigma(v(x), v(x))} \leq \frac{\exp(\alpha(z - x)) - 1}{\alpha} \quad (10)$$

holds for all $z \geq 0$. By construction, (10) holds at $z = x$. We distinguish two cases:

1. CASE: Let $z < x$. Divide both sides of (10) by $z - x < 0$, which gives

$$\frac{\frac{v(x) - v(z)}{x - z}}{v'(x)} \frac{\sigma(v(z), v(x))}{\sigma(v(x), v(x))} \geq \frac{\exp(\alpha(x - x)) - \exp(\alpha(z - x))}{\alpha}. \quad (11)$$

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 (11) 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 (11) is strictly larger than 1. In sum, we conclude that, for any $z < x$, Condition (10) is satisfied for any $\alpha > 0$.

2. CASE: Let $z > x$. Since both sides of (10) are zero at $z = x$, we can re-write (10) 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

$$\int_x^z \exp(\alpha(w-x)) - \left[\frac{v'(w)}{v'(x)} \frac{\sigma(v(w), v(x))}{\sigma(v(x), v(x))} + [v(w) - v(x)] \frac{v'(w)}{v'(x)} \frac{\frac{\partial}{\partial v(w)} \sigma(v(w), v(x))}{\sigma(v(x), v(x))} \right] dw \geq 0. \quad (12)$$

A sufficient condition for (12) 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

$$\alpha \geq \frac{\ln \left(\frac{v'(w)}{v'(x)} \frac{\sigma(v(w), v(x))}{\sigma(v(x), v(x))} + [v(w) - v(x)] \frac{v'(w)}{v'(x)} \frac{\frac{\partial}{\partial v(w)} \sigma(v(w), v(x))}{\sigma(v(x), v(x))} \right)}{w - x}, \quad (13)$$

which has to hold for all relevant $w > x$. The right-hand side of (13) 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 (13) 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

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

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 (12) 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 1, 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

$$[\tilde{u}(z) + \beta] \leq [\tilde{u}(x) + \beta] \exp(\alpha(z - x)). \quad (14)$$

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 (14), 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 stop immediately. $\square$

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 stop immediately. 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

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

where the inequality follows from the assumption that the naïve salient thinker with a linear value function immediately stops the process $(X_t)_{t \in \mathbb{R}_{\geq 0}}$. Consequently, the naïve salient thinker with a value function $v(\cdot)$ immediately stops 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. Fix an initial wealth level $x \in \mathbb{R}_{>0}$ and a volatility ν . Consider the two processes with drift parameters $0 \geq \mu' > \mu$. We have to show that *if a naïve salient thinker does not stop the process with drift μ immediately, then he does not stop the process with drift μ' immediately.*

As before, let $\Phi_\mu(z) := \mathbb{P}_0[X_T \leq z | X_0 = x, \tau_{a,b} \geq T]$. If a naïve salient thinker does not stop the process with drift μ immediately, there exists a stopping time $\tau_{a,b}$ such that

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

whereby the probability $p = p(a, b, \mu)$ is defined in Eq. (4). By Lemma 2, an increase in the drift of the process improves the distribution of $X_{\tau_{a,b}}$ in terms of first-order stochastic dominance, and by Lemma 3 (e) the same is true for the distribution of X_T conditional on reaching the expiration date. Hence, by Proposition 1 in Dertwinkel-Kalt and Köster (forthcoming), both $(\star)$ and $(\star\star)$ monotonically increase in the drift μ . We have to distinguish three cases.

1. Case: Suppose that, for the process with drift μ , there exists a stopping time $\tau_{a,b}$ such that $(\star)$ and $(\star\star)$ are non-negative. Since $(\star)$ and $(\star\star)$ monotonically increase in μ , the same stopping strategy is more attractive than stopping immediately also for the process with a drift μ' .

2. Case: Suppose that, for the process with drift μ , any attractive stopping time $\tau_{a,b}$ implies that $(\star)$ is negative, which in turn implies that $(\star\star)$ is positive.

If $\mu' < 0$, then take an attractive stopping time $\tau_{a,b}$ and let $T \geq -\frac{(x-a)}{\mu'} =: \hat{T}$. Then, by Lemma 3 (f), we have $\frac{\partial}{\partial \mu} \mathbb{P}_0[\tau_{a,b} < T] < 0$ for any $\mu < \mu'$ and therefore $[(\star) - (\star\star)] \cdot \frac{\partial}{\partial \mu} \mathbb{P}_0[\tau_{a,b} < T] > 0$

for any $\mu < \mu'$. But this, together with the fact that $(\star)$ and $(\star\star)$ monotonically increase in μ , implies that, for any $T \geq \hat{T}$,

$$\begin{aligned} \mathbb{P}_0[\tau_{a,b} < T] & [p(v(a) - v(x))\sigma(v(a), v(x)) + (1 - p)(v(b) - v(x))\sigma(v(b), v(x))] \\ & + \mathbb{P}_0[\tau_{a,b} \geq T] \int_{(a,b)} (v(z) - v(x))\sigma(v(z), v(x)) d\Phi_\mu(z) \end{aligned}$$

monotonically increases in the drift μ . Obviously, as the drift of the process increases, the sign of $(\star) - (\star\star)$ can change. But, since $(\star)$ and $(\star\star)$ monotonically increase in μ , we would need $(\star) > 0$ for this to happen, which would then bring us back to the first case. Hence, we conclude that, for any $T \geq \hat{T}$, the same stopping time $\tau_{a,b}$ is more attractive than stopping immediately also for the process with a drift μ' .

Now let $\mu' = 0$. By Lemma 3 (g), at $\mu = \mu'$, we have $\frac{\partial}{\partial \mu} \mathbb{P}_0[\tau_{a,b} < T] = 0$ for any T . Since $(\star)$ and $(\star\star)$ strictly increase in μ and since $(\star)$ is independent of T , continuity of the salience-weighted utility implies that there exists some $\epsilon > 0$ and $T' < \infty$ such that, for $T > T'$,

$$\begin{aligned} \mathbb{P}_0[\tau_{a,b} < T] & [p(v(a) - v(x))\sigma(v(a), v(x)) + (1 - p)(v(b) - v(x))\sigma(v(b), v(x))] \\ & + \mathbb{P}_0[\tau_{a,b} \geq T] \int_{(a,b)} (v(z) - v(x))\sigma(v(z), v(x)) d\Phi_\mu(z) \end{aligned}$$

monotonically increases in the drift on $[0, \epsilon)$. The claim follows by setting $\hat{T} := \max\{T', -\frac{(x-a)}{\epsilon}\}$.

3. Case: Suppose that, for the process with drift μ , any attractive stopping time $\tau_{a,b}$ implies that $(\star)$ is positive, but $(\star\star)$ is negative. Then, also for the process with a drift μ' , $(\star)$ is positive. Moreover, for a fixed stopping time $\tau_{a,b}$, $(\star)$ is independent of T . By Lemma 3 (c), it holds that $\lim_{T \rightarrow \infty} \mathbb{P}_0[\tau_{a,b} < T] = 1$, which then implies that, in the limit of T approaching infinity,

$$\begin{aligned} \mathbb{P}_0[\tau_{a,b} < T] & [p(v(a) - v(x))\sigma(v(a), v(x)) + (1 - p)(v(b) - v(x))\sigma(v(b), v(x))] \\ & + \mathbb{P}_0[\tau_{a,b} \geq T] \int_{(a,b)} (v(z) - v(x))\sigma(v(z), v(x)) d\Phi_\mu(z) > 0. \end{aligned}$$

Since the left-hand side of the above expression is continuous in T , we conclude that the inequality already holds for a finite T . More precisely, there exists a $\hat{T} < \infty$ such that for any $T \geq \hat{T}$ the same stopping strategy is more attractive than stopping immediately also for μ' . $\square$

Proof of Proposition 3. 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

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

where the probability $p = p(a, b, \mu)$ is defined as in Eq. (4). 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 inequality follows from the fact that the drift is non-positive in combination with Jensen's Inequality, and diminishing sensitivity of the salience function. The weak inequality holds by Lemma 3 (h), and the last inequality follows again from the concavity of the value function, which implies $v(x_t + z) - v(x_t) \leq v(x_t) - v(x_t - z)$ for any $z > 0$. $\square$

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 5 (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, never gambles 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 stops immediately.*

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 5, the threshold stopping time $\tau_{a,b}$ cannot constitute an equilibrium. In contrast, at any initial wealth level $x \in \mathbb{R}_{>0}$, stopping immediately can be supported as an equilibrium outcome: given that all future selves will stop immediately, the current self is indifferent between stopping immediately and continuing 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. (4) 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 3.

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 3 (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 not stop immediately. By the same arguments as for the case of $T = \infty$, stopping immediately can be supported as an equilibrium outcome, which proves the claim. $\square$

Appendix C: Additional Experimental Details

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 1. Figure 8 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 1: Lotteries used to elicit skewness preferences 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 8: The figure illustrates the layout of the static choices in the experiment (in German).

C.2: Experiment 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 the graph that is depicted in Figure 3.]

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 8.]

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 8.]

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: Saliency Predictions on Static Skewness Preferences

In this section, we extend a result from Dertwinkel-Kalt and Köster (forthcoming) on a salient thinker's skewness preferences 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 (forthcoming) 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, forthcoming). *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 4 and Figure 6). A positive correlation between the two measures indicates that a subject revealing skewness preferences consistent with saliency 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,

$$\pi \frac{v \left(E + \sqrt{\frac{V}{\pi}} \right) - v(E)}{v(E) - v \left(E - \sqrt{V\pi} \right)} > \frac{\sigma \left(v \left(E - \sqrt{V\pi} \right), v(E) \right)}{\sigma \left(v \left(E + \sqrt{\frac{V}{\pi}} \right), v(E) \right)}, \quad (15)$$

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 (15) 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 (15) monotonically increases in π . Second, we observe that the left-hand side monotonically decreases in the likelihood ratio π :

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

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 (15) 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 (forthcoming), 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 (forthcoming) 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 9: 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 10: 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 11: 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%	Before	Loss-Exit	72.65%	0%
	Gain-Exit	12.01%	14.37%		Gain-Exit	3.48%	23.87%

Figure 12: 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 13: The figure depicts the share of subjects choosing each of the lotteries depicted in Table 1 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,²² 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),

$$U(x) = \begin{cases} (x - r)^\alpha & \text{if } x \geq r, \\ -\lambda \cdot (r - x)^\alpha & \text{if } x < r, \end{cases} \quad (16)$$

with parameters $\alpha \in (0, 1]$ and $\lambda > 1$.²³ 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):²⁴

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

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

$$\begin{aligned} CPT(X) &:= \int_{\mathbb{R}_+} w^+(\mathbb{P}[U(X) > y]) dy - \int_{\mathbb{R}_-} w^-(\mathbb{P}[U(X) < y]) dy \\ &= \int_{\mathbb{R}_+} w^+(\mathbb{P}[X > r + y^{1/\alpha}]) dy - \int_{\mathbb{R}_-} w^-(\mathbb{P}[X < r - (-y)^{1/\alpha}/\lambda]) dy, \end{aligned} \quad (17)$$

where the second equality holds due to a (weakly) S-shaped value function in Eq. (16).

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

²²In 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.

²³As argued in Wakker (2010, p. 270), the model is ill-specified when taking different α for gains and losses.

²⁴The bounds on the parameters are taken from Dhami (2016, p. 122).

with this random variable is given by

$$\begin{aligned} CPT(X_{T \wedge \tau_{a,b}}) &= \int_{(0, (b-r)^\alpha)} w^+(\mathbb{P}_t[X_{T \wedge \tau_{a,b}} > r + y^{1/\alpha}]) dy \\ &\quad - \int_{(-\lambda(r-a)^\alpha, 0)} w^-(\mathbb{P}_t[X_{T \wedge \tau_{a,b}} < r - (-y)^{1/\alpha}/\lambda]) dy. \end{aligned} \quad (18)$$

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

$$CPT(X_{T \wedge \tau_{a,b}}) = \int_{((a-r)^\alpha, (b-r)^\alpha)} w^+(\mathbb{P}_t[X > r + y^{1/\alpha}/\lambda]) dy + (a-r)^\alpha, \quad (19)$$

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

$$CPT(X_{T \wedge \tau_{a,b}}) = - \int_{(-\lambda(r-a)^\alpha, -\lambda(r-b)^\alpha)} w^-(\mathbb{P}_t[X < r - (-y)^{1/\alpha}/\lambda]) dy - \lambda(r-b)^\alpha. \quad (20)$$

At time $t < T$ with a current wealth level $x_t \in \mathbb{R}_{>0}$, we consider the following class of threshold stopping times:²⁵ 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 14 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 (naïvely) right-skewed 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 14 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 14 further 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

²⁵These strategies are similar to those used in the proof of Theorem 1 in Ebert and Strack (2015).

Figure 14: The figure depicts $CPT(X_{T \wedge \tau_{a_{t,k}, b_{t,k}}})$ 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 $\gamma = 0.61$.

(2015)—which is 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 becomes smaller, the CPT value of the depicted, loss-exit strategies increases (before it eventually falls toward zero). This follows from the fact that, for a loss-exit strategy that induces a sufficiently skewed return distribution, the distribution of $X_{T \wedge \tau_{a_k, b_k}}$ improves in terms of first-order stochastic dominance as the remaining time until the expiration date decreases. To see why, notice that

$$\begin{aligned} \frac{\partial}{\partial T} \mathbb{P}_t[X_{T \wedge \tau_{a,b}} \leq z] &= -p(a, b, x_t) \int_z^b \frac{\partial}{\partial T} q(y, T-t | X_t = x_t) dy \\ &\quad + (1 - p(a, b, x_t)) \int_a^z \frac{\partial}{\partial T} q(y, T-t | X_t = x_t) dy, \end{aligned}$$

where, for any $\tau_{a,b}$ with $a = x_t - kp$ and $b = x_t + k(1-p)$ for a fixed $p = p(a, b, x_t) \in (0, 1)$, we have $\frac{\partial}{\partial T} q(y, T | X_t = x_t) > 0$. As p approaches zero, it follows that $\frac{\partial}{\partial T} \mathbb{P}_t[X_{T \wedge \tau_{a,b}} \leq z] > 0$. This result is illustrated in Figure 15. But notice that the ranking in terms of first-order stochastic dominance is not perfect since the depicted strategy is not skewed enough.

Importantly, both patterns discussed above are robust to using different CPT specifications (e.g., a piece-wise linear value function with $\alpha = 1$ or a reference point of $r_t = 0$). This suggests

Figure 15: The figure depicts the CDF of $X_{T \wedge \tau_{a,b}}$ for the time invariant strategy with a lower bound $a = 99.9$ and an upper bound $b = 109.9$ for a current wealth level of $x_t = 100$ and a remaining time until the expiration date $T - t \in \{0.5, 1, 10\}$. We assume a drift of $\mu = -2$ and a volatility of $\nu = 5$.

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.