Pre-Analysis Plan:
Optimal Stopping in a Dynamic Salience Model

Markus Dertwinkel-Kalt* Jonas Frey† Mats Köster‡

January 2020

1 Aim of the Study

Models of non-linear probability weighting, such as cumulative prospect theory (henceforth: CPT; Tversky and Kahneman, 1992), predict time-inconsistent behavior when applied to a dynamic context (e.g., Machina, 1989). While time-inconsistent behavior is indeed widespread, common specifications of CPT have (too) extreme implications in certain setups (e.g., Ebert and Strack, 2015, 2018). In this study, we compare the implications of (exogeneous) probability weighting as assumed in CPT to those of (endogenous) probability weighting as proposed in salience theory of choice under risk (Bordalo *et al.*, 2012), under the (testable) assumption that the decision maker is naive about his time-inconsistency.

We propose a dynamic salience model to study the choice of when to optimally stop an arithmetic brownian motion with non-positive drift. Our salience model predicts that the optimal stopping behavior is sensitive to the drift of the process; namely, a naive agent will gamble if the drift of the process is slightly negative, but will stop immediately if the drift becomes too negative. This prediction is arguably more plausible than those of CPT, which under common specifications predicts excessive gambling irrespective of the drift of the process, and EUT (with a concave utility function), which predicts no gambling at all.

2 A Dynamic Version of Salience Theory of Choice under Risk

2.1 Static Model

Consider an agent choosing from a choice set $\mathcal{C}$ that contains exactly two non-negative random variables (or *lotteries*), X and Y. Let $S \subseteq \mathbb{R}^2_{\geq 0}$ be the support of the joint distribution of X and Y. We denote the corresponding joint cumulative distribution function as F. Moreover, if a random variable (or lottery) is degenerate, we call it a *safe* option.

According to salience theory of choice under risk (Bordalo *et al.*, 2012), the agent evaluates a random variable by assigning a subjective probability to each state of the world $s \in S$ that

*Frankfurt School of Finance & Management. Email: m.dertwinkel-kalt@fs.de.
†Said Business School, University of Oxford. Email: Jonas.Frey.DPHIL2017@said.oxford.edu.
‡HHU Düsseldorf (DICE). Email: mats.koester@hhu.de.

depends on the state's objective probability and on its salience. The salience of a state is assessed by a *salience function*, which is defined as follows.

Definition 1. *A symmetric, bounded, and absolutely continuous 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:*

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

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

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

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

We say that a given state of the world $(x, y) \in S$ is the more salient the larger its salience value $\sigma(x, y)$ is. Ordering implies that a state is the more salient the more the attainable outcomes in this state differ. In this sense, ordering captures the well-known *contrast effect* (e.g., Schkade and Kahneman, 1998), whereby large contrasts (in outcomes) attract a great deal of attention. Diminishing sensitivity reflects *Weber's law* of perception and it implies that the salience of a state decreases if the outcomes in this state uniformly increase.[1]

A *salient thinker* evaluates monetary outcomes via a linear value function, $u(x) = x$, and chooses from the set $\mathcal{C} = \{X, Y\}$ as to maximize his *salience-weighted utility* defined as follows, whereby the salience-weighted probabilities are normalized so that they sum to one (e.g., Bordalo *et al.*, 2012; Dertwinkel-Kalt and Köster, forthcoming).

Definition 2. *The* salience-weighted utility *of a lottery X evaluated in the choice set $\mathcal{C} = \{X, Y\}$ is*

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

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

2.2 Dynamic Model

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

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

with an intial value $X_0 = x$, a drift $\mu \in \mathbb{R}$, a volatility $\nu \in \mathbb{R}_{>0}$, and a standard Brownian Motion $(W_t)_{t \in \mathbb{R}_{\geq 0}}$. Following Ebert and Strack (2015, 2018) we abstract from discounting.

[1] Bordalo *et al.* (2012) also allow for lotteries with negative outcomes and therefore add a third property of a salience function to ensure that diminishing sensitivity (with respect to zero) reflects to the negative domain: by *reflection*, for any $w, x, y, z \geq 0$, we have $\sigma(x, y) > \sigma(w, z)$ if and only if $\sigma(-x, -y) > \sigma(-w, -z)$.

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

The set of stopping strategies. The central feature of a Markov diffusion is that the distribution of future wealth does not depend on the history of the process. Following Ebert and Strack (2018), we thus restrict attention to *Markovian* (*pure*) *strategies*, according to which the agent conditions the decision whether to stop at time t only on the payoff-relevant variables: the current wealth level, X_t, and the time distance to the expiration date, $T - t$. A (*stopping*) *strategy* $s : \mathbb{R}_{\geq 0} \times [0, T] \to \{\text{stop}, \text{ continue}\}$ is a deterministic function of the payoff-relevant variables.

For tractability, we derive certain results under the restriction to *time-invariant* strategies, which satisfy $s(y, T - t) = s(y)$ for any wealth level $y \in \mathbb{R}_{\geq 0}$ and any point in time $t \in [0, T)$; that is, irrespective of the distance to the expiration date, the current wealth level determines the agent's stopping decision. For a time-invariant strategy, s, there exists a single stopping set

$$\mathcal{S} := \{y \in \mathbb{R}_{\geq 0} : s(y) = \text{stop}\} \subseteq \mathbb{R}_{\geq 0}$$

such that an agent playing according to strategy s stops at time $t \in [0, T)$ if and only if $X_t \in \mathcal{S}$. Thus, conditional on stopping before the expiration date T, following stopping strategy s at time t is equivalent to *choosing* a distribution over future wealth with support $\mathcal{S}$.

Notice that the process $(X_t)_{t \in \mathbb{R}_{\geq 0}}$ has continuous paths by assumption. Thus at any time t, conditional on stopping before the expiration date, any time-invariant strategy s can be represented as a binary lottery with outcomes $a_t = \sup\{y \in \mathcal{S} : y \leq x_t\}$ and $b_t = \inf\{y \in \mathcal{S} : y \geq x_t\}$ (Ebert and Strack, 2018, p. 13). But, depending on the thresholds a_t and b_t, the agent may not stop before the expiration date T, which gives rise to a random wealth level $X_{T \wedge \tau_{a_t, b_t}}$ with

$$\tau_{a_t, b_t} = \inf\{r > t : s(X_r) = \text{stop}\} = \inf\{r > t : X_r \notin (a_t, b_t)\}$$

being the first leaving time of the interval (a_t, b_t).

Solution concept. Since non-linear probability weighting implies that an agent's optimal strategy at time t might no longer be optimal at some later point in time (e.g., Machina, 1989), the stopping behavior under salience theory depends on whether the salient thinker is aware of this time-inconsistency or not. We follow Ebert and Strack (2015) and assume that the agent is naive about his time-inconsistency.

As in Ebert and Strack (2015) we assume that "at every point in time the naive [salient thinker] looks for a [...] strategy s that brings her higher [salience-weighted utility] than stopping immediately. If such a strategy exists, [he] holds on to the investment—irrespective of [his] earlier plan." We thereby assume that the naive salient thinker evaluates each stopping strategy in isolation, meaning that he compares it only to the alternative of stopping immediately. While this assumption entails a loss of generality, we do not regard it as far-fetched, because, at a given point in time t, the agent simply decides whether to stop the process or not, and only conditional on not stopping the process the exact stopping strategy matters. We think that the agent's reasoning on whether to stop at time t is well captured by comparing a given stopping strategy to the alternative of stopping immediately. Assuming that the naive salient thinker gambles if and only if he strictly prefers to do so, the naive decision rule then reads as follows.

Definition 3 (Naive Decision Rule)**.** *A naive salient thinker* continues at time *t with a current wealth level x_t if there exists a stopping time τ_{a_t,b_t} with $a_t < x_t < b_t$, such that he chooses the random variable $X_{T\wedge\tau_{a_t,b_t}}$ from the set $\{X_{T\wedge\tau_{a_t,b_t}}, x_t\}$. Otherwise, the naive salient thinker* stops at time *t.*

3 Stopping Behavior of a Naive Salient Thinker

We characterize the stopping behavior of a naive salient thinker by translating the dynamic decision when to stop the process into a static choice between a (binary) lottery and a safe option that pays the current wealth level. More precisely, we start out from the case of a fair gamble and ask whether there exists some time-invariant stopping strategy that is more attractive than stopping immediately. Subsequently, we extend our findings to processes with a slightly negative drift. In a last step, we analyze the limits of naive gambling under salience theory.

Naive gambling under salience theory. Consider a salient thinker with a current wealth level of x_t. For the sake of illustration, let $T = \infty$, in which case any pure Markov strategy is time-invariant, and let $\mu = 0$, which implies that the gamble is "fair". Suppose that at time t the salient thinker chooses a stopping strategy s that can be represented by the stopping time $\tau_{a,b}$ and therefore induces a binary lottery over wealth, $X_{\tau_{a,b}} = (a, p; b, 1-p)$, where $p = \frac{b-x}{b-a}$ gives the probability that the downside payoff a is realized. Since the process has zero drift by assumption, the expected value of this binary lottery is given by $\mathbb{E}[X_{\tau_{a,b}}] = x_t$.

By construction, if both states are equally salient, the salient thinker behaves as if he was risk neutral; that is, he is indifferent between the lottery $X_{\tau_{a,b}}$ and the safe option x_t. In fact, the salient thinker strictly prefers the binary lottery $X_{\tau_{a,b}}$ over the safe option paying the lottery's expected value x_t if and only if the lottery's upside payoff, b, is more salient than its downside payoff, a; that is, if and only if $\sigma(b, x_t) > \sigma(a, x_t)$ holds. Due to $\sigma(b, x_t) > \sigma(x_t, x_t)$ as well as continuity of the salience function, we can find, for any x_t, outcomes a and b so that the salient thinker strictly prefers the binary lottery $X_{\tau_{a,b}}$ over the safe option x_t. Since the salient thinker can choose a and b independently of each other, he can always find a stopping time $\tau_{a,b}$ that yields a strictly higher salience-weighted utility than stopping immediately. Hence, by Definition 3, if $T = \infty$, a naive salient thinker never stops a process with zero drift. Since the salience-weighted utility is continuous in the probability $p = p(a, b, \mu)$ that the lower outcome a is realized, which in turn is continuous in the drift μ (see Lemma 1 in the Appendix), the preceding arguments (partly) extend to processes with a slightly negative drift. The following proposition further shows that the results still hold in setups with a finite expiration date.

Proposition 1. *Fix an initial wealth level $x \in \mathbb{R}_{>0}$ and an expiration date $T \in \mathbb{R}_{>0} \cup \{\infty\}$.*

- (a) *If the drift is non-negative, then the naive salient thinker never stops before the expiration date T.*
- (b) *For any initial wealth x and any volatility ν, there is a constant $\hat{\mu} < 0$ so that for any drift $\mu > \hat{\mu}$, a naive salient thinker starts to gamble* (*but does not necessarily continue until the expiration date*)*.*

The role of skewness in naive gambling. Dertwinkel-Kalt and Köster (forthcoming) show that in static choices a salient thinker takes a binary risk (over its expected value) if and only if it is sufficiently right-skewed (or positively skewed); that is, if and only if the binary lottery offers a large, but unlikely upside and a likely downside that is close to the expected value. A preference for positive skewness is also what drives a salient thinker's dynamic gambling behavior. To make this point explicit, we introduce the notion of a *naively right-skewed* strategy.

Definition 4. *A time-invariant stopping strategy $s(y,t) = s(y)$ is* naively right-skewed *at time t if and only if the corresponding stopping time $\tau_{a,b}$ satisfies $b - x_t > x_t - a$.*[2]

Similar to the case of static choices, as analyzed in Dertwinkel-Kalt and Köster (forthcoming), ordering and diminishing sensitivity together imply that a salient thinker gambles according to a certain stopping strategy only if it is naively right-skewed. As an illustration, consider again the case without an expiration date and assume a zero drift. If $b - x_t < x_t - a$, then not only is the payoff level lower in the state (a, x_t) compared to the state (b, x_t), but also the contrast in outcomes is larger in this state. Hence, the downside a is more salient than the upside b, which makes this stopping strategy unattractive to a salient thinker. The argument extends to processes with a negative drift as well as to setups with a finite expiration date. Conversely, due to diminishing sensitivity, $b - x_t > x_t - a$ does not imply that the upside b is more salient than the downside a, so that a salient thinker does not choose any naively right-skewed strategy.

Proposition 2. *A naive salient thinker chooses a time-invariant strategy only if it is naively right-skewed.*

On the limits of naive gambling. Next, we explore the limits on the gambling behavior of a naive salient thinker, under the restriction to time-invariant strategies.

Proposition 3. *Suppose that the agent can only choose time-invariant strategies. Then, for any initial wealth level $x \in \mathbb{R}_{>0}$ and any volatility ν, there exists a constant $\tilde{\mu} < 0$ such that for any drift $\mu < \tilde{\mu}$ a naive salient thinker stops immediately.*

Since the salience function is bounded, the (naive) skewness created by a time-invariant stopping time $\tau_{a,b}$ is not enough to make up for the fact that the expected value will be close to a for a sufficiently negative drift, irrespective of the initial wealth level x. As a consequence, a naive salient thinker immediately stops any process with a sufficiently negative drift.

If we assume that the expiration date is sufficiently large, our model makes the stronger prediction that a salient thinker's stopping behavior depends on the drift in a montonic way.

Proposition 4. *Suppose that the agent can only choose time-invariant strategies, and fix an initial wealth level $x \in \mathbb{R}_{>0}$ as well as a volatility ν. For any two drifts μ and μ' with $0 \geq \mu' > \mu$, there exists some threshold value $\hat{T} \in \mathbb{R}_{>0}$ such that, for any expiration date $T \geq \hat{T}$, a naive salient thinker immediately stops the process with drift μ if he does so for the process with drift μ'.*

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

To establish the intuition for this stronger result, we start again with the case of no expiration date. Recall that, for $T = \infty$, any stopping time $\tau_{a,b}$ induces a binary distribution over wealth, $X_{\tau_{a,b}}$, and notice that an increase in the drift μ improves this binary distribution in terms of first-order stochastic dominance. As a consequence, a salient thinker's certainty equivalent to $X_{\tau_{a,b}}$ monotonically increases in the drift of the process (Proposition 1 in Dertwinkel-Kalt and Köster, forthcoming). If we allow for $T < \infty$, we observe that also the distribution of X_T, conditional on reaching the expiration date with the stopping time $\tau_{a,b}$, is monotonic (in terms of first-order stochastic dominance) in the drift, which makes the result extend to finite expiration dates.

4 An Experiment on Dynamic Gambling Behavior

4.1 Experimental Design and Implementation

We conduct a lab experiment in which subjects have to decide at which price to sell an asset. Subjects make their selling decisions in (approximately) continuous time and they can hold the asset for a maximum duration of 10 seconds. If a subject does not sell the asset within 10 seconds, it is automatically sold at the price reached at the expiration date. The selling price follows an ABM with a drift parameter $\mu \in \{0, -0.1, -0.3, -0.5, -1, -2\}$ and a volatility $\nu = 5$, and it is updated every tenth of a second, which implies $T = 100$.[3] The asset's initial price is equal to $x = 100$ Taler, an experimental currency that is converted into Euros at a ratio of 10:1.

We restrict the strategy space to all time-invariant strategies, as it is illustrated in Figure 1: at every point in time subjects can choose an upper and a lower stopping threshold. Once the price of the asset reaches either threshold, subjects can decide whether to sell the asset at this price or to adjust the thresholds and continue the process (see the lower left panel). In addition, subjects can pause the process at any point in time to adjust the thresholds (see upper right panel). But, importantly, subjects can set only one upper bound and one lower bound at a time and thus observe each stopping strategy in isolation. Before starting the process, subjects can further decide to sell the asset immediately (see the upper left panel). At the beginning, the upper and lower bound are centered symmetrically around the initial price. In order to start the process, subjects have to move each bound at least once.

Overall, subjects make six selling decisions, one decision for each of the drift parameters. The order of drifts is randomized at the subject level. To explain the drift of an ABM to the subjects, they have to draw three sample paths from the underlying process and, in addition, they see an overview of ten additional sample paths of this process before making a decision (Figure 2). The sample paths are randomly drawn at the subject level; that is, different subjects see different sample paths of the same underlying process.

After making the six selling decisions, subjects make twelve choices between a binary lottery and a safe option paying the lottery's expected value. The binary lotteries are the same as the ones used in Experiment 1 of Dertwinkel-Kalt and Köster (forthcoming): more specifically, we use two sets of lotteries with the same expected value and the same variance, but different levels

[3] Notice that the drift of an ABM is additive over time. To help subjects understand what the drift of a process is, we present them with aggregated drifts per second $\mu \in \{0, -1, -3, -5, -10, -20\}$ in the experiment.

Figure 1: *Screenshots of the decision screen for the process with zero drift (in German). The red lines indicate the upper and lower stopping bounds. 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 displays the final selling price.*

of skewness. The order of lotteries is randomized at the subject level. Finally, subjects answer five CRT-questions and five financial literacy questions.

At the end of the experiment, one of the six selling decisions will be randomly drawn by the computer to be payoff-relevant. In addition, we will randomly draw one subject in each session for whom also one of the twelve binary choices is randomly chosen to be payoff-relevant. Subjects will be further rewarded for correctly answered CRT and financial literacy questions (1 Taler per correctly answered question). All subjects receive an additional 4 Euros for their participation in the experiment.

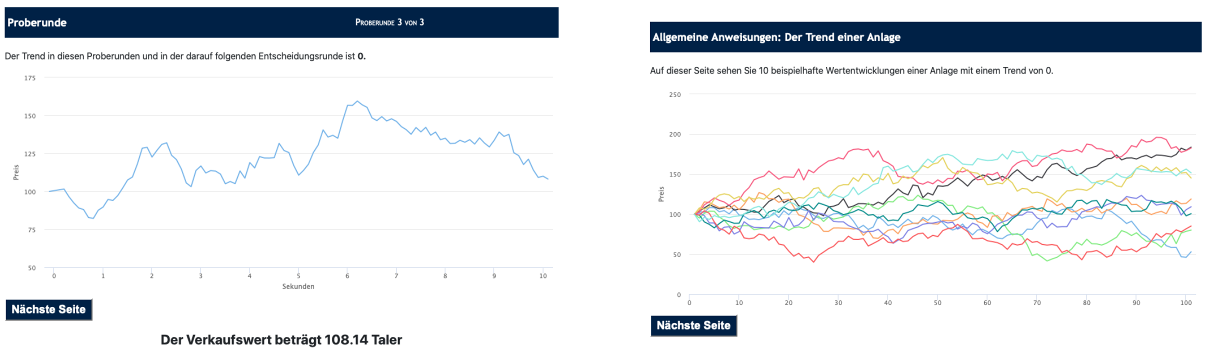

Figure 2: *Screenshots of the sampling screens for the process with zero drift (in German).*

We plan to conduct 5 sessions with a total number of $n = 150$ subjects. The sessions will take place in January/February 2020 at the experimental laboratory of the University of Cologne.

4.2 Salience Predictions

In this section we specify the predictions that guided our experimental setup. The first three predictions follow directly from our model, under the assumption that subjects are heterogeneous with respect to their salience function. The last prediction is not a direct implication.

Within-subjects predictions. By Proposition 1 (a), a naive salient thinker will never stop a process with zero drift before the expiration date. This gives rise to our first prediction:

Prediction 1. *If $\mu = 0$, subjects hold the asset until the expiration date.*

Proposition 2 further implies that a naive salient thinker will only gamble according to naively right-skewed strategies. This gives rise to our second prediction:

Prediction 2. *Subjects choose naively right-skewed strategies.*

Between-subjects predictions. By Propositions 3 and 4, we expect a monotonic relationship between the drift of the process and a subject's stopping behavior. Assuming that the subjects are heterogeneous with respect to their salience functions, we obtain our third prediction:

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

Since our model extends a theory of static choice under risk to a dynamic setup, we will further study the relationship between the subjects' static and dynamic risk preferences. For that, we classify the choice of a binary lottery in one of the twelve static decisions, which each subject has to make, as a *skewness seeking choice* if and only if this lottery is right-skewed.

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

Testing the naivete assumption. As we show in Appendix B, our design allows us to test the assumption of naivete about the own time-inconsistency within the salience framework. A sophisticated salient thinker (without committment), who can only choose time-invariant strategies, immediately sells any asset with a non-positive drift. Hence, naivete is a necessary assumption to rationalize gambling (within the salience framework) in our experiment.

4.3 Analysis

Descriptive analysis. Based on Predictions 1 and 2, we will classify subjects into three groups: For each subject we count how many of the choices in the six selling decisions are consistent with salience theory, CPT, and EUT. We then divide subjects into three groups depending on which of the theories best predicts their behavior (in the sense of maximizing the number of choices that are consistent with a theory). We are interested in the share of subjects whose behavior is best explained by salience theory, compared to the alternative models. As a robustness check, we will restrict attention to the subset of subjects who behave in line with one of the three theories in *all* choices (in particular, all strategies they choose) they make in the six selling tasks.

Statistical analysis. In the following, we denote the process with the k-th largest drift as process $k \in \{1, 2, \dots, 6\}$ and we index subjects by $i \in \{1, 2, \dots, n\}$.

To test for Prediction 3, we will construct for each subject i and each process k a binary indicator, $S_{i,k}$, that takes a value of one if the subject has stopped the process immediately and a value of zero otherwise. Then, using OLS, we regress $S_{i,k}$ on a continuous variable μ_k that indicates the drift of the process. We cluster the standard errors at the subject level. According to Prediction 3, the coefficient on the drift should be negative and statistically significant.

To test for Prediction 4, we will calculate, for each subject i, the share of naively right-skewed strategies, NRS_i, chosen in the six (dynamic) selling decisions. In addition, we calculate for each subject i the share of skewness seeking choices (as defined in the preceding subsection), SSC_i, in the twelve static decisions. We then regress NRS_i on SSC_i using OLS, with the standard errors being clustered at the subject level. If static and dynamic skewness preferences have the same driver, we expect the coefficient on SSC_i to be positive and statistically significant. As a robustness check, we run the same regression on the subsample of subjects who reveal in the static decisions a unique switching point, consistent with salience theory, for each set of six lotteries with the same expected value and the same variance (i.e., we use only those subjects who choose, for a given expected value and variance, a binary lottery over its expected value if and only if it is sufficiently right-skewed).

References

Bordalo, P., Nicola Gennaioli and Shleifer, A. (2012). Salience theory of choice under risk. *Quarterly Journal of Economics*, **127** (3), 1243–1285.

Cox, D. R. and Miller, H. D. (1977). *The theory of stochastic processes*. Chapman and Hall Ltd.

Dertwinkel-Kalt, M. and Köster, M. (forthcoming). Salience and skewness preferences. *Journal of the European Economic Association*.

Ebert, S. (2015). On skewed risks in economic models and experiments. *Journal of Economic Behavior & Organization*, **112**, 85–97.

— and Strack, P. (2015). Until the bitter end: On prospect theory in a dynamic context. *American Economic Review*, **105** (4), 1618–1633.

— and — (2018). Never, ever getting started: On prospect theory without commitment. `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2765550`, unpublished.

Karni, E. and Safra, Z. (1990). Behaviorally consistent optimal stopping rules. *Journal of Economic Theory*, **51** (2), 391–402.

Machina, M. J. (1989). Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature*, **27** (4), 1622–1668.

Revuz, D. and Yor, M. (1999). *Continuous martingales and Brownian motion*. Springer-Verlag, Berlin.

Schkade, D. A. and Kahneman, D. (1998). Does living in California make people happy? A focusing illusion in judgments of life satisfaction. *Psychological Science*, **9** (5), 340–346.

Tversky, A. and Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, **5** (4), 297–323.

Appendix A: Proofs

A.1: Preliminary Results on Arithmetic Brownian Motions

Lemma 1. *Fix a current wealth level $x_t \in \mathbb{R}_{>0}$ and a non-zero drift μ. Then, for any stopping time $\tau_{a,b}$ with $a < x_t < b$, we have*

$$\mathbb{P}_t[X_{\tau_{a,b}} = a] = \frac{\exp(-(2\mu/\nu^2)b) - \exp(-(2\mu/\nu^2)x_t)}{\exp(-(2\mu/\nu^2)b) - \exp(-(2\mu/\nu^2)a)}. \tag{1}$$

In particular, an increase in the drift of the process improves the distribution of $X_{\tau_{a,b}}$ in terms of first-order stochastic dominance.

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

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

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

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

which yields the claim. The second part follows immediately from taking the partial derivative of the right-hand side of Eq. (1) with respect to μ, which is strictly negative. □

Lemma 2. *Fix an initial wealth level $x \in \mathbb{R}_{>0}$ and an expiration date $T \in \mathbb{R}_{>0}$. Consider a stopping time $\tau_{a,b}$ with $a < x < b$ and take the perspective of period $t = 0$.*

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

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

where the integrand is given by

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

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

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

(*d*) *For any wealth level $y \in \mathbb{R}_{\geq 0}$ and point in time $t \in (0, T]$, the partial derivative $\frac{\partial}{\partial x}q(y,t|X_0 = x)$ exists and is bounded. Therefore, also $\frac{\partial}{\partial x}\mathbb{P}_0[\tau_{a,b} \geq T|X_0 = x]$ exists and is bounded.*

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

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

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

(*g*) $\lim_{\mu \to \infty} \mathbb{P}_0[X_T \le x | X_0 = x, \tau_{a,b} \ge T] = 0 = \lim_{\mu \to -\infty} \mathbb{P}_0[X_T > x | X_0 = x, \tau_{a,b} \ge T]$.

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

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

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

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

PART (b). Follows from $q(y,t|X_0 = x)$ being continuous in x and $\sin(n\pi) = 0$ for any $n \in \mathbb{Z}$.

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

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

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

PART (d). Straightforward computations show that, if $\frac{\partial}{\partial x}q(y,t|X_0 = x)$ exists, then

$$\begin{aligned}
\frac{\partial}{\partial x}q(y,t|X_0 = x) &= -\frac{\mu}{\nu^2}q(y,t|X_0 = x) \\
&\quad + \frac{2\exp\left(\frac{\mu(y-x)}{\nu^2}\right)\exp\left(-\frac{T}{2}\frac{\mu^2}{\nu^2}\right)}{(b-a)^2} \sum_{n=1}^{\infty} \left\{ n\pi \cos\left(\frac{\pi n(x-a)}{b-a}\right) \sin\left(\frac{\pi n(y-a)}{b-a}\right) \exp\left(-\frac{t}{2}\frac{n^2\pi^2\nu^2}{(b-a)^2}\right) \right\}.
\end{aligned}$$

Now, since $\left|\cos\left(\frac{\pi n(x-a)}{b-a}\right)\sin\left(\frac{\pi n(y-a)}{b-a}\right)\right| \le 1$ and since $\exp\left(-\frac{t}{2}\frac{n^2\pi^2\nu^2}{(b-a)^2}\right) \le \exp\left(-\frac{t}{2}\frac{n\pi\nu^2}{(b-a)^2}\right)$ and since $q(y,t|X_0 = x) \le 1$, it is sufficient to show that

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

As $\mathbb{P}_0[\tau_{a,b} \ge T | X_0 = x]$ is given by the integral of $q(y,t|X_0 = x)$ over the interval $[a,b]$ and as it is bounded by one, we conclude that $\frac{\partial}{\partial x}\mathbb{P}_0[\tau_{a,b} \ge T | X_0 = x]$ exists and that it is bounded.

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

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

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

or, equivalently, $\mathbb{E}_0[X_T|X_0 = x, \tau_{a,b} \ge T, X_T \le z] \le \mathbb{E}_0[X_T|X_0 = x, \tau_{a,b} \ge T]$, which is true.

PART (g). Follows immediately from Part (d) and the fact that $\lim_{\mu\to\infty}\exp\left(\frac{\mu(y-x)}{\nu^2}\right) = 0$ if and only if $y < x$, while $\lim_{\mu\to-\infty}\exp\left(\frac{\mu(y-x)}{\nu^2}\right) = 0$ if and only if $y > x$.

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

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

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

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

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

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

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

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

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

A.2: Stopping Behavior of a Naive Salient Thinker

Proof of Proposition 1. PART (a). Let $\mu \geq 0$. We have to show that *for any point in time $t < T$ with a current wealth level $X_t = x_t$ there exists some stopping time $\tau_{a,b}$ such that* $U^s\big(X_{\tau_{a,b} \wedge T} | \mathcal{C}\big) > x_t$. Since, for any $a > 0$, the process is a Submartingale by assumption (even conditional on stopping before the expiration date), we have $U^s\big(X_{\tau_{a,b} \wedge T} | \mathcal{C}, \tau_{a,b} < T\big) > x_t$ if $\sigma(b, x_t) > \sigma(a, x_t)$. It is easy to see that for any x_t and any $a \in (0, x_t)$ there exists some $b^* = b^*(a, x_t)$ such that $\sigma(b^*, x_t) = \sigma(a, x_t)$ and that, for any $b > b^*$, $\sigma(b, x_t) > \sigma(a, x_t)$ holds. Since $\sigma(a, x_t)$ decreases in a for any $a < x_t$ and since $\sigma(b, x_t)$ increases in b for any $b > x_t$, we know that $b^*(a, x_t)$ decreases in a. The claim follows from the fact that, for any fixed $t < T$, we have $\lim_{a \to x_t} \mathbb{P}_t[\tau_{a,b^*+\epsilon} < T] = 1$.

PART (b). Follows immediately from Part (a) and continuity. $\square$

Proof of Proposition 2. Consider a stopping strategy that corresponds to the stopping time $\tau_{a,b}$ with $a := x - \epsilon - \epsilon'$ and $b := x + \epsilon$ for some $\epsilon, \epsilon' > 0$ and $\epsilon + \epsilon' \leq x$, and that is therefore not naively right-skewed. Moreover, we denote as

$$\Phi_\mu(z) := \mathbb{P}_0[X_T \leq z | X_t = x_t, \tau_{a,b} \geq T]$$

the CDF of X_T conditional on reaching the expiration date when choosing the above strategy.

Denote $\mathbb{P}_t[\tau_{a,b} < T] := \mathbb{P}_t[\tau_{a,b} < T | X_t = x_t]$. Then, it follows that

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

where the probability $p = p(a, b, \mu)$ is defined as in Eq. (1), the first inequality follows from ordering and diminishing sensitivity as well as the construction of $\tilde{\Phi}_\mu$, which is defined as $\tilde{\Phi}_\mu(z) := \Phi_\mu(z)$ for any $z \geq x - \epsilon$ and $\tilde{\Phi}_\mu(z) := 0$ for any $z < x - \epsilon$, the second inequality follows from the fact that the drift of the process is non-positive and by diminishing sensitivity, and the weak inequality holds by Lemma 2 (i) and by diminishing sensitivity. $\square$

Proof of Proposition 3. The proof proceeds in two steps. In a first step, we assume $T = \infty$, where any Markov strategy implements a binary lottery. In a second step, we consider $T < \infty$.

1. STEP: Let $T = \infty$ and $\mu < 0$. We have to show that *there exists some $\tilde{\mu} \in \mathbb{R}$ such that for any $\mu < \tilde{\mu}$ no attractive stopping strategy exists; that is, for any $\mu < \tilde{\mu}$, it has to hold that*

$$\sup_{(a,b)\in[0,x)\times(x,\infty)} p(a,b,\mu)(a-x)\sigma(a,x) + (1-p(a,b,\mu))(b-x)\sigma(b,x) \leq 0. \tag{2}$$

Since the salience function is bounded (away from zero), we obtain

$$\begin{aligned}\sup_{(a,b)\in[0,x)\times(x,\infty)} & p(a,b,\mu)(a-x)\sigma(a,x) + (1-p(a,b,\mu))(b-x)\sigma(b,x) \\ & \leq \sup_{(a,b)\in[0,x)\times(x,\infty)} p(a,b,\mu)(a-x)\underline{\sigma} + (1-p(a,b,\mu))(b-x)\overline{\sigma},\end{aligned}$$

where $\overline{\sigma} := \sup_{(x,y)\in\mathbb{R}^2_{\geq 0}} \sigma(x,y)$ and $\underline{\sigma} := \inf_{(x,y)\in\mathbb{R}^2_{\geq 0}} \sigma(x,y)$. We proceed by distinguishing, for a fixed drift $\mu < 0$ and a fixed upper stopping threshold $b > x$, different cases depending on the lower stopping threshold $a < x$ that maximizes this upper bound on the difference in salience-weighted utility, $\Delta(a,b,\mu) := p(a,b,\mu)(a-x)\underline{\sigma} + (1-p(a,b,\mu))(b-x)\overline{\sigma}$.

<u>1. Case:</u> For the sake of a contradiction, let the upper bound $\Delta(a,b,\mu)$ be maximized at $a^* = a^*(b,\mu) \in (0,x)$. A necessary condition for this to be true is that $\frac{\partial}{\partial a}\Delta(a,b,\mu)\big|_{a=a^*(b,\mu)} = 0$ holds. This leaves us with a unique candidate for the optimal lower stopping threshold,

$$a^*(b,\mu) = -\frac{\nu^2}{2\mu}\left[\text{LambertW}\left(-\exp\left(-1+2\frac{\mu}{\nu^2}(b-x)(\overline{\sigma}/\underline{\sigma}-1)\right)\right)+1\right] + x + \frac{\overline{\sigma}}{\underline{\sigma}}(b-x), \tag{3}$$

where LambertW : $(-1/e,\infty) \to (-1,\infty)$ denotes the principle branch of the Lambert W function. For any $b > x$ and any $\mu < 0$, we have $a^*(b,\mu) > x$; a contradiction. Hence, $\Delta(a,b,\mu)$ is either maximized at $a = 0$ or in the limit of a approaching x. We analyze both cases separately.

<u>2. Case:</u> Suppose that $\Delta(a,b,\mu)$ is maximized at $a = 0$. Then, we have

$$\sup_{(a,b)\in[0,x)\times(x,\infty)} p(a,b,\mu)(a-x)\sigma(a,x) + (1-p(a,b,\mu))(b-x)\sigma(b,x) \leq \sup_{b\in(x,\infty)} \Delta(0,b,\mu).$$

Now it is sufficient to show that *there exists some $\tilde{\mu} \in \mathbb{R}$ such that for any $\mu < \tilde{\mu}$ it holds that* $\sup_{b\in(x,\infty)} \Delta(0,b,\mu) = \lim_{b\to x} \Delta(0,b,\mu) = 0$. Denote $\tilde{\mu} := -\frac{\nu^2}{2}\frac{\overline{\sigma}}{\underline{\sigma}}\frac{1}{x}$. Then, it follows that

$$\begin{aligned}\frac{\partial}{\partial b}\Delta(0,b,\mu) &= -\big(x\underline{\sigma} + (b-x)\overline{\sigma}\big)\left(\frac{\partial}{\partial b}p(0,b,\mu)\right) + \big(1-p(0,b,\mu)\big)\overline{\sigma} \\ &= \big(1-p(0,b,\mu)\big)\left[2\frac{\mu}{\nu^2}\big(x\underline{\sigma} + (b-x)\overline{\sigma}\big)\left(\frac{\exp\big(-2\frac{\mu}{\nu^2}b\big)}{\exp\big(-2\frac{\mu}{\nu^2}b\big)-1}\right) + \overline{\sigma}\right] \\ &< \underbrace{\big(1-p(0,b,\mu)\big)}_{>0}\underbrace{\left[2\frac{\mu}{\nu^2}x\underline{\sigma} + \overline{\sigma}\right]}_{<0 \text{ for } \mu<\tilde{\mu}},\end{aligned}$$

where the inequality holds by the fact that $\mu < 0$ and $b > x$. This proves the claim.

3. Case: Suppose a maximizer does not exists: $\sup_{a\in[0,x)}\Delta(a,b,\mu)=\lim_{a\to x}\Delta(a,b,\mu)=0$, where the second equality follows from the fact that $\lim_{a\to x}p(a,b,\mu)=1$. But then we have

$$\begin{aligned}\sup_{(a,b)\in[0,x)\times(x,\infty)} & p(a,b,\mu)(a-x)\sigma(a,x)+(1-p(a,b,\mu))(b-x)\sigma(b,x)\\ &= \sup_{b\in(x,\infty)}\sup_{a\in[0,x)} p(a,b,\mu)(a-x)\sigma(a,x)+(1-p(a,b,\mu))(b-x)\sigma(b,x)\\ &\le \sup_{b\in(x,\infty)}\sup_{a\in[0,x)}\Delta(a,b,\mu)=0,\end{aligned}$$

irrespective of the drift of the process, which proves our claim.

2. STEP: Let $T<\infty$ and fix an initial wealth $x>0$. Suppose the agent is restricted to time-invariant strategies. Let $\tau_{a,b}$ be a stopping time with $a<x<b$. As before, denote as

$$\Phi_\mu(z):=\mathbb{P}_0[X_T\le z|X_0=x,\tau_{a,b}\ge T]$$

the cumulative distribution function of X_T conditional on reaching the expiration date.

In order to prove the statement, we have to show that *there exists some $\tilde{\mu}\in\mathbb{R}$ such that for any $\mu<\tilde{\mu}$ no attractive stopping strategy exists; that is, for any $\mu<\tilde{\mu}$, it has to hold that*

$$\sup_{(a,b)\in[0,x)\times(x,\infty)}\Big\{\mathbb{P}_0[\tau_{a,b}<T]\cdot\big[p(a-x)\sigma(a,x)+(1-p)(b-x)\sigma(b,x)\big]\\ +\mathbb{P}_0[\tau_{a,b}\ge T]\cdot\int_{(a,b)}(z-x)\sigma(z,x)\,d\Phi_\mu(z)\Big\}\le 0,$$

Again, we will distinguish, for a fixed drift $\mu<0$ and a fixed upper stopping threshold $b>x$, different cases depending on the lower stopping threshold $a<x$ that maximizes the upper bound on the difference in salience-weighted utility, which is now given by

$$\begin{aligned}\tilde{\Delta}(a,b,\mu):=\mathbb{P}_0[\tau_{a,b}<T]\cdot&\big[p(a-x)\underline{\sigma}+(1-p)(b-x)\overline{\sigma}\big]\\ &+\mathbb{P}_0[\tau_{a,b}\ge T]\cdot\left[\underline{\sigma}\int_{(a^*,x)}(z-x)\,d\Phi_\mu(z)+\overline{\sigma}\int_{(x,b)}(z-x)\,d\Phi_\mu(z)\right]\end{aligned}$$

1. Case: Suppose $\tilde{\Delta}(a,b,\mu)$ is maxmized at $a^*=a^*(b,\mu)$ with $\lim_{\mu\to-\infty}a^*(b,\mu)<x-\epsilon$ for some $\epsilon>0$. By the first step (i.e., $T=\infty$), there exists some $\tilde{\mu}\in\mathbb{R}$ such that for any $\mu<\tilde{\mu}$,

$$\tilde{\Delta}(a^*,b,\mu)\le\mathbb{P}_0[\tau_{a^*,b}\ge T]\cdot\left[\underline{\sigma}\int_{(a^*,x)}(z-x)\,d\Phi_\mu(z)+\overline{\sigma}\int_{(x,b)}(z-x)\,d\Phi_\mu(z)\right].$$

Now it is sufficient to show that

$$\begin{aligned}\lim_{\mu\to-\infty}\underline{\sigma}\int_{(a^*,x)}&(z-x)\,d\Phi_\mu(z)+\overline{\sigma}\int_{(x,b)}(z-x)\,d\Phi_\mu(z)\\ &<\lim_{\mu\to-\infty}\underline{\sigma}\int_{(x-\epsilon,x)}(z-x)\,d\Phi_\mu(z)+\overline{\sigma}\int_{(x,b)}(z-x)\,d\Phi_\mu(z)<0,\end{aligned}$$

which follows immediately from the fact that, by Lemma 2 (g), we have $\lim_{\mu\to-\infty}\Phi_\mu(x)=1$.

2. Case: Suppose that $\tilde{\Delta}(a,b,\mu)$ is maximized in the limit of a approaching x. Then,

$$\sup_{(a,b)\in[0,x)\times(x,\infty)} \tilde{\Delta}(a,b,\mu) = \sup_{b\in(x,\infty)} \lim_{a\to x} \tilde{\Delta}(a,b,\mu) = 0.$$

since, for any fixed b and any fixed $T>0$, by Lemma 2 (b) it follows that $\lim_{a\to x} \mathbb{P}_0[\tau_{a,b} < T] = 1$.

3. Case: Suppose, for the sake of a contradiction, that a maximizer $a^* = a^*(b,\mu)$ of $\tilde{\Delta}(a,b,\mu)$ exists and that any maximizer satisfies $\lim_{\mu\to-\infty} a^*(b,\mu) = x$. Notice that, for any fixed $b > x$, the mapping $\tilde{\Delta}(a,b,\mu)$ is continuous in both a and μ. Next, notice that the set of feasible a is not constrained by the parameter μ, that is, any stopping threshold $a \in [0,x)$ is feasible, irrespective of the drift μ. Hence, the set of feasible thresholds a is trivially continuous, because constant, in μ. Obviously, $[0,x)$ is not compact. But, as we assume that the set of maximizers is non-empty, Berge's Maximum Theorem still implies that this set has to be upper hemicontinuous.

Now, for any fixed $a < x < b$, we have $\lim_{\mu\to-\infty} \mathbb{P}_0[\tau_{a,b} < T] = 1$, and thus it follows that

$$\lim_{\mu\to-\infty} \tilde{\Delta}(a,b,\mu) = \lim_{\mu\to-\infty} \big[p(a-x)\underline{\sigma} + (1-p)(b-x)\overline{\sigma}\big].$$

We already know that, if $p(a-x)\underline{\sigma} + (1-p)(b-x)\overline{\sigma}$ has a maximizer, then it is unique and given by the expression in Eq. (3). Since the set of maximizers is non-empty by assumption and, in particular, upper hemicontinuous, we know that, in the limit of μ approaching negative inifinity, it has to include the limit of Eq. (3), which is given by $x + \frac{\overline{\sigma}}{\underline{\sigma}}(b-x) > x$; a contradiction. $\square$

Proof of Proposition 4. Fix an initial wealth level $x \in \mathbb{R}_{>0}$ and a volatility ν. Consider the two processes with drift parameters $0 \geq \mu' > \mu$. We have to show that *if a naive salient thinker does not stop the process with drift μ immediately, then he does not stop the process with drift μ' immediately.*

Without loss of generality, we can assume that $\mu' < 0$, as we have already seen that for $\mu' = 0$ there is a stopping strategy that is more attractive than stopping immediately (Proposition 1). As before, let $\Phi_\mu(z) := \mathbb{P}_0[X_T \leq z | X_0 = x, \tau_{a,b} \geq T]$. If a naive salient thinker does not stop the process with drift μ immediately, there exists a stopping time $\tau_{a,b}$ such that

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

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

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

2. Case: Suppose that, for the process with drift μ, any attractive stopping time $\tau_{a,b}$ implies that $(\star)$ is negative, which in turn implies that $(\star\star)$ is positive. Take an attractive stopping time $\tau_{a,b}$ and let $T \geq -\frac{(x-a)}{\mu'} =: \hat{T}$. Then, by Lemma 2 (h), we have $\frac{\partial}{\partial\mu}\mathbb{P}_0[\tau_{a,b} < T] < 0$ for any $\mu < \mu'$ and therefore $\big[(\star) - (\star\star)\big] \cdot \frac{\partial}{\partial\mu}\mathbb{P}_0[\tau_{a,b} < T] > 0$ for any $\mu < \mu'$. But this, together with the fact that $(\star)$ and $(\star\star)$ monotonically increase in μ, implies that, for any $T \geq \hat{T}$,

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

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

3. Case: Suppose that, for the process with drift μ, any attractive stopping time $\tau_{a,b}$ implies that $(\star)$ is positive, but $(\star\star)$ is negative. Then, also for the process with a drift μ', $(\star)$ is positive. Moreover, for a fixed stopping time $\tau_{a,b}$, $(\star)$ is independent of T. By Lemma 2 (c), it holds that $\lim_{T\to\infty}\mathbb{P}_0[\tau_{a,b} < T] = 1$, which then implies that, in the limit of T approaching infinity,

$$\mathbb{P}_0[\tau_{a,b} < T]\big[p(a-x)\sigma(a,x) + (1-p)(b-x)\sigma(b,x)\big] \\ + \mathbb{P}_0[\tau_{a,b} \geq T]\int_{(a,b)}(z-x)\sigma(z,x)\, d\Phi_{\mu'}(z) > 0.$$

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

Appendix B: Sophisticated Stopping Behavior Without Commitment

Solution concept. A sophisticated salient thinker differs from her naive counterpart in that she anticipates her future selves to act differently which she takes into account when making her stopping decision. A sophisticated salient thinker who lacks commitment then behaves as if she is playing a game with her future selves (see, e.g., Karni and Safra, 1990). In order to solve this game, we adopt the following equilibrium concept of Ebert and Strack (2018):

Definition 5 (Equilibrium)**.** *A stopping strategy s constitutes an* equilibrium *if and only if at every point in time t it is optimal to take the decision $s(X_t, t)$, given that all future selves follow this strategy.*

Avoiding unfair gambles. Assuming $\mu \leq 0$, we solve for the set of equilibria defined above. The next results shows that, in the context of our experiment, naivete is a necessary assumption to explain (unfair) casino gambling within the salience framework.

Proposition 5. *Suppose that the agent can only choose time-invariant strategies. Fix an initial wealth level $x \in \mathbb{R}_{>0}$, and suppose that the process has a non-positive drift $\mu \leq 0$. Then, in any equilibrium, the sophisticated salient thinker stops immediately.*

As an illustration, let us assume that $T = \infty$. In this case any Markov strategy is time-invariant and therefore can be represented by a stopping time $\tau_{a,b}$. For any such stopping time, there exists some wealth level $y' \in (a, b)$ such that the downside of the binary lottery $X_{\tau_{a,b}}$ is salient when evaluated in the choice set $\mathcal{C} = \{X_{\tau_{a,b}}, y'\}$. Moreover, if the process has a non-positive drift, then, at any wealth level y, we have $\mathbb{E}[X_{\tau_{a,b}}] \leq y$. Since a salient thinker values a binary lottery with a salient downside below its expected value, the sophisticated agent anticipates to stop no later than at wealth level y'. Thus, by Definition 5, a strategy s that satisfies $s(X_t) = continue$ for any $X_t \in (a, b)$ cannot constitute an equilibrium. In contrast, at any initial wealth level $x \in \mathbb{R}_{>0}$, stopping immediately can be supported as an equilibrium outcome: given that all future selves will stop immediately, the current self is indifferent between stopping immediately and continuing the process, so that it is indeed optimal to stop at every point in time. The following proof extends this line of reasoning to finite expiration dates.

Proof of Proposition 5. Fix an initial wealth level x and a drift $\mu \leq 0$. The proof proceeds in two parts: in Part (a) we assume that $T = \infty$, while in Part (b) we consider the case with $T < \infty$.

PART (a). Let $T = \infty$. Throughout the proof we will compare a given stopping strategy (i.e., a potentially risky option) to the alternative of stopping immediately (i.e., a safe option). Notice that, for a given stopping time $\tau_{a,b}$, a change in the drift μ induces a first-order stochastic dominance shift in $X_{\tau_{a,b}}$. Hence, by Proposition 1 in Dertwinkel-Kalt and Köster (forthcoming), it is sufficient to consider the case of a zero drift. We first prove the following lemma.

Lemma 3. *Let $\mathcal{C}_t = \{L_{s,t}, x_t\}$, where s is a stopping strategy and x_t denotes the safe option that pays the current wealth level. The stopping strategy s constitutes an equilibrium in the sense of Definition 5 if and only if gambling behavior is described by the first leaving time $\tau_{a,b}$ of an interval (a, b) such that*

$$U^s(X_{\tau_{a,b}}|\mathcal{C}_t) \geq U^s(x_t|\mathcal{C}_t) \quad \textit{for all } t < \tau_{a,b}, \tag{4}$$

where $X_{\tau_{a,b}}$ gives the lottery induced by stopping strategy s with leaving time $\tau_{a,b}$.

Proof. Since, for $T = \infty$, any stopping strategy s implements a binary lottery $X_{\tau_{a,b}}$ over future wealth and since a sophisticated salient thinker is aware of her time-inconsistency, she anticipates that her future selves will follow through with a stopping strategy s only if Eq. (4) holds. In order to see that Eq. (4) is also sufficient for describing equilibrium behavior, consider the stopping strategy s that satisfies $s(y) = continue$ if and only if $y \in (a, b)$. By construction, it is optimal to continue at any wealth level $y \in (a, b)$. Since the process has continuous paths, it is also optimal to stop at any $y' \notin (a, b)$, given that all future selves follow through with the stopping strategy s: deviating to $s(y') = continue$ would not change the outcome, as the process is stopped with probability one in any arbitrarily short period of time anyhow (Ebert and Strack, 2018, p. 14). As a consequence, the stopping strategy s constitutes an equilibrium and Eq. (4) therefore fully describes equilibrium behavior. $\square$

For any stopping time $\tau_{a,b}$, with $0 \leq a < x < b$, there exists some wealth level $y' \in (a, b)$ such that $\sigma(a, y') > \sigma(b, y')$. In addition, since the process has zero drift and the censoring in zero does not play a role due to $a \geq 0$, at any wealth level y, we have $\mathbb{E}[X_{\tau_{a,b}}] = y$. Since a salient thinker values a binary lottery with a salient downside below its expected value, the sophisticated salient thinker anticipates stopping at wealth level y'. Thus, by Definition 5, a strategy s that satisfies $s(X_t) = \textit{continue}$ for any $X_t \in (a,b)$ cannot constitute an equilibrium. In contrast, at any initial wealth level $x \in \mathbb{R}_{\geq 0}$, stopping immediately can be supported as an equilibrium outcome: given that all future selves will stop immediately, the current self is indifferent between stopping immediately and continuing the process, so that it is indeed optimal to stop at every point in time. This completes the proof of this part.

PART (b). Let $T < \infty$. We have to show that *it cannot be an equilibrium to play according to a strategy that implements the stopping time* $\tau_{a,b}$ (*conditional on stopping before the expiration date*). At any point in time t with a wealth level $X_t = y \in (a,b)$, a salient thinker follows the strategy that implements $\tau_{a,b}$ (conditional on stopping before the expiration date) if and only if

$$\mathbb{P}_t[\tau_{a,b} < T] \cdot \big[p(a-y)\sigma(a,y) + (1-p)(b-y)\sigma(b,y)\big] \\ + \mathbb{P}_t[\tau_{a,b} \geq T] \cdot \int_{(a,b)} (z-y)\sigma(z,y)\, d\Phi_\mu(z) \geq 0,$$

where the probability $p = (a,b,\mu)$ is defined as in Eq. (1) and where Φ_μ denotes the CDF of X_T conditional on reaching the expiration date with this strategy. Notice that $\sigma(a,y) > \sigma(b,y)$ holds for any wealth level y sufficiently close to b. Also, $\mathbb{E}_t[X_{\tau_{a,b}} | X_t = y] \leq y$ due to the non-positive drift of the process. This implies that $p(a-y)\sigma(a,y) + (1-p)(b-y)\sigma(b,y) < 0$ for any wealth level y sufficiently close to b. Since, for any fixed time t, we have $\lim_{y \to b} \mathbb{P}_t[\tau_{a,b} < T] = 1$, we thus conclude that for any $\tau_{a,b}$ there is some $y' \in (a,b)$ such that

$$\mathbb{P}_t[\tau_{a,b} < T] \cdot \big[p(a-y)\sigma(a,y) + (1-p)(b-y)\sigma(b,y)\big] \\ + \mathbb{P}_t[\tau_{a,b} \geq T] \cdot \int_{(a,b)} (z-y)\sigma(z,y)\, d\Phi_\mu(z) < 0,$$

which was to be proven. Finally, by the same arguments as in Part (a), *stopping immediately can be supported as an equilibrium outcome*. This completes the proof. $\square$