INTERNATIONAL ECONOMIC REVIEW
Vol. 65, No. 2, May 2024 DOI: 10.1111/iere.12681

OPTIMAL STOPPING IN A DYNAMIC SALIENCE MODEL*

By Markus Dertwinkel-Kalt and Jonas Frey

University of Münster, Germany; Max Planck Institute for Collective Goods, Germany; Said Business School, University of Oxford, U.K.

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

1. INTRODUCTION

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

Since skewness seeking in static and dynamic decisions plausibly stems from the same cognitive mechanism, the existing literature has adapted static models of choice under risk—such as expected utility theory (EUT) or cumulative prospect theory (CPT; Tversky and Kahneman, 1992)—to dynamic decision making. We follow this approach by studying, both theoretically and experimentally, the dynamic implications of Bordalo et al.'s (2012) salience theory. In salience theory, skewness seeking in static problems originates from the idea that outcomes, which are extreme relative to a reference point—such as the relatively large upside of a right-skewed risk—are particularly salient. These outcomes attract a disproportionate amount of attention and their probabilities are overweighted (Bordalo et al., 2012). Thus, salience theory

*Manuscript received October 2022; revised August 2023.
The experiments reported in this article were preregistered in the AEA RCT registry as trials AEARCTR-0005359 and AEARCTR-0007398, respectively. First of all, we thank Mats Köster for his continuous support throughout the duration of the project. We also thank the editor, Masaki Aoyagi, two anonymous referees as well as Sebastian Ebert, Nicola Gennaioli, Katrin Gödker, Paul Heidhues, Alex Imas, Joshua Miller, Frank Schlütter, Philipp Strack, and Peter Wakker for helpful comments and suggestions. We gratefully acknowledge financial support by the German Research Foundation (DFG project No. 462020252). Please address correspondence to: Markus Dertwinkel-Kalt, University of Münster, Am Stadtgraben 9, 48143 Münster, Germany. E-mail: *markus.dertwinkel-kalt@wiwi.uni-muenster.de*

exhibits the core prediction that agents seek skewness (Dertwinkel-Kalt and Köster, 2020). At the same time, the salience distortions are bounded, which places a tight limit on the degree of probability weighting and therefore also on skewness seeking. Models that imply much stronger skewness seeking than salience theory (such as CPT) could make reasonable predictions in static settings but make unreasonably extreme predictions in some important dynamic settings (Ebert, 2015). Hence, salience theory represents a natural candidate for modelling skewness seeking in dynamic choices.

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

In Section 3, we adapt salience theory to analyze dynamic choice under risk at the hand of standard optimal stopping problems. We ask when a naïve[1] salient thinker stops an arithmetic Brownian motion (ABM) with a nonpositive drift and a finite expiration date. EUT with a concave utility function cannot explain gambling when the agent loses money in expectation, that is, when the ABM's drift is negative. With a specific stopping strategy in mind, a *salient thinker*, however, inflates the probabilities of those realizations that "differ most" from his current wealth level (*contrast effect*), which might render gambling attractive. Adopting the naïve decision rule proposed in the literature (Ebert and Strack, 2015), we assume that the (naïve) salient thinker continues to gamble as long as he can find at least one stopping strategy that is more attractive to him than stopping.

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

Section 5 presents a laboratory experiment that is designed to test our salience-based predictions on stopping behavior. Participants in the experiment have to decide when to stop ABMs with different nonpositive drifts. Subjects stop the process by defining an upper and a lower bound and the process is stopped if it reaches either bound. If a process is stopped, subjects can either sell it or restart it by moving the bounds. This design allows us to test whether subjects choose loss-exit strategies (i.e., strategies with a salient upside) and whether they

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

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

revise their initial strategies as predicted by the model. We validate our approach of adapting the static salience model to an optimal stopping problem by further eliciting skewness seeking in static choices. Generalizing results from Dertwinkel-Kalt and Köster (2020), we show that, for a fixed expected value and variance, a salient thinker chooses a binary lottery over the safe option paying its expected value with certainty if and only if the lottery's skewness exceeds some threshold. If salience is indeed the psychological mechanism driving skewness seeking in general, it should coherently explain revealed attitudes toward skewness in such static choices as well as in the optimal stopping problems.

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

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

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

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

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

2. RELATED LITERATURE ON BEHAVIORAL STOPPING

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

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

For our purpose of testing salience theory's predictions, the setups with naïvete, but without commitment are relevant (Barberis, 2012; Ebert and Strack, 2015): naïvete is a more plausible assumption than sophistication (e.g., sophisticates should not start gambling with non-positive drifts), and only the absence of commitment allows to test whether predictions on time-inconsistent behavior hold true. In the seminal paper by Barberis (2012), it is numerically shown that in finite discrete time setups, naïve CPT agents without commitment mostly choose loss-exit (as compared to gain-exit) strategies and start to gamble (at least for a wide range of parameters), but revise their strategies, so that ex post they exhibit gain-exit behavior. The reason is that close to the expiration date, agents cannot choose strongly skewed return distributions anymore, and therefore, exit earlier than intended.[3] In continuous time setups, this mechanism is not at work: agents can always choose strongly skewed return distributions and, as a consequence, naïve CPT-agents never stop (Ebert and Strack, 2015). This also holds true with an indefinite end date: for most empirically relevant cumulative prospect theory parameter values, a naïve agent does not stop with probability 1 at any loss level (He et al., 2019). This never-stopping prediction can only be avoided by allowing for randomized stopping strategies and thereby some form of commitment (Henderson et al., 2017), or by imposing different functional forms for dynamic choices than those typically used for static choices (Duraj, 2020; Huang et al., 2020). We show in how far the predictions of salience theory differ from those by other models such as CPT, and show, in particular, that dynamic salience theory does not yield the (too) extreme never-stopping prediction that Ebert and Strack (2015) have derived for CPT.

We also contribute to the small, but growing experimental literature on behavioral stopping (Imas, 2016; Nielsen, 2019; Strack and Viefers, 2021). Unlike us, these papers focused on the question in how far stopping decisions are path-dependent; in particular, in how far the realization of previous gains and losses affects behavior. Closest related to us is the contemporaneous paper by Heimer et al. (2023), who study optimal stopping behavior using a process consisting of repeated (fair) coin tosses. Similarly to us, they focus on stop-loss and take-profit strategies, and they find, both in laboratory experiments as well as in observational brokerage data, that subjects ex ante choose loss-exit strategies, but then deviate by revealing disposition-effect-like behavior. Both findings are also reflected in our data. In contrast to Heimer et al. (2023), our article is focused on salience theory and establishes a novel link between static and dynamic skewness seeking.

3. A DYNAMIC VERSION OF SALIENCE THEORY OF CHOICE UNDER RISK

3.1. *Static Model.* Consider an agent who has to choose from some set $\mathcal{C}$ that contains exactly two nonnegative random variables (or *lotteries*), X and Y, with a joint cumulative distribution function (CDF) $F : \mathbb{R}^2_{\geq 0} \to [0, 1]$. A state of the world here refers to a tuple of

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

outcomes, $(x, y) \in \mathbb{R}^2_{\geq 0}$. We denote the state space by $\mathcal{S} \subseteq \mathbb{R}^2_{\geq 0}$. If a random variable is degenerate, we call it a *safe* option.

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

DEFINITION 1 (SALIENCE FUNCTION). We say that a symmetric, bounded, and twice differentiable function $\sigma : \mathbb{R}^2_{\geq 0} \to \mathbb{R}_{>0}$ is a salience function if and only if it satisfies the following two properties:[4]

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

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

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

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

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

A salient thinker is intrinsically (weakly) risk-averse but may, depending on the salience of outcomes, sometimes behave as if being risk-seeking. He evaluates monetary outcomes via a strictly increasing, (weakly) concave, and twice differentiable value function $v : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$, and forms an "expectation" by assigning each state of the world a (probability) weight that depends on how a given option compares to the alternative at hand in this state. More specifically, a salient thinker behaves as if maximizing a *salience-weighted utility*, which is defined as follows:

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

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

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

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

[4] Bordalo 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 $\sigma(-x, -y) > \sigma(-w, -z)$.

3.2. *Dynamic Model.*

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

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

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

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

Stopping strategies As in Ebert and Strack (2015), we represent the set of stopping strategies by the set of stopping times, where each stopping time τ refers to a deterministic plan of when to stop the process. The central feature of a stopping time is that it is based on past information only: that is, any τ is adapted to the natural filtration $(\mathcal{F}_t)_{t\in\mathbb{R}_{\geq 0}}$ of the process $(X_t)_{t\in\mathbb{R}_{\geq 0}}$. For a fixed expiration date $T \in \mathbb{R}_{>0} \cup \{\infty\}$, choosing a stopping time $\tau \leq T$ (with probability 1) implements a random wealth level X_τ with a CDF denoted by F_τ.

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

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

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

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

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

4. STOPPING BEHAVIOR OF A NAÏVE SALIENT THINKER

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

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

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

4.2. *Main Theoretical Result.* We are interested in how general the never-stopping result derived in the previous subsection is. By Definition 3, a salient thinker continues (or starts) to gamble if he can find a strategy that gives him strictly higher utility than not gambling. We will now show that there are two reasons why a salient thinker cannot find such a strategy and hence stops before the expiration date: either the drift of the process is sufficiently negative, or the salient thinker is intrinsically risk-averse. More precisely, while a naïve salient thinker with a linear value function also holds processes with a slightly negative drift until the expiration date, a salient thinker with a sufficiently concave value function does not start even a fair process, and this holds irrespective of his intrinsic risk-aversion (i.e., irrespective of how

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

concave his value function is). This last prediction distinguishes salience theory from models like CPT (see Ebert and Strack, 2015), and it constitutes our main theoretical result.

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

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

where F_τ denotes the CDF of the induced wealth level X_τ. Fixing the initial value $X_0 = x$, we define an *auxiliary utility function* $\tilde{u}(z) := (v(z) - v(x))\sigma(v(z), v(x))$, which is strictly increasing and differentiable in $z \in \mathbb{R}_{\geq 0}$. By construction, the condition derived in Equation (1) is equivalent to

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

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

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

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

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

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

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

That salience theory produces more realistic predictions than CPT is due to the boundedness of the salience function; if the salience function was unbounded, similar predictions as in CPT would prevail. While Bordalo et al. (2012) did not psychologically motivate the boundedness of the salience function they assumed, it is in line with the well-known fact that humans have difficulties interpreting numbers outside of the range they commonly experience, namely, very small and very large numbers (see Resnick et al., 2017, for a review). Salience distortions stem from large payoff contrasts that attract attention. But when humans cannot well understand the difference in magnitude between two large numbers, then it is natural to assume that increasing a large contrast even further does not induce a (strong) behavioral reaction, and therefore, should also not distort salience weights much further (as salience theory wants to well-describe actual behavior). This is precisely the effect that the boundedness of the salience function produces.[6]

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

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

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

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

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

DEFINITION 5 (LOSS-EXIT STRATEGY). A stop-loss and take-profit strategy $\tau_{a,b}$ (with b denoting the upper and a the lower threshold) is a loss-exit strategy at the wealth level x_t if $b - x_t > x_t - a$.

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

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

expiration date and calculate the skewness of the resulting empirical distribution (see Online Appendix A.6).

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

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

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

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

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

What happens if we allow for a finite expiration date instead? Because the drift of the process affects the probability of stopping before the expiration date, it is not clear, in general, whether the distribution of $X_{T\wedge\tau_{a,b}}$ improves in terms of first-order stochastic dominance as the drift increases. For loss-exit strategies, however, an increase in the drift does improve the distribution of $X_{T\wedge\tau_{a,b}}$ in terms of first-order stochastic dominance (Lemma 2(d) in Online Appendix A). Hence, by Proposition 2, we can again invoke Proposition 1 in Dertwinkel-Kalt and Köster (2020) to establish that a salient thinker's gambling behavior is monotonic in the drift of the process.

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

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

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

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

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

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

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

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

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

with $\theta > 0$. In Online Appendix A.5, we examine whether this salience specification predicts the disposition effect and find that this is not the case. Hence, in order for salience theory to predict the disposition effect, other salience specifications are needed: either other salience functions or other value functions (e.g., a piece-wise linear value function reflecting loss aversion, as used in the Online Appendix of Bordalo et al., 2012).

5. AN EXPERIMENT ON DYNAMIC GAMBLING BEHAVIOR

In this section, we present and discuss our experimental design.[8]

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

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

FIGURE 1

SCREENSHOTS OF THE DECISION SCREEN FOR THE PROCESS WITH ZERO DRIFT (IN GERMAN)

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

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

As it is illustrated in Figure 1, we restricted the choice set to all stop-loss and take-profit strategies: at every point in time, subjects could choose an upper and a lower stopping threshold. Once the price of the asset reached either threshold, subjects could decide whether to sell the asset at this price or to adjust the thresholds and continue the process (see the lower left panel). Therefore, the strategies were nonbinding, which ruled out any form of commitment.

[9] Notice 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 in the experiment, so that a drift of, for example, $\mu = 1$ was displayed as a drift of 10 per second.

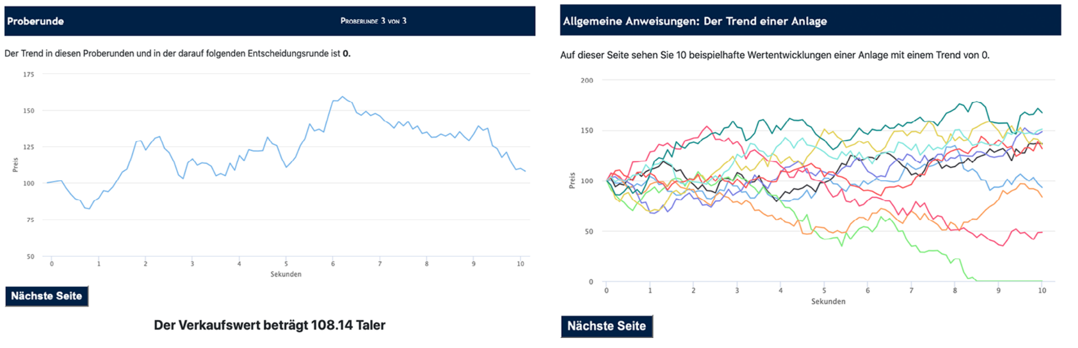

FIGURE 2

SCREENSHOTS OF THE SAMPLING SCREENS FOR THE PROCESS WITH ZERO DRIFT (IN GERMAN)

Subjects could pause the process at any point in time to adjust the thresholds (see the upper right panel). But, importantly, subjects could set only one upper threshold and one lower threshold at a time, and thus observed a stopping strategy in isolation. At the beginning, the upper and lower thresholds were centered symmetrically around the initial price (see the upper left panel). To start the process, subjects had to move each threshold at least once. Before starting the process, subjects could decide to sell the asset immediately (see also the upper left panel).

Overall, subjects made six selling decisions, one decision for each of the drift parameters. The order of the drifts was randomized at the subject level. It is important for our analysis to make sure that subjects understand before the start of the decision round that a non-positive drift of a process means that they will, on average, not gain money from gambling with this process. If they do not understand this, they might start gambling (even if they have EUT preferences), because they expect to earn money; then they would stop as soon as they have learned from observing the process during the decision round that they are losing money over time.

We have two approaches to rule out such learning about the drift during the decision round: First, at the beginning of each round, we inform the subjects about this round's drift. If subjects fully understand the implications of a given drift, this would be sufficient to rule out learning. However, subjects may not understand the meaning of a drift, despite our explanations in the instructions. Therefore, we also let subjects watch the development of three sample paths from the underlying process for 10 seconds each. Moreover, we show them an overview of 10 additional sample paths of the process (see Figure 2). These 13 sample paths that subjects see before making a decision (which were randomly drawn at the subject level, meaning that different subjects saw different sample paths of the same underlying process) should give them a quite good understanding of the process and its drift. Thus, we feel confident that substantial learning during the decision round while watching the change of the process is not a driver of our results.

After completing the six stopping problems, subjects faced a series of 12 (static) choices between a binary lottery and the safe option paying the lottery's expected value with certainty. We used two sets of lotteries with the exact same expected value (either € 30 or € 50) and the exact same variance, but different levels of skewness (see Table 2 in Online Appendix C for an overview). The order of the lotteries was randomized at the subject level. Finally, subjects answered five questions of a modified cognitive reflection test (CRT; closely aligned to Primi et al., 2016), and the five financial literacy questions proposed by Lusardi and Mitchell (2011). All additional questions are listed in Online Appendix C.

At the end of the experiment, for each subject, one of the six selling decisions was randomly drawn by the computer to be payoff-relevant. We randomly selected one subject in each session for whom, in addition, 1 of the 12 static choices was randomly chosen to be payoff-relevant. Subjects were further rewarded for correctly answering CRT and financial literacy questions (1 Taler per question), and they received an additional € 4 for their participation.

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

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

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

> "*In this experiment you will see assets of varying profitability. How profitable an asset is in the long run is described by the drift of the asset. The drift denotes the average change in the value of the process per second.*
> *A positive drift implies that the asset will increase in value in the long run, while a negative drift implies that the asset will decrease in value in the long run. Notice that the value of the asset varies. Hence, even an asset with a negative drift sometimes increases in value.*"

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

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

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

This may raise the concern that subjects could think (at least if they did not carefully read the previous part of the instructions, stating that a negative drift gives, on average, a decrease in value) that even assets with negative drifts are, on average, a profitable investment. In this case, the total of 13 sample paths that subjects see for each drift before making their selling decision should give subjects a rough understanding of the expected value of the process.

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

Features of the process To make our theory testable, we deviate from Ebert and Strack (2015) in two ways: First, since it is impossible to implement a process that can run forever with probability one, we implemented—similar as Heimer et al. (2023)—a finite expiration date. Alternatively, we could have implemented a random termination rule, according to which, at time t, the asset is automatically sold with probability $\omega_t \in [0, 1]$. A finite expiration date makes a theoretical analysis of stopping behavior feasible, whereas with a random termination rule, the probability distribution associated with a given stop-loss and take-profit

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

strategy would not be tractable anymore. A finite expiration date is also easier to explain to the subjects, which we regard—given the complexity of the experiment—as a major advantage. Second, to ensure incentive compatability, we make the process absorbing in zero.[11] We further restrict the drift of the process to be nonpositive because processes with a positive drift do not allow us to separate between the predictions of different models such as EUT, CPT, and salience theory.

Duration of the process A potential concern of our experimental design is that the process only runs for 10 seconds, which is significantly shorter than the time horizon of our motivating examples such as stock trading or job search. However, during the 10 seconds processes already change considerably, and this rather short time span makes it easy to visually follow the development of the process. Notably, skewness both intuitively and theoretically affects short- and long-run behavior likewise. Hence, abstracting from concerns about reaction times (which are discussed below), skewness effects can be studied with short and long processes alike. Notably, subjects can pause the process at any time and thus slow down their decision making. We view the short duration of the process even as an advantage as it allows subjects to stay focused on the task, which would become increasingly difficult for longer horizons.

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

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

[11] 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.

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

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

[14] 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. But, even if subjects adopted other strategies, the test of our main theoretical result—namely, Theorem 1—would still be valid, as here we did not impose any restriction on the choice set.

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

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

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

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

A sophisticated salient thinker differs from her naïve counterpart in that she anticipates her future selves to act in a different way than her present self does, which she takes into account when making her stopping decision. A sophisticated salient thinker who lacks commitment then behaves as if she was playing a game with her future selves (as in, 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 Online Appendix B, a sophisticated salient thinker, who lacks commitment and chooses from the set of all stop-loss and take-profit strategies, does not start any process with a nonpositive drift. Consequently, (partial) naïvete is a necessary assumption to rationalize gambling in the context of our experiment within the salience framework.

5.3. *Experimental Predictions.* We now state the precise predictions of salience theory that guided our experimental design. We slightly deviate from our preregistration, which was based on the salience model with a linear value function: due to the weakly concave value function, Prediction 2 differs from what was preregistered, whereas the remaining predictions are identical to the preregistered ones.

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

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

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

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

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

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

Finally, since we extend a theory of "static" choice under risk to a dynamic setup, we are interested in the empirical relationship between a subject's attitude toward static and dynamic risks. If salience is indeed the psychological mechanism underlying our results, it should coherently explain behavior revealed in static and dynamic choices. As we show in Online Appendix D, a salient thinker chooses a binary lottery, with a fixed expected value and a fixed variance, over the safe option paying its expected value if and only if the lottery's skewness exceeds a certain threshold. By Proposition 2, this preference for positive skewness is also what drives a salient thinker's stopping behavior. We therefore classify both static and dynamic choices into being *skewness-seeking* or not. We say that a static choice is skewness-seeking if the subject chooses a right-skewed lottery over the safe option or the safe option over a left-skewed or symmetric lottery (see Table 2 in Online Appendix C for the exact lotteries and classification). We further classify a stopping strategy as being skewness-seeking if it is a loss-exit strategy, and thus, induces a right-skewed return distribution. Based on Proposition 2, we expect a positive correlation between the share of skewness-seeking choices in static and dynamic decisions.

TABLE 1
THE TABLE SHOWS DESCRIPTIVE STATISTICS FOR OUR EXPERIMENTAL DATA

Drift:	0.0	−0.1	−0.3	−0.5	−1.0	−2.0
% Sold Immediately	5.70%	9.49%	13.29%	22.15%	27.22%	41.14%
% Never Sold	18.99%	8.86%	9.49%	5.06%	3.80%	1.27%
Termination Value	95.45	92.34	87.98	81.89	73.97	82.74
	(40.19)	(38.99)	(33.63)	(32.59)	(34.03)	(26.70)
Stopping Time	6.31	5.03	3.77	2.99	2.04	0.74
	(3.72)	(3.69)	(3.72)	(3.59)	(3.16)	(1.85)
Upper Bound	136.92	133.57	124.97	124.93	122.64	122.14
	(27.99)	(28.47)	(29.51)	(25.22)	(26.63)	(26.80)
Lower Bound	71.82	75.88	68.40	71.65	66.42	77.75
	(31.23)	(33.82)	(31.38)	(33.21)	(32.24)	(30.04)
Distance Lower Bound to Value	19.02	16.80	19.38	18.59	19.99	14.94
	(18.85)	(17.80)	(20.61)	(20.34)	(22.52)	(21.07)
Distance Upper Bound to Value	33.76	28.86	26.74	24.55	25.47	24.99
	(27.74)	(27.02)	(22.40)	(25.98)	(27.02)	(31.30)

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

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

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

6. EXPERIMENTAL RESULTS ON DYNAMIC GAMBLING BEHAVIOR

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

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

Table 1 shows descriptive statistics for the data from our experiment. We can see that the share of subjects who sell the asset immediately increases as the drift becomes more negative. The share of subjects who do not sell the asset before it expires and the average time the asset was held decreases as the drift becomes more negative. Moreover, the upper bounds are, on average, further away from the current value of the process than the lower bounds for all

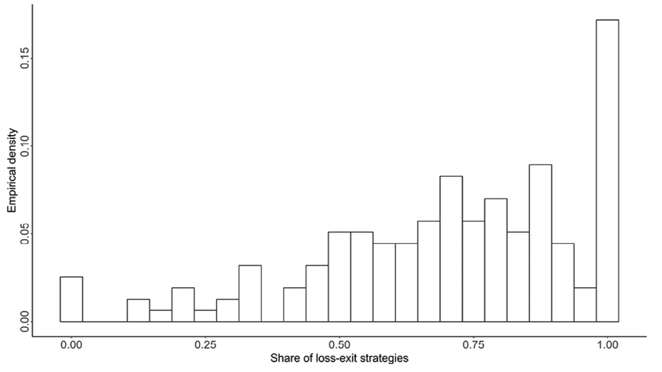

NOTES: The share is calculated on the subject level by taking all strategies, including both initial and revised strategies, aggregating across different drifts, and determining the percentage of those strategies that are loss-exit strategies.

FIGURE 3

THE FIGURE DEPICTS THE EMPIRICAL DISTRIBUTION OF THE SHARE OF LOSS-EXIT STRATEGIES ACROSS SUBJECTS

drifts. The termination value also decreases with the drift except for the comparison between the drifts −1 and −2. The higher termination value for a drift of −2 is likely driven by the higher fraction of subjects selling the asset immediately, at a value of 100).

6.2. *Main Test of our Salience Predictions.* First, we show that, consistent with Prediction 1, a majority of subjects initially choose a loss-exit strategy. This result on initial strategies holds across all the different drifts that we considered (see Figure 12 in Online Appendix E).

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

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

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

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

Overall, these results suggest that the majority of subjects are skewness seeking as predicted by our model. The fact that not all selected strategies are right-skewed can partially be explained by inherent noise in experimental data collection, and partially by subject heterogeneity (meaning that not all subjects are salient thinkers).

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

To address Prediction 2, we look into stopping behavior for the fair process with zero drift. Around a fifth of all subjects hold the asset with a drift of zero until the expiration date,

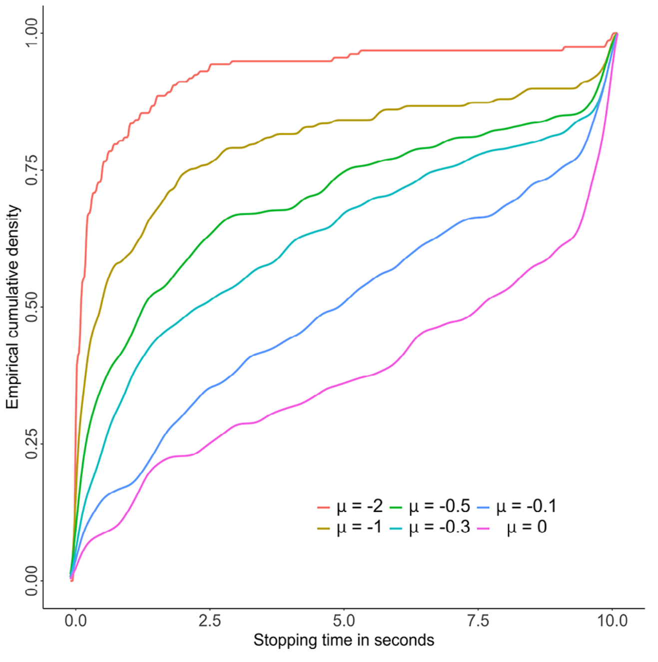

FIGURE 4

THE FIGURE DEPICTS THE SMOOTHED EMPIRICAL CUMULATIVE DISTRIBUTION FUNCTIONS OF STOPPING TIMES, ONE FOR EACH OF THE DIFFERENT DRIFTS

whereas 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. Moreover, we calculate the median of the subjects' stopping times, which gives our next result:

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

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

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

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

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

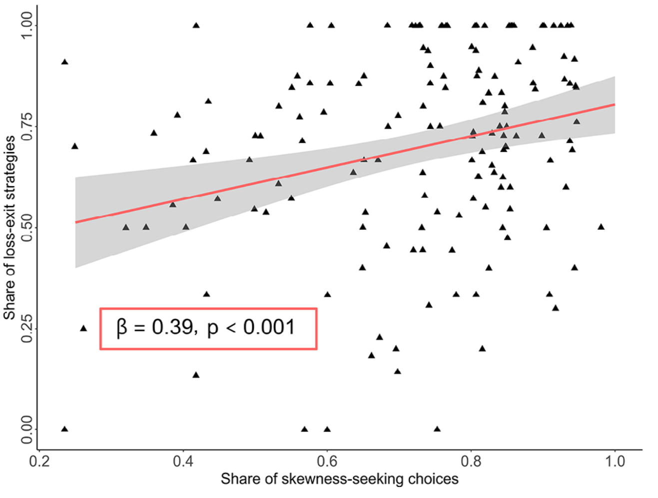

NOTES: We further provide the estimated slope coefficient of the depicted linear regression, which is significantly larger than zero. The data points are scattered for illustrative purposes.

FIGURE 5

THE FIGURE DEPICTS THE RELATIONSHIP BETWEEN STATIC AND DYNAMIC GAMBLING BEHAVIOR

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

RESULT 3 (B). *Subjects stop earlier for processes with more negative drifts.*

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

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

Finally, we study whether skewness seeking in static and dynamic decisions is related. As depicted in Figure 5, subjects behave quite consistently in the static and the dynamic decision problems. To test for the link between static and dynamic skewness seeking, we regress the share of loss-exit strategies among all strategies chosen throughout the six dynamic problems on the share of skewness-seeking choices in the 12 static problems. We find a positive and statistically significant correlation, which gives our fourth result:

[15] This 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.

RESULT 4. *The share of skewness-seeking choices by a subject in the 12 static decisions is positively correlated with the share of loss-exit strategies this subject chooses in the six selling decisions.*

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

6.3. *On the Salience Mechanism: Frequency and Direction of Strategy Adjustments.* Consistent with our model, strategy revisions are ubiquitous and follow precise patterns. Altogether, (i) more than 93% of the subjects (147 out of 158) revised their initial strategy in at least one of the six selling tasks, (ii) conditional on starting, subjects adjust their strategies 1.6 times per task, and (iii) about 70% of the strategy adjustments happen in an attempt to prolong gambling after the process has hit one of the previously chosen thresholds. Moreover, if a subject chooses a loss-exit strategy and the process hits a threshold, the subject is—conditional on not stopping the process—more than six times as likely to again choose a loss-exit rather than a gain-exit strategy (see the left table of Figure 13 in Online Appendix E),[16] which is consistent with Prediction 1.

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

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

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

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

[16] Strategy adjustments conditional on not hitting a threshold follow a similar pattern. Suppose that in the moment of pausing the process, the currently played strategy is a loss-exit strategy. Then, our subjects were more than 10 times as likely to select another loss-exit strategy and compared to a gain-exit strategy (right table in Figure 13 in Online Appendix E).

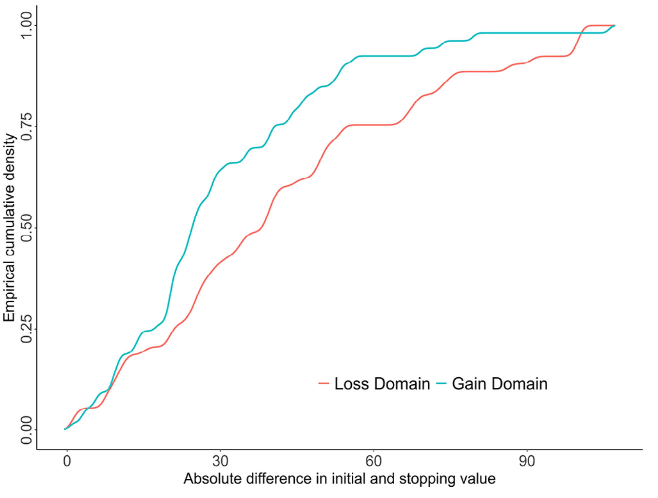

NOTES: We consider only fair processes with a drift of zero for which the initial strategy was adjusted at least once.

FIGURE 6

THE FIGURE DEPICTS THE SMOOTHED EMPIRICAL CDF OF STOPPING AT A GIVEN DISTANCE TO THE INITIAL VALUE OF THE PROCESS, SEPARATELY FOR PROCESSES THAT HAVE GAINED AND THAT HAVE LOST IN VALUE

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

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

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

6.5. *Limitations of the Model in Explaining the Data.* Although the majority of subjects choose mostly strategies inducing right-skewed return distributions, strategies inducing left-skewed return distributions are, unlike what our model predicts, also chosen. This can be explained by noise in the data and by heterogeneity in subjects' susceptibility to salience.

Another issue is that, despite the good fit of our model and our data, we cannot conclusively show that salience indeed is the mechanism that drives our results. Given the strong correlation between skewness effects revealed in static and in dynamic choices, there arguably is

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

one cognitive mechanism that drives all of these skewness effects, but this does not have to be salience. In any case, we think of the model as a useful "as if" model.

7. DISCUSSION OF ALTERNATIVE MODELS

7.1. *Expected Utility Theory.* 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 (Bernoulli, 1738; Rothschild and Stiglitz, 1970). Under this assumption, however, EUT predicts that all assets with a nonpositive drift will be immediately sold, and it thus cannot explain why subjects start to gamble in our experiment (see Result 2 in Section 6).

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

7.2. *Cumulative Prospect Theory.* Abstracting from a finite expiration date, Ebert and Strack (2015) have shown that under empirically weak assumptions on the probability weighting function, a CPT agent will never stop an ABM, irrespective of how negative its drift is. This stark never-stopping result follows from the fact that the preference for positive skewness induced by common CPT-specifications is so strong that the naïve CPT agent can always find a stop-loss and take-profit strategy that is more attractive than not starting. As we numerically show in Online Appendix F, at the example of the representative CPT agent proposed by Tversky and Kahneman (1992),[18] the never-stopping result extends to processes with a finite expiration date. Consequently, common specifications of CPT can neither rationalize the fact that subjects stop a process with zero drift before the expiration date (Result 2) nor that stopping behavior is sensitive to the drift of the process (Result 3).[19] 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 1: because a CPT agent overweights the tails of a probability distribution, he likes the right-skewed distribution generated by loss-exit strategies (this has also been shown in Barberis, 2012; Ebert and Strack, 2015; Heimer et al., 2023).[20]

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

[19] CPT 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.

[20] The most striking difference between salience theory and alternative approaches—such as EUT and CPT—is that it predicts behavior to be context-dependent in the sense that the evaluation of a given option depends on the alternatives at hand. We ran an additional (preregistered) experiment that documents context-dependent stopping

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

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

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

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

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

or, equivalently,

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

A sufficient condition for Equation (5) to hold is that $(1-p)(b - x_0) \leq p(x_0 - a)$ or, equivalently, $\mathbb{E}[X_{\tau_{a,b}}] \leq x_0$, which is satisfied for any process with a nonpositive drift. We conclude that an agent with realization utility à la Barberis and Xiong (2012) would immediately sell any asset in our experiment; that is, their model can neither account for Result 2 nor Result 3.[23]

More generally, the preceding analysis highlights that some form of nonlinear probability weighting is necessary to explain our results on skewness seeking, not only in the dynamic

behavior in line with salience theory. To focus on our main results, and not to disrupt the flow of the main text, we decided to relegate this additional experiment to Online Appendix G, however.

[21] Barberis 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.

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

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

selling decisions, but also in the static choices, which we analyze in Online Appendix E (see Figure 14). The former point is made in an informal way also in Heimer et al. (2023). Adding nonlinear probability weighting to the model by Barberis and Xiong (2012) would yield a model that is essentially equivalent to the ones studied in Barberis (2012) or Ebert and Strack (2015), which we have already discussed in detail in the previous subsection.

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

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

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

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

As illustrated in Gul (1991), we need to assume $\beta > 0$ 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 nonpositive drift immediately, which contradicts both our results on dynamic (Result 2) and static choices (see Figure 14 in Online Appendix E).

So, let us assume not only that $\beta > 0$, 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 (Result 2). All the preceding arguments carry over to the case of a finite expiration date. In sum, we conclude that a model of disappointment aversion cannot coherently explain the findings on skewness seeking in static and dynamic settings.

8. CONCLUSION

Although we find that people take up symmetric gambles if they can obtain skewed return distributions through the choice of their stopping strategies, theoretical considerations suggest similar behavior in case the underlying process is skewed itself. On the one hand, even if the

[24] Disappointment 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).

underlying process is negatively skewed, the return distribution associated with the "right" loss-exit strategy is again positively skewed (Ebert, 2020). And if the process itself is positively skewed—which is indeed the case in many real-world applications—our results are likely to be amplified.

A first example refers to processes underlying many casino gambles (as discussed in Ebert and Strack, 2015, Section V) and many asset values, which are not symmetric, but positively skewed. Skewness seeking, as modeled by salience theory, then suggest that consumers gamble or overinvest 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 about teenagers or young adults who decide whether to pursue the career of a professional athlete, actor, or musician. Although the probability of actually making it to the professional level is small, it requires substantial investments of time and other resources to take the shot at becoming a superstar. A teenager who practices excessively for a particular sport, for instance, might as a result neglect school or studies, thereby lowering the attainable wage in the likely case that he fails to become a professional athlete. Now suppose that, as suggested by our model, this teenager adopts the following strategy: each year, he hopes for a breakthrough, but plans to quit on sports and instead study otherwise. This strategy generates a positively skewed return distribution, which can be particularly appealing due to the skewness that is inherent to the process of becoming a superstar. After each failure, however, the teenager revises his plans and decides to try it for *one more year*, as this way he can again experience a right-skewed distribution of returns. This idea of excessively pursuing a career is not only consistent with our model, but it is also supported by empirical studies (Choi et al., 2017; Grove et al., 2021). A similar type of argument applies to the classical problem of searching for a job, one of our introductory examples from the classical stopping literature. Here, skewness seeking can explain why people pass on too many mediocre jobs, thereby forgoing pay over a longer time horizon, in the hope of finding one of very few outstanding jobs with excellent pay. Also, in this example, the skewness of the return distribution that results from the chosen stopping strategy is complemented by the skewness of the process itself. In sum, skewness seeking can explain time-inconsistent behavior in trying to reach an elusive goal.

ACKNOWLEDGMENTS

Open access funding enabled and organized by Projekt DEAL.

DATA AVAILABILITY STATEMENT The data that support the findings of this study are openly available in the Open Science Foundation data repository at https://www.doi.org/10.17605/OSF.IO/5BJQW.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Figure 7: Skewness of the Outcome Distribution Induced by $\tau_{60,180}$
Figure 8: Salience-Weighted Utility of Gambling over Time
Table 2: Lotteries Used to Elicit Skewness Seeking in Static Choices
Figure 9: The Figure Illustrates the Layout of the Static Choices in the Experiment (in German)
Figure 10: The Left Panel Depicts the Share of Subjects Holding the Asset until the Expiration Date, Separately for the Different Drifts
Figure 11: The Left Panel Depicts the Share of Subjects Holding the Asset until the Expiration Date, Separately for the Different Drifts and Below- and Above-Median Subjects in Terms of Cognitive Skills

REFERENCES

BARBERIS, N., "A Model of Casino Gambling, " *Management Science* 58 (2012), 35–51.

———, and W. XIONG, "What Drives the Disposition Effect? An Analysis of a Long-Standing Preference-Based Explanation, " *Journal of Finance* 64 (2009), 751–84.

———, and ———, "Realization Utility, " *Journal of Financial Economics* 104 (2012), 251–71.

BERNOULLI, D., "Specimen Theoriae Novae de Mensura Sortis," *Comentarii Academiae Scientiarum Imperialis Petropolitanae* 5 (1738), 175–92 [translated by L. Sommer (1954) in *Econometrica*, 22, 23–36].

BORDALO, P., N. GENNAIOLI, and A. SHLEIFER, "Salience, " *Annual Review of Economics* 14 (2022), 521–44.

———, NICOLA GENNAIOLI, and ———, "Salience Theory of Choice under Risk, " *Quarterly Journal of Economics* 127 (2012), 1243–85.

CERREIA-VIOGLIO, S., D. DILLENBERGER, and P. ORTOLEVA, "Cautious Expected Utility and the Certainty Effect, " *Econometrica* 83 (2015), 693–728.

CHEN, D., M. SCHONGER, and C. WICKENS, "oTree - An Open-Source Platform for Laboratory, Online, and Field Experiments, " *Journal of Behavioral and Experimental Finance* 9 (2016), 88–97.

CHOI, D., D. LOU, and A. MUKHERJEE, "Superstar Firms and College Major Choice," Discussion Paper No. 12296, Centre for Economic Policy Research, 2017.

DERTWINKEL-KALT, M., and M. KÖSTER, "Salience and Skewness Preferences, " *Journal of the European Economic Association* 18 (2020), 2057–107.

DHAMI, S., *The Foundations of Behavioral Economic Analysis* (Oxford: Oxford University Press, 2016).

DURAJ, J., "Optimal Stopping with General Risk Preferences," Discussion Paper, Department of Economics, University of Pittsburgh, 2020.

EBERT, S., "On Skewed Risks in Economic Models and Experiments, " *Journal of Economic Behavior & Organization* 112 (2015), 85–97.

———, "On Taking a Skewed Risk More Than Once," Discussion Paper, Heidelberg University -, 2020.

———, and P. STRACK, "Until the Bitter End: On Prospect Theory in a Dynamic Context, " *American Economic Review* 105 (2015), 1618–33.

———, and ———, "Never, Ever Getting Started: On Prospect Theory Without Commitment," Discussion Paper, Heidelberg University & Yale University, 2018.

FISCHBACHER, U., G. HOFFMANN, and S. SCHUDY, "The Causal Effect of Stop-Loss and Take-Gain Orders on the Disposition Effect, " *Review of Financial Studies* 30 (2017), 2110–29.

GREINER, B., "Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE, " *Journal of the Economic Science Association* 1 (2015), 114–25.

GROVE, W. A., M. JETTER, and K. L. PAPPS, "Career Lotto? Labor Supply in a Superstar Market, " *Journal of Economic Behavior & Organization* 183 (2021), 362–76.

GUL, F., "A Theory of Disappointment Aversion, " *Econometrica* 59 (1991), 667–86.

HE, X., S. HU, J. OBLOJ, and X. ZHOU, "Optimal Exit Time from Casino Gambling: Strategies of Pre-Committed and Naive Gamblers, " *SIAM Journal on Control and Optimization* 57 (2019), 1854–68.

HEIMER, R., Z. ILIEWA, A. IMAS, and M. WEBER, "Dynamic Inconsistency in Risky Choice: Evidence from the Lab and Field," https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3600583 (2023).

HENDERSON, V., D. HOBSON, and S. ALEX, “Randomized Strategies and Prospect Theory in a Dynamic Context, ” *Journal of Economic Theory* 168 (2017), 287–300.

EBERT, S., and P. STRACK, “Probability Weighting, Stop-Loss and the Disposition Effect, ” *Journal of Economic Theory* 178 (2018), 360–97.

HUANG, Y.-J., A. NGUYEN-HUU, and X. Y. ZHOU, “General Stopping Behaviors of Naïve and Noncommitted Sophisticated Agents, with Application to Probability Distortion, ” *Mathematical Finance* 30 (2020), 310–40.

IMAS, A., “The Realization Effect: Risk-Taking after Realized versus Paper Losses, ” *American Economic Review* 106 (2016), 2086–109.

KAHNEMAN, D., and A. TVERSKY, “Prospect Theory: An Analysis of Decision Under Risk, ” *Econometrica* 47 (1979), 263–91.

KARNI, E., and Z. SAFRA, “Behaviorally Consistent Optimal Stopping Rules, ” *Journal of Economic Theory* 51 (1990), 391–402.

LUSARDI, A., and O. MITCHELL, *Financial Literacy and Planning: Implications for Retirement Wellbeing* (Oxford: Oxford University Press, 2011).

MACHINA, M. J., “Dynamic Consistency and Non-Expected Utility Models of Choice under Uncertainty, ” *Journal of Economic Literature* 27 (1989), 1622–68.

NIELSEN, K., “Dynamic Risk Preferences under Realized and Paper Outcomes, ” *Journal of Economic Behavior & Organization* 161 (2019), 68–78.

ODEAN, T., “Are Investors Reluctant to Realize Their Losses?, ” *Journal of Finance* 53 (1998), 1775–98.

PRIMI, C., K. MORSANYI, F. CHIESI, M. A. DONATI, and J. HAMILTON, “The Development and Testing of a New Version of the Cognitive Reflection Test Applying Item Response Theory (IRT), ” *Journal of Behavioral Decision Making* 29 (2016), 453–69.

QUIGGIN, J., “A Theory of Anticipated Utility, ” *Journal of Economic Behavior & Organization* 3 (1982), 323–43.

RESNICK, I., N. S. NEWCOMBE, and T. F. SHIPLEY, “Dealing with Big Numbers: Representation and Understanding of Magnitudes Outside of Human Experience, ” *Cognitive Science* 41 (2017), 1020–41.

ROTHSCHILD, M., and J. E. STIGLITZ, “Increasing Risk: I. A Definition, ” *Journal of Economic Theory* 2 (1970), 225–43.

SCHKADE, D. A., and D. KAHNEMAN, “Does Living in California make People Happy? A Focusing Illusion in Judgments of Life Satisfaction, ” *Psychological Science* 9 (1998), 340–46.

SHEFRIN, H., and M. STATMAN, “The Disposition to Sell Winners Too Early and Ride Losers Too Long: Theory and Evidence, ” *Journal of Finance* 40 (1985), 777–90.

STRACK, P., and P. VIEFERS, “Too Proud To Stop: Regret in Dynamic Decisions, ” *Journal of the European Economic Association* 19 (2021), 165–99.

TVERSKY, A., and D. KAHNEMAN, “Advances in Prospect Theory: Cumulative Representation of Uncertainty, ” *Journal of Risk and Uncertainty* 5 (1992), 297–323.

WEBER, M., and C. F. CAMERER, “The Disposition Effect in Securities Trading: An Experimental Analysis, ” *Journal of Economic Behavior & Organization* 33 (1998), 167–84.

XU, Z., and X. ZHOU, “Optimal Stopping under Probability Distortion, ” *Annals of Applied Probability* 23 (2013), 251–82.