Stubborn by Design: Neurobiological Foundation of Maladaptive Risk-Taking at Downturn Onsets*

Elise Payzan-LeNestour[†1], Yunshen Yang[1], Samuel Thelaus[2], and Bernard Balleine[1]

[1]University of New South Wales
[2]London School of Economics

February 9, 2025

Abstract

We document a spillover of past rewards into subsequent gambling decisions in a controlled laboratory experiment: after receiving rewards from a profitable gambling strategy, participants persist with it despite being explicitly informed it has become unprofitable. Through experimental manipulations and computational modelling, we isolate a Pavlovian mechanism driving this maladaptive behavior: past rewards create an automatic attraction to gambling that can derail rational choice. This Pavlovian influence intensifies under uncertainty, suggesting it extends beyond gambling for monetary rewards to include "sign-tracking"—an attraction to gambling cues independent of monetary outcomes. These findings may help explain behavioral persistence in financial markets and suggest dual-focused interventions for problem gambling: both recalibrating outcome expectations and redirecting attention away from gambling cues.

Keywords: Risk-taking; Pavlovian influences; Investor inertia; Problem gambling

1 Introduction

Recent research highlights a disconnect between investors' beliefs and their actions, with changes in beliefs often failing to translate into corresponding adjustments in investment behavior (Giglio et al., 2021). This apparent inertia has prompted diverse explanations, including behavioral factors

*We thank the many colleagues who provided insightful feedback on the study, including Peter Bossaerts, Giorgio Coricelli, Brice Corgnet, Mateus Joffily, Etienne Koechlin, Marie-Claire Villeval, and seminar participants at The 2022 Asian Meetings of the Econometric Society, The 2023 Economic Science Association World Meeting, UNSW School of Psychology, Ecole Normale Supérieure de Paris, GATE-Lab at the University of Lyon, and the University of Chicago. The analyses and conclusions set forth are those of the authors. Funding from UNSW Business School is gratefully acknowledged.

[†]Corresponding author: Email: elise@elisepayzan.com

like investor risk aversion, underconfidence, and inattention (Giglio et al., 2021), as well as imprecise beliefs (Charles et al., 2024). The well-documented "status quo" bias (Samuelson and Zeckhauser, 1988; Dean et al., 2017) may also be relevant, alongside institutional factors such as agency frictions and passive fund strategies (Gabaix and Koijen, 2021).

This paper presents a complementary perspective grounded in human neurobiology, examining how a key aspect of brain function—the so-called "*Pavlovian system*"—may explain an apparent paradox in investor behavior: clinging to formerly successful strategies despite being aware of their current harm. Neuroscientists call this "*Pavlovian misbehavior*," where conditioned responses derail rational decision-making.[1]

To illustrate, consider the human brain as a car with two drivers: a rational one who plans the optimal route (referred to by neuroscientists as the "deliberative" or "instrumental" system), and an impulsive backseat driver who reacts reflexively to environmental cues, disregarding future implications (the Pavlovian system). The backseat driver frequently derails planned routes—for example, insisting on staying on the current road to pass by a familiar restaurant associated with past enjoyment, despite a faster route being available. This illustrates how the two systems fundamentally differ in action initiation: the instrumental system takes actions *for* their anticipated results, whereas the Pavlovian system reacts to the situation *despite* potential consequences (Dayan et al., 2006).

In the investor's brain, the two systems clash when market conditions shift: Even when an investor rationally recognizes that a strategy has become unfavorable, their Pavlovian system may still crave it, driven by recent positive outcomes. When this reflexive urge overrides rational analysis, it produces our hypothesized anomaly: clinging to a previously rewarding but now-suboptimal strategy. Such apparent 'stubbornness' mirrors our driving metaphor—the impulsive backseat driver compels persistence on the familiar route, despite the rational driver's awareness of a better alternative.

We find evidence for this form of Pavlovian-driven inertia in a controlled laboratory experiment. Participants engaged in a gambling task consisting of 120 sessions, each comprising five trials. In each trial, they decided whether to bet on a slot machine for a 70-cent fee or pass. The session background color, either blue or yellow, indicated positive or negative expected values for spinning the slot machine, respectively. To ensure complete transparency, the task instructions explicitly detailed the underlying statistics and every aspect of the experimental design, preventing participants from forming arbitrary assumptions. The instructions also emphasized performance-based rewards (up to $110) and potential losses (up to $50), providing strong incentives for optimal behavior.

The main "Test" Treatment exposed participants to a series of blue background sessions, known to have a positive expected value, followed by a yellow background session, known to have a negative expected value. As anticipated, we observed near-universal betting in blue sessions, consistent with rational decision-making under perfect information and strong incentives for optimal behavior.

[1]For example, Breland and Breland (1961), Dayan and Balleine (2002), and Dayan et al. (2006).

Crucially, we also observed a negative spillover effect: participants' betting behavior extended into the yellow sessions, consistent with the core paradox proposed in this study.

Demonstrating the paradox requires strict control over participant beliefs, a challenge inherently difficult in real-world scenarios. Our experimental task, however, achieved this through full disclosure of all relevant information to participants, along with other design features detailed in Section 2.1.3. This approach, advocated by Shmaya and Yariv (2016), maintained rigorous control over participant beliefs throughout the experiment and minimized the risk that the observed spillover stems from incomplete information or misunderstanding of the task structure.

Besides allowing us to control for participant beliefs, the laboratory experimental approach also afforded us a high degree of control to isolate the proposed Pavlovian factor from other potential spillover sources. Our strategy incorporated four key elements:

1. Incentivizing participants through a "double-edged" incentive structure. Our experiment featured a distinctive approach to participant motivation: a "double-edged" incentive structure where participants faced both genuine financial losses during the experimental task (up to $50) and significant reward potential (up to $110). This design choice, combining both downside risk and upside potential at meaningful scales, aimed to replicate real-world incentives in the lab to the fullest extent possible. This was crucial as a lack of consequential incentives would be an obvious confound for the expression of the hypothesized anomaly in participant behavior.

2. Model comparison analyses. Our design enabled a reliable test of the Pavlovian hypothesis against a prominent alternative explanation for the observed spillover: the "*Law of Effect*" or "*model-free*" behavior.[2] This theory posits that action value is directly tied to recent outcomes, reinforcing successful actions and not unsuccessful ones, in contrast to "*model-based*" behavior where actions are chosen based on expected values. To distinguish between model-free and model-based explanations of participant behavior in our task, we incorporated two critical design features: no feedback on outcomes of foregone bets when participants choose to pass, preventing value updating for the unchosen betting option, together with a relatively high loss probability (0.2) in positive-expected-value sessions, leading model-free agents to frequently avoid betting despite the positive expected value. In contrast, model-based agents seldom pass in these sessions.

Participants' behavior unambiguously aligned with the model-based prediction. This alignment, along with other behavioral aspects (detailed in Section 3.4), invalidates the Law of Effect as an explanation for the observed spillover anomaly. Instead, incorporating the hypothesized Pavlovian bias into our model-based framework significantly improved fit to participant behavior. Our Pavlovian model specifically captures the observed tendency to persist in selecting a previously rewarding betting option even after it becomes disadvantageous, a behavioral pattern that neither pure model-based nor model-free accounts can explain.

3. Control Treatment. To isolate the hypothesized Pavlovian influence from other potential spillover sources, we introduced a Control Treatment that accounted for factors such as "*loss tolerance*" (Brown et al., 2024; Chapman, 2024), faulty cognition (Benjamin, 2019), behavioral inat-

[2]For example, Daw (2014), Payzan-LeNestour and Bossaerts (2015), and Barberis and Jin (2023).

tention (Gabaix, 2019), "*probability matching*" (where participants match their betting frequency to the probability of success; e.g., Lo (2017) and Martínez-Marquina et al. (2019)), and psychological factors like thrill-seeking (Conlisk, 1993) or boredom avoidance (Schüll, 2012). This Control Treatment replicated the Test Treatment conditions, with one crucial alteration: yellow sessions preceded blue ones. This sequence reversal ensured that at the group level, yellow session betting in the Control Treatment could not stem from Pavlovian influences, thus establishing a Pavlovian-free baseline for comparison. If the above non-Pavlovian factors were driving yellow session betting, we would expect similar betting rates across treatments. Instead, we found "excess" betting in the Test Treatment's yellow sessions compared to the Control baseline, suggesting Pavlovian influences on participant decision-making.

4. Exploiting a Pavlovian signature. To further substantiate Pavlovian effects and probe the nature of the underlying mechanisms, we tested for intensification of the observed spillover anomaly with uncertainty. This pattern, not predicted by non-Pavlovian explanations, uniquely characterizes "*sign-tracking,*" a specific form of Pavlovian influence where conditioned agents treat reward cues as inherently rewarding, transferring their response from the unconditioned stimulus (monetary reward) to the cue itself (slot machine).[3] In our experimental task, participants exhibiting sign-tracking tendencies ("sign trackers") would develop an attraction to the slot machine, treating the game as an end in itself, rather than a means to monetary outcomes. Unlike primary monetary addiction (referred to as "*Pavlovian goal-tracking*" by neuroscientists; e.g., Russell and Robinson (2019)), sign-tracking intensifies when the relationship between the Pavlovian cue and the reward becomes more unpredictable. We tested for this signature by manipulating uncertainty in the slot machine-reward association while preserving rational behavior. In half the sessions, chosen slot machine spins were guaranteed, while in the other half, there was a 50% probability of spin denial. We found that uncertainty positively impacted participant betting rates, providing additional evidence for Pavlovian influences on decision-making.

Relation to literature. This paper synthesizes insights from anthropology, experimental economics, and psychology to expand our understanding of investor inertia. Building on Giglio et al. (2021)'s seminal study, recent research points to multiple causes, including institutional factors (Gabaix and Koijen, 2021), cognitive imprecision (Charles et al., 2024), and boundedly rational ("model-free") investor learning (Barberis and Jin, 2023). Our study controls for these and proposes a complementary explanation rooted in human neurobiology. We argue that even ideal decision-makers—fully rational and motivated "*active thinkers*"[4]—can succumb to inertia due to Pavlovian disturbances, much like the rational driver disturbed by the impulsive backseat driver in our earlier metaphor. Crucially, these disturbances affect everyone, regardless of knowledge or expertise. As such, they may have important implications for asset pricing. Future research could translate our experimental findings into testable market predictions, potentially revealing new patterns in asset pricing.

[3]For example, Brown and Jenkins (1968), Hearst and Jenkins (1974), and Anselme and Güntürkün (2019).
[4]For example, Szollosi and Newell (2020) and Szollosi et al. (2023).

While Pavlovian influences on human behavior have been extensively explored in psychology and neuroscience,[5] their examination in economic settings remains scarce. Two exceptions are Bushong et al. (2010), who provide evidence for Pavlovian influences on consumer choice, and Payzan-LeNestour and Doran (2024), who investigate analogous Pavlovian disturbances on investor behavior. Bushong et al. (2010) reveal that physically presenting an appealing food item can elicit a Pavlovian response that is substantially diminished or even absent when that item is merely shown symbolically (as an image) or abstractly (as text). Payzan-LeNestour and Doran (2024) highlight the "pitfall of negative skew:" investors' tendency to favor investments with frequent small gains but occasional large losses—resulting in high Pavlovian value despite negative expected value. There, the rational and Pavlovian systems conflict. Our study reveals another such conflict where Pavlovian responses can derail investor decision-making: behavior during market downturn onsets. Here, previously profitable strategies now carry negative expected value yet retain strong Pavlovian appeal. When the Pavlovian response dominates, it leads to the paradoxical behavior captured by our experimental design: persistence in a strategy despite recognizing its shift to negative expected value.

Our study's implications extend beyond finance, providing insight into why problem gamblers persist despite mounting losses. While harmful gambling is often attributed to faulty cognition (gamblers misjudging unfavorable odds), our experiment, by controlling for this factor, suggests that the phenomenon cannot be explained by cognitive errors alone. This view, already articulated by Redish (2013), aligns with recent anthropological observations of casino gamblers entering "the machine zone" (Schüll, 2012), a state of absorbed focus on gaming machines with apparent disregard for monetary outcomes. As discussed above, this behavior closely parallels Pavlovian sign-tracking, a phenomenon extensively documented by neuroscientists in animal studies. This synthesis of anthropological and neuroscientific evidence suggests that purely motivational factors, referred to by neuroscientists as "*irrational wanting*,"[6] work alongside faulty cognition ("*miswanting*") in driving problem gambling.

A distinctive feature of this study is its attempt to examine the specific form that Pavlovian influences take in gambling decisions. Following recent studies that identified Pavlovian influences on human behavior, and responding to their explicit call for investigating underlying mechanisms,[7] we address a key unresolved question: Is it simply a strong pull toward the unconditioned stimulus—akin to money addiction—or does it also involve an attraction to the conditioned stimulus itself—sign-tracking?

This question is crucial for protecting both recreational gamblers and vulnerable investors. Indeed, the choice of appropriate protective measures depends critically on the underlying mechanism. For those gambling primarily for money, interventions should prioritize educational approaches to recalibrate their outcomes expectations, such as simulations of probable long-term outcomes. Sign-trackers, on the other hand, show heightened sensitivity to sensory stimuli in gambling venues and

[5] See Payzan-LeNestour and Doran (2024) for references and discussion.
[6] For example, Winkielman and Berridge (2003) and Berridge and O'Doherty (2014).
[7] For example, Dorfman and Gershman (2019) and Payzan-LeNestour and Doran (2024).

to gamified elements in trading platforms. As such, they will benefit more from interventions targeting cue reduction—minimizing environmental triggers and gamification elements—and attention training to redirect focus away from gambling-associated cues. By suggesting that Pavlovian influences on gambling decisions involve sign-tracking beyond mere attraction to money, our findings underscore the importance of environmental design.

This emphasis on environmental cues aligns with Bushong et al. (2010)'s aforementioned insights [that not all associative cues operate as conditioned cues capable of triggering a Pavlovian response; 'real' stimuli appear particularly powerful]. This suggests that the Pavlovian bias captured in our computerized experiment may be amplified in traditional casinos where the Pavlovian cue and its associated reward are fully 'real' (e.g., physically receiving coins from an actual slot machine), rather than represented symbolically onscreen as in online gambling platforms and the current task. Bushong et al. (2010)'s findings further suggest that Pavlovian cues are "fragile" in that their association with the unconditioned stimulus depends heavily on contextual factors.[8] Combined with our findings, this insight may help motivate the design of environments to mitigate Pavlovian disturbances on gambling decisions.

2 Generating the paradox of informed inertia in the laboratory

2.1 Experimental design

The study's hypotheses, methods, and analysis plan were preregistered on the Open Science Framework (OSF) prior to data collection. All study materials—preregistration document, task instructions script and demo of their online delivery, task demo, and post-task debriefing—are available here.

2.1.1 Task interface.

We designed a web-based video game for our experiment. The game consists of 630 trials and takes approximately 30 minutes to complete. On each trial, participants must choose between spinning a slot machine for a 70-cent fee or passing. The slot machine cue remains visible on screen throughout the decision phase. Opting to pass advances the game to the next trial without an outcome. Choosing to spin results in either a 1-second animated spin followed by a 1-second bimodal outcome display (visual outcome paired with outcome-specific sound), or a spin denial screen (Fig. 1A). In the event of a denial, the fee is refunded.

The game is structured into sequences of 5-trial "sessions", which vary in winning amount ($1 or $2), gambling expected value (positive or negative), and play availability (always guaranteed or uncertain). Before each session, a 3-second pre-session screen displays visual cues revealing the context to the participant (Fig. 1A). Coin stimuli indicate the reward amount in the session (1 coin: $1; 2 coins: $2). Background color signifies gambling expected value (blue: positive, 0.8 winning

[8] For instance, introducing a physical barrier between the subject and the physical food item dramatically reduced the Pavlovian response in their experiment.

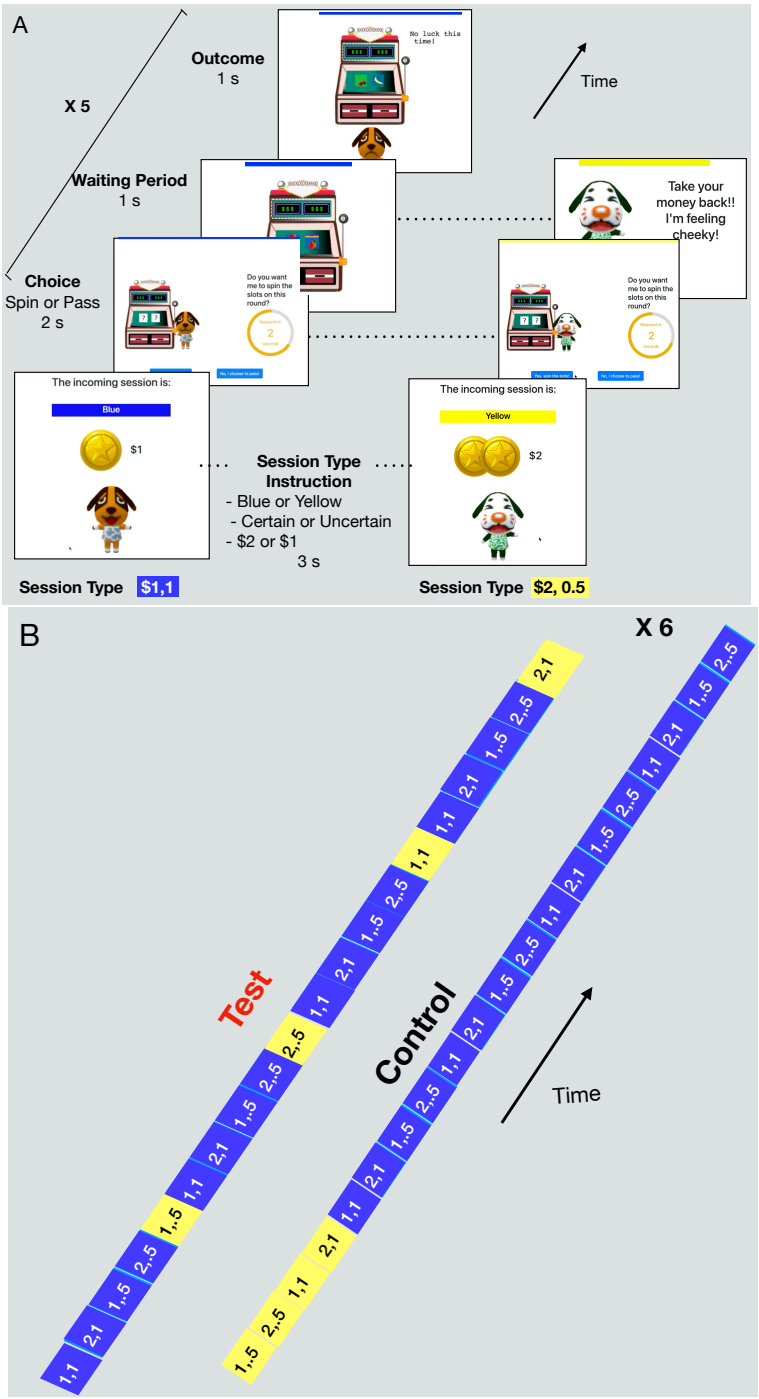


Figure 1. Experimental design. (A) Task interface. The pre-session screen displays the context for the upcoming 5-trial block ("session"). Background color indicates the expected value of gambling in the session (blue: positive; yellow: negative). A dog's face signals play availability (happy brown: guaranteed; cheeky white: uncertain). Coin stimuli show potential win amounts (1 coin: \$1; 2 coins: \$2). On each trial, choosing to spin leads to either a 1-second slot animation followed by a 1-second outcome display (left sequence), or a play denial screen (right sequence). **(B) Task structure.** Each square represents one session. Blue: positive expected value for gambling. Yellow: negative expected value for gambling. In the Test Treatment, yellow sessions always follow a series of blue sessions. In the Control Treatment, yellow sessions come first. Session types: $(2, 0.5)$: \$2 for winning spin, 50% chance of play. $(1, 0.5)$: \$1 for winning spin, 50% chance of play. $(2, 1)$: \$2 for winning spin, 100% chance of play. $(1, 1)$: \$1 for winning spin, 100% chance of play. Session type is counterbalanced across colors and session order is randomized within each color as shown in Appendix Appendix B.

probability; yellow: negative, 0.2 winning probability). A dog's face represents play availability (happy brown face: always guaranteed; "cheeky" white face: 50% chance of spin denial on each trial). Participants undergo a thorough 25-minute instruction phase, ensuring a full grasp of all cues and game mechanics. By eliminating any ambiguity, this comprehensive pre-task briefing (detailed in Section 3.1.1) ensures participants make fully informed decisions on each trial.

Eliciting participant beliefs. A pre-game quiz assesses participants' initial strategies and perceived winning probabilities for yellow and blue background sessions. During the game, after betting during a yellow session, participants engage in a 1-minute real-money lottery aimed to elicit their perceived winning probability in yellow sessions. This lottery, presented up to three times per participant to maintain game flow, uses a numerical probability format initially, then a pie chart format if presented again,[9] to accommodate different comprehension styles (e.g., Hoffrage et al. (2000)). Post-game, participants who placed at least one bet during yellow sessions are asked again to report the winning probabilities for both session types, mirroring the pre-game question.

2.1.2 Payoff rule.

Task earnings are calculated as 24 times the difference between a participant's final accumulated outcomes and a preset threshold, capped at \$110. For instance, with a \$100 threshold and \$102 in accumulated outcomes, a participant earns $(\$102 - \$100) \times 24 = \$48$. Negative differences result in losses, capped at \$50. The threshold is determined before the task begins but revealed only upon completion. Calibrated using conservative forecasts of participants' accumulated outcomes, it ensures average earnings of at least \$40—honoring commitments made in the task instructions (detailed in Section 3.1.1).

2.1.3 Task structure.

The task's basic structure is shown in Fig.1B. The Test Treatment consists of 6 sequences, each comprising 4 repetitions of a series of positive expected value (blue) sessions followed by a negative (yellow) session. A 1-minute break follows each sequence. The Control Treatment is identical but reverses session order (yellow sessions precede blue). In both treatments, we counterbalanced play availability and reward size factors across colors, ensuring each session type appeared equally often in blue and yellow sessions. We randomized session order within each color while maintaining the balanced representation of session types (see Appendix Appendix B for a visual representation of this procedure). We enhanced this structure by adding a blue session at the start of each sequence and randomly inserting another blue session within the Control Treatment's yellow block, a crucial addition to boost statistical power, see Section 2.2.[10]

[9] Screenshots are provided here.

[10] We added the same blue sessions to the Test Treatment's sequences to maintain structural equivalence and identical trial counts across treatments.

2.2 Rationale behind this design

We aimed to design a controlled and highly incentivized lab environment for our tests.

2.2.1 Task interface.

The task is "neutral" in the sense of Alekseev et al. (2017), avoiding evocative elements, especially financial undertones. This approach ensures accessibility for all participants, regardless of financial literacy (Lusardi and Mitchell, 2007). It also prevents unintended associations with familiar investing contexts or financial stereotypes, crucial for controlling participant beliefs throughout the experiment as advocated by Shmaya and Yariv (2016).

Eliciting participants' beliefs about winning probabilities and their initial game strategies helps control for the faulty cognition factor: betting in yellow sessions due to misperceived odds. We applied a pre-registered filter based on these elicited beliefs, identifying the 'cognizant' group: participants with accurate beliefs and a clear pre-game strategy (more in Section 3.1.3). By focusing our analyses exclusively on this group, we increase the likelihood that any observed harmful betting stems from behavioral, not cognitive, causes—ruling out misunderstandings of game mechanics.

2.2.2 Task structure.

Main treatment effect. A positive mean betting rate in the Test Treatment's yellow sessions, averaged across participants, would demonstrate our proposed paradox: gambling persists despite awareness of its newly negative expected value. This spillover could stem from various behavioral factors (Fig.2) or unobservable experimental 'quirks' (e.g., color-induced betting). The Control Treatment isolates the Pavlovian factor from these other potential causes: The Pavlovian factor is engineered into the Test Treatment only—our design ensures its absence in the Control Treatment, while these other factors exist in both treatments alike. Any excess harmful betting in Test compared to Control would thus indicate the Pavlovian factor's contribution to the spillover effect.

Our treatment effect is precisely this: the difference in average betting rates during yellow sessions between Test and Control Treatments.

Fully counterbalancing play availability and reward size factors across colors while randomizing session type order within each color aims to prevent bias in measuring the treatment effect and eliminate session order artifacts, enhancing result generalizability. The addition of a couple of blue sessions to this basic design serves to boost our power in detecting the treatment effect by addressing a potential experimental artifact in the Control Treatment: participants might bet during the block of yellow sessions to alleviate boredom from consecutive skipping (20 trials in a row). Interspersing a couple of blue sessions within the yellow block injects variability and breaks the monotony of the task. This strategy mitigates the confounding factor of boredom-induced

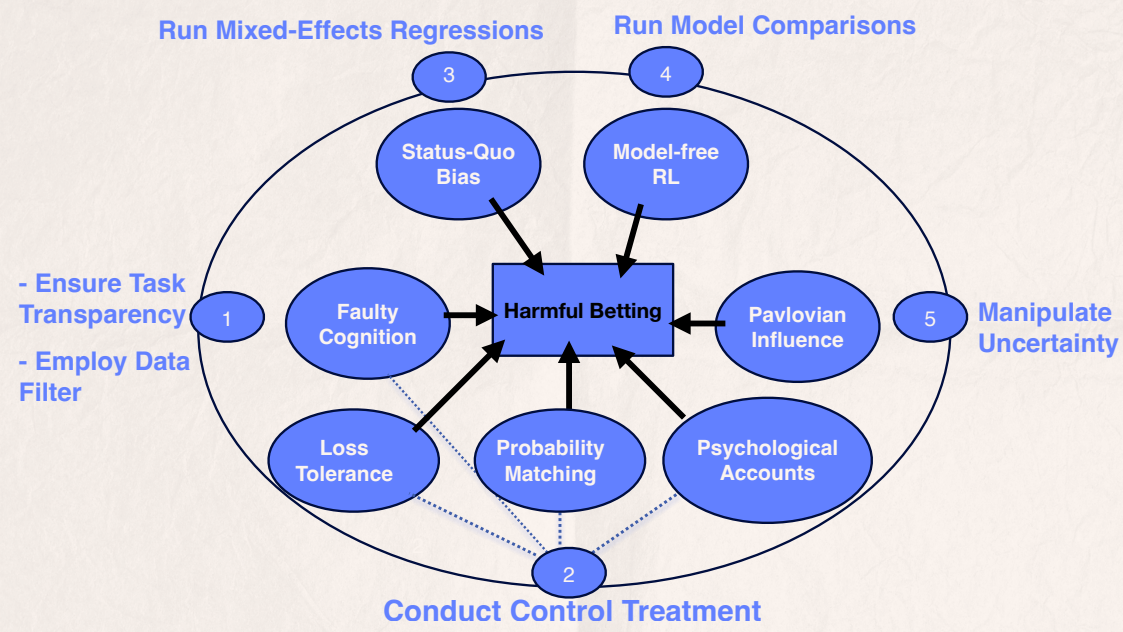


Figure 2. Controlling for alternative explanations of maladaptive betting. Participants might bet maladaptively for several reasons beyond the purported Pavlovian factor: misunderstanding the negative expected value due to unclear instructions or limited attention/memory ("faulty cognition"); seeking gambling thrills or avoiding boredom ("psychological accounts"); following model-free Reinforcement Learning (Barberis and Jin, 2023) or probability matching (Lo, 2017; Martínez-Marquina et al., 2019); being loss tolerant (Brown et al., 2024; Chapman, 2024); or exhibiting the status quo bias (Samuelson and Zeckhauser, 1988; Dean et al., 2017). Our experimental strategy addressed each possibility through: (1) transparent pre-task instructions and strict data filtering; (2) Treatment comparisons to control for all factors except model-free RL and status quo bias; (3) mixed-effects regressions to control for status quo bias; (4) model comparisons to control for model-free RL; and (5) Uncertainty manipulation designed to test for a Pavlovian signature (amplification of betting with uncertainty).

betting, thereby enhancing the validity of our Control condition.[11]

Uncertainty effect. Introducing the play availability factor aims to manipulate uncertainty in the slot machine-reward association while, crucially, not affecting rational behavior (as shown in Table 4, "Base"). This manipulation tests whether the Pavlovian influence extends beyond monetary attraction ("goal-tracking") to include an attraction to the gambling cue itself ("sign-tracking"). Unlike goal-tracking, whose strength shows no established relationship with uncertainty, sign-tracking has been shown to intensify with uncertainty.[12] Therefore, finding a stronger anomaly in high-uncertainty sessions would suggest the presence of sign-tracking.

[11] This enhancement of validity holds only if this strategy does not inadvertently induce significant Pavlovian effects at the group level. We verified this is not the case by simulating our Pavlovian model (described p.16) with parameters within empirically supported ranges. We found negligible betting rates within these yellow blocks, confirming the validity of our strategy.

[12] For example, Hearst and Jenkins (1974), Anselme and Güntürkün (2019), Collins et al. (1983), Berridge and Robinson (2003), Anselme et al. (2013), and Robinson et al. (2014).

Beyond probing the nature of the mechanisms underlying the hypothesized Pavlovian bias, this uncertainty test serves to strengthen the evidence for the Pavlovian bias itself: intensification with uncertainty is by and large uniquely Pavlovian—*not* predicted by non-Pavlovian explanations.[13]

2.2.3 Payoff rule.

Unlike conventional "pay-one" rules that randomly select a subset of trials for payment (Charness et al., 2016), we implement a "pay-all" rule: rewarding correct decisions and penalizing suboptimal ones (betting during yellow background sessions) on each trial. This design choice maximizes the earnings gap between optimal play and suboptimal strategies in our setting, as validated through task simulations (see Appendix Appendix A for details). We kept the preset threshold undisclosed until task completion to mitigate "wealth effects" and maintain consistent motivation throughout.[14] As such, it is an important aspect of our incentivization scheme (detailed in Section 3.1.1).

3 Results

3.1 Experimental procedures

The experimental protocol was approved by the University of New South Wales Human Research Ethics Committee (project number HC210411). Informed consent was obtained from all participants.

3.1.1 Incentivization scheme and transparency of task instructions

Participants were 215 undergraduate students from the University of New South Wales (mean age: 21.6, range $18-30$; 120 males). Prior to the session, they reviewed online task instructions introducing the video gambling game and its rules. These instructions explained all aspects of the task, including color cues for gambling expected value (blue: positive; yellow: negative), conveyed through animated simulations to effectively communicate statistical information (e.g., Hoffrage et al. (2000) and Sanborn and Chater (2016)).[15]

Our recruitment ad emphasized the experiment's unique incentivization scheme (detailed in Section 2.1, "Payoff rule"), presenting the task as "special" with the potential for significant rewards or losses based on performance. The online instructions further highlighted these stakes:

> *"How much can I expect to earn? Optimal play yields median earnings of* \$75*, meaning* 50% *of optimal players get more than* \$75*; mean is* \$55*. Note we designed the task and set the payment rule taking into account the fact that not all participants will play*

[13] A few notable exceptions exist, such as uncertainty-enhanced arousal, which could in principle generate similar patterns. However, we show in Section 3.4.2 that sign-tracking provides a superior account of our data, offering both better predictive power and greater parsimony.

[14] We borrowed this design feature from prior experimental work that used similar payoff rules. For example, Payzan-LeNestour et al. (2013), Payzan-LeNestour and Bossaerts (2015), and Payzan-LeNestour (2018).

[15] The instruction script is available here, alongside a video demonstration of its online delivery (link).

optimally. Our intended average payment across participants is around \$40*, which is about twice the average used in BizLab. We think the higher average payment is fair given Point 2 above, namely, the fact that some participants will likely end up losing money at the end of the game."*

This emphasis on stakes in the pre-task instructions aims to replicate real-world incentives in the lab, complementing our experimental design where optimal performance yields substantially better outcomes (as explained in Section 2.2.3). Effectively conveying this distinctive aspect was essential to ensure participants' full engagement.

To comply with UNSW Human Research Ethics Committee requirements, we guaranteed a positive final payout for all participants through a post-task "wheel of fortune" that recovered any incurred losses. This recovery mechanism was not disclosed in the task instructions, preserving the intended incentive structure while adhering to ethics requirements and the *"no-deception rule"*.[16]

Participants were randomly assigned to the Test or Control Treatment upon lab arrival, unaware of the two-treatment design. After reviewing the online instructions, they were verbally reminded of the high stakes: significant earnings for top performers and possible losses from the game for poor performers.

The instruction phase included a practice session for task interface familiarization and a multiple-choice questionnaire (MCQ) to verify comprehension. Participants were warned that more than one mistake on the MCQ would result in immediate dismissal without payment. The experimenters verified all answers, allowing only one error for potential terminology misunderstandings.

Immediately before the task began, the experimenter verbally reiterated all key features, adhering to Shmaya and Yariv (2016)'s transparency criterion to eliminate ambiguity and prevent participants from forming arbitrary conjectures. Just prior to game initiation, the computer posed two questions: 1. Chance of a winning spin in a yellow session? $< 10\%$ / $10 - 30\%$ / $> 30\%$ / Not sure; 2. Game strategy? None / Have one, unsure / Fairly confident / Think it's right.

3.1.2 Matching procedure

To ensure an unbiased comparison of participant behavior across the two treatments, we matched participants' key demographic characteristics (age, gender, academic background). Participants completed a brief questionnaire, selecting from predefined categories for academic major (business school, science, other), gender, and age. We maintained balanced proportions across the two treatments for each characteristic (see Appendix Appendix C for details).

[16] Omitting the wheel of fortune feature in the task instructions qualifies as *"deception by omission"*, a practice widely accepted in economics as compatible with the no-deception rule when viable alternatives are lacking (Ortmann, 2020)—precisely our situation in this study. We conducted thorough post-task debriefings, explaining our omission of this feature and emphasizing the exceptional nature of such omissions in behavioral economics research.

3.1.3 Preregistered filter

As outlined in Section 2.2, we controlled for the faulty cognition factor through a strict filtering process. We excluded participants who failed any of these criteria: answering at least 5 out of 6 MCQ questions correctly; passing the pre-task gambling odds quiz (and post-task quiz when administered); reporting having a strategy for the task; and demonstrating correct beliefs in the majority of in-game lottery choices. Given the high incentives for task engagement, we anticipated minimal data loss through this filter, a prediction supported by the analyses reported in Section 3.2.

3.2 Summary statistics and preliminary tests

Our filter excluded 16 participants, yielding a final sample of $N = 199$. Table 1 and Fig.3 summarize participant behavior across treatments. Participants bet nearly universally in blue background sessions (mean: 96%, median: 100%), consistent with optimal behavior. The Test Treatment group exhibited the hypothesized spillover effect, with a positive betting rate in yellow background sessions [participant-level t-test ($N = 99$): $t(98) = 5.94$, Cohen's D $= 0.59$, $p < 0.001$].

Table 1. Participant betting rates by session color and treatment.

	Blue		Yellow	
	Control	Test	Control	Test
Mean	.969	.963	.054	.098
Median	1	1	0	.008
SD	.097	.084	.113	.164
Min	.500	.583	0	0
Max	1	1	.558	.742
Skew	-3.795	-3.059	2.387	2.097

In the Test Treatment, 53% (53/99) of participants bet at least once during yellow background sessions, with 44% betting twice or more. Comparatively, in the Control Treatment, 36% (36/100) bet at least once, and 33% bet multiple times. Participant earnings show the expected bimodal distribution: 52% received \$110, while 15% earned \$10 or less (41% $\leq$ \$15; mean: \$65).

Turning to our main hypotheses, the raw data show participants in the Test Treatment had higher average betting rates during yellow background sessions, exceeding the Control Treatment by 5 percentage points (10% vs 5%, see Fig.4A), and betting increased with uncertainty (Fig.4B). These preliminary results support both our hypotheses: first, that participants exhibited irrational betting patterns driven by a Pavlovian bias and, second, that this bias partly operated through a sign-tracking mechanism.

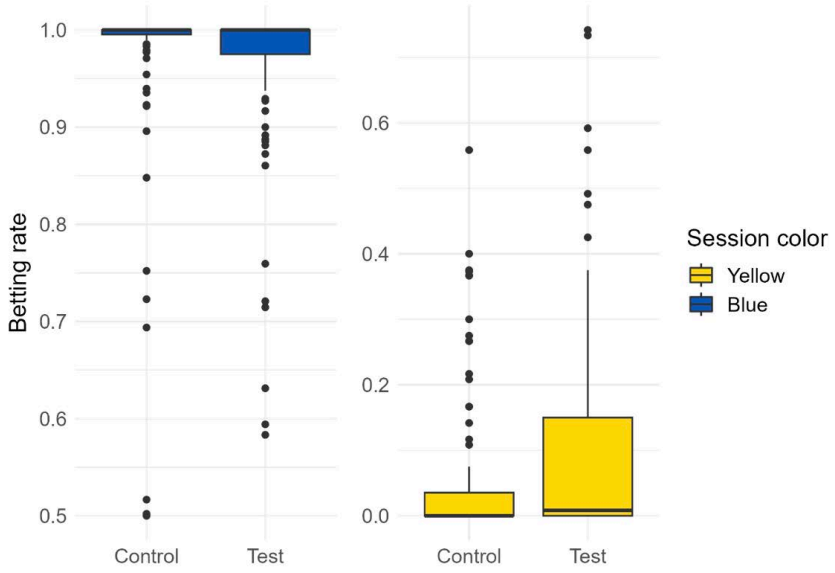


Figure 3. Distributions of participant betting rates in the experiment.

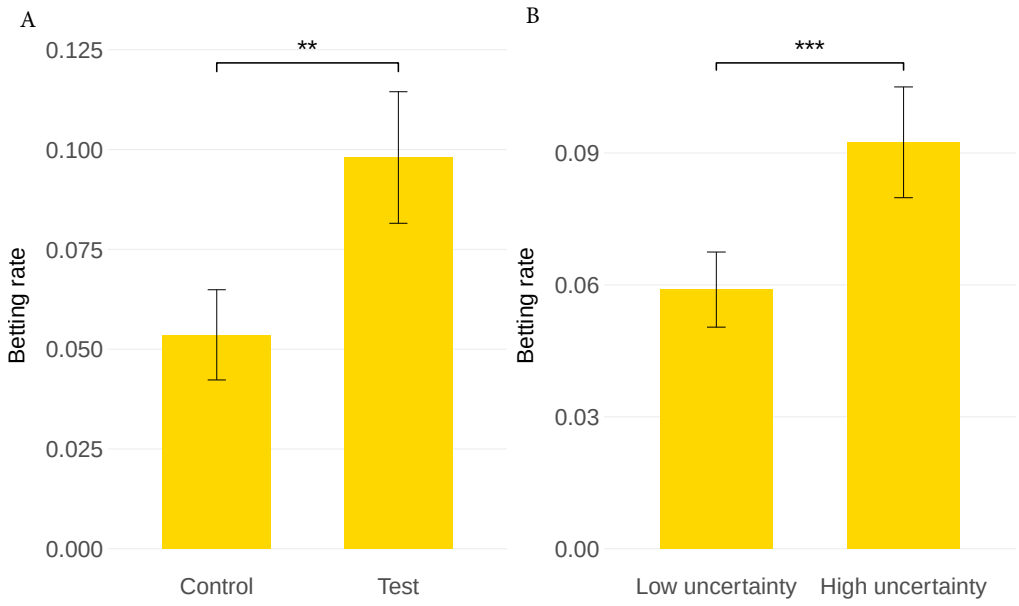


Figure 4. Treatment and uncertainty effects. (A) Two-sample t-test comparing participant betting rates in yellow background sessions between Control and Test Treatments. Ho: average betting rate in Test Treatment $\leq$ Control Treatment. $T(173.792) = 2.35$, effect size (Cohen's D) $= 0.33$, $p = 0.009$. Due to data skewness (1.402, outside the ideal $[-1, 1]$ range) and non-normality confirmed by Shapiro-Wilks testing, we ran the test on Box-Cox transformed data ($\lambda = -13.5$), which reduced skew to 0.738. Non-parametric test on raw data confirms conclusions [Wilcoxon test, $V = 5874$, $p = 0.006$, effect size (rank-biserial) $= 0.19$, one-tailed]. **(B)** Paired t-test comparing betting rates between high-uncertainty and low-uncertainty trials during yellow background sessions, at the participant level. H0: participant betting rate in high-uncertainty trials $\leq$ low-uncertainty trials. $T(198) = 3.25$, effect size (Cohen' s D)$= 0.23$, $p < 0.001$. Data were Box-Cox transformed to reduce skew ($\lambda = -10$). Non-parametric test on raw data confirms conclusions [paired Wilcoxon test, $V = 2912$, $p < 0.001$, effect size (rank-biserial) $= 0.56$, one-tailed].

3.3 Main regression results

Logistic mixed-effects regressions confirm both treatment and uncertainty effects while accounting for potential confounds, including status-quo bias through the participant's previous decision (Tables 2 A and 2 B, "Previous choice" variable). These regressions, which include detailed robustness checks (summarized in Appendix Appendix D), further reveal that harmful betting decreases over the course of the task (see Table 2 B, "Sequence number" variable). This decline aligns with patterns reported in recent studies, which suggest that Pavlovian influences diminish as participants gain task experience and develop a greater sense of control over outcomes.[17]

Table 2. Logistic mixed-effects regressions. * < 0.05, ** < 0.01, *** < 0.001. Standard errors in parentheses. Beta coefficients show log odds ratios, with effect sizes categorized as small (0.519), medium (1.24), and large (1.90) per Chen et al. (2010). Models include by-participant random intercepts for subject-wise error control. All continuous variables are z-scored. Penalized log-likelihood addresses limited sample sizes in some participant subgroups (Firth, 1993). R^2 values follow Nakagawa & Schielzeth's conditional metric for mixed-effects models (Nakagawa and Schielzeth, 2013). **(A) Whole sample analysis.** Uses all participants, all trials. Dependent variable: participant decision (0: Pass, 1: Spin). Independent variables are demographics (gender, age, and academic discipline), reward value (0: low, 1: high), uncertainty (0: low, 1: high), Treatment (0: Control, 1: Test), session color (0: blue, 1: yellow), session color × Treatment interaction, previous choice (0: Pass, 1: Spin), reward exposure (sum of rewards from previous trials), sequence number (1-6), and reaction time. Key variables of interest: Treatment × session color interaction and uncertainty. Positive coefficients support pre-registered hypotheses. **(B) Maladaptive bettors' analysis.** Yellow sessions only, participants with minimum two bets. Dependent variable: participant decision (0: Pass, 1: Spin). Independent variables: as defined in (A). Key variables of interest: Treatment and uncertainty. Positive coefficients support pre-registered hypotheses.

A

	Dependent Variable: betting
Intercept	2.911 (.268) ***
Uncertainty	.154 (.037) ***
Treatment × Color	1.495 (.135) ***
Treatment	-.831 (.227) ***
Colour	-8.527 (.117) ***
Previous choice	1.901 (.039) ***
Reward value	1.561 (.042) ***
Reward exposure	-.017 (.024)
Sequence number	-.030 (.019)
Reaction Time	.085 (.019) ***
Nakagawa R^2	.837
N	119,096

B

	Dependent Variable: harmful betting
Intercept	-4.203 (.310) ***
Uncertainty	1.285 (.110) ***
Treatment	.822 (.292) **
Uncertainty × Treatment	-.879 (.137) ***
Previous choice	1.172 (.067) ***
Reward value	1.436 (.072) ***
Reward exposure	-.065 (.052)
Sequence number	-.208 (.034) ***
Reaction time	.428 (.035) ***
Nakagawa R^2	.472
N	9224

[17] For example, Dorfman and Gershman (2019), Gershman et al. (2021), and Payzan-LeNestour and Doran (2024).

3.4 Model comparison analyses

3.4.1 Further evidence of Pavlovian bias.

Behavioral models We compared three models of participant behavior: a rational-choice base model, its Pavlovian-augmented variant, and a model-free Reinforcement Learning (RL) model that behaves according to the Law of Effect.

Base model. The rational agent adjudicates between the gambling (G) and pass options based on gambling value:

$$U(G) = p(x - 0.7)^{\alpha} - (1 - p)\lambda \cdot 0.7^{\alpha}, \tag{1}$$

where p denotes the probability of a winning spin (blue sessions: 0.8; yellow: 0.2), and x denotes the reward for a winning spin (1 or 2). $\lambda \geq 1$ is *Prospect Theory*'s loss aversion parameter and α governs the agent's risk attitude ($0 < \alpha \leq 1$). When $\lambda = \alpha = 1$, the agent is risk/loss neutral, and in this case, $U(G)$ reduces to the expected value: $p \cdot x - 0.7$.

To model decision, we use the popular logit/softmax rule (chiefly for simplicity, and without loss of generality). The agent chooses to gamble with probability:

$$P(G) = \frac{e^{\beta U(G)}}{1 + e^{\beta U(G)}}. \tag{2}$$

The parameter β governs choice randomness or "trembling" ($\beta = 0$ gives completely random choice; $\beta \to \infty$ gives deterministic choice).

Pavlovian model. Let t denote the current trial, and $1_x(k)$ denote the indicator function taking value 1 if the agent wins x at trial k, and 0 otherwise. The Pavlovian system's attraction to gambling, denoted by $DA(t)$, is the following function of the Pavlovian cumulative reward value $V(t)$:[18]

$$DA(t) = \max\left(f\left(V(t)\right), 0\right), \text{ with } \begin{cases} f : u \to \frac{1}{1+e^{-\kappa_1 u}} - \kappa_2 \\ V(t) = \sum_{k=1}^{t} \theta^{t-k} 1_x\left(k\right) \end{cases}, \tag{3}$$

where $0 \leq \kappa_1 \leq 1$, $0.5 \leq \kappa_2 \leq 1$, and $0.9 \leq \theta \leq 1$. The probability to gamble at trial $t + 1$ is distorted as shown in Fig.5. Formally:

$$P(G)^{\text{Pavlovian}}(t + 1) = DA(t) + (1 - DA(t))P(G), \tag{4}$$

where $P(G)$ is the gambling probability that prevails in the absence of Pavlovian influence (Eq. (2)).

[18] For details on model specification and parameter interpretations, see Payzan-LeNestour and Doran (2024).

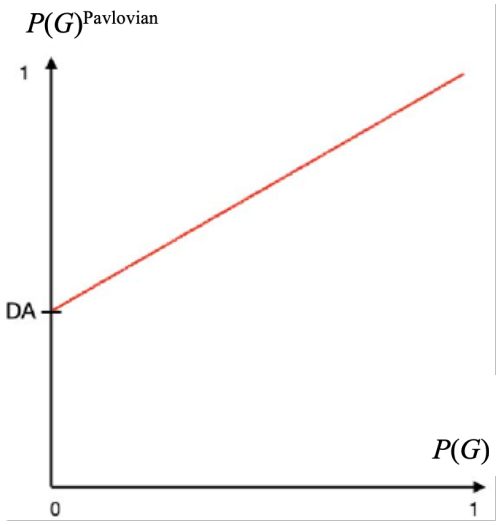


Figure 5. Pavlovian bias. The factor DA(t) defined in Eq. (3) takes value "DA" in this example. X-axis: $P(G)$, the rational agent's probability to gamble. Y-axis: $P(G)^{\text{Pavlovian}}$, the distorted probability to gamble (Eq. (4)). For $P(G) < 1$, the Pavlovian influence increases the probability of gambling.

Crucially, while the Pavlovian system's attraction to gambling ($DA(t)$) depends on the cumulative reward value ($V(t)$), it is not equivalent to it. This model feature (inspired by Zhang et al. (2009)) means that the same Pavlovian value can lead to different levels of attraction to gambling depending on the agent's motivational state, captured by the κ parameters. Specifically, higher values of κ_1 (or lower values of κ_2) strengthen this attraction. Research by Dorfman and Gershman (2019) and Gershman et al. (2021) points to perceived controllability as being a key determinant of these κ parameters. Lower perceived control over outcomes increases κ_1, which generates declining Pavlovian influences over the course of experimental tasks (as discussed in Section 3.3). Additional determinants of the κ parameters may include agents' current appetite for monetary rewards (which increases with factors such as deprivation), with higher appetite leading to higher κ_1, and, as we propose in Section 3.4.2, uncertainty levels, which when elevated also increase κ_1.

RL model. Rather than computing the gambling expected value, the RL agent updates the gambling value from observed outcomes when gambling is chosen (the value estimate remains unchanged after pass choices), starting from an initial value estimate $U_0(G)$:

$$U_t(G) = U_{t-1}(G) + \eta \cdot (x_t - U_{t-1}(G)),$$

where x_t denotes the observed gambling outcome at trial t and η denotes the learning rate ($0 < \eta < 1$; higher rates place more weight on observed outcomes relative to the starting-point value). For a given gambling value, the RL agent makes choices like the base model (Eq. (2)) where the value of skipping is 0.

Comparing predicted and observed betting rates. Task simulations reveal stark contrasts between model predictions and average participant behavior (Table 3). The RL model proves inadequate, underestimating blue session betting while overestimating yellow session betting—reflecting RL agents' tendency to bet early but cease after initial losses. The base model accurately predicts near-universal blue session betting but misses the key behavioral anomaly: 5 percentage points excess betting in Test Treatment yellow sessions versus Control (10% vs 5%). Only the Pavlovian model captures both patterns through its Pavlovian mechanism: high win probability (0.8) in blue sessions enhances the Pavlovian value of betting in the subsequent yellow session, a bias absent in the Control Treatment where yellow sessions precede blue ones by design.

Table 3. Average betting rates by treatment, color, and agent type. Simulations used 99 agents per treatment for each model: rational choice ("base"), Pavlovian model ("Pavlovian"), and Reinforcement Learning ("RL"). Agents generated trial-by-trial action values, computed choice probabilities using Eq. (2) (base and RL models) or Eq. (4) (Pavlovian model), and made trial-specific decisions. Model parameters were calibrated to match the mean betting rates observed across lab participants. Since simulated RL agents could not simultaneously match the observed mean betting rates in both blue and yellow sessions, we used two parameter sets: Type 1 calibrated to match blue session betting rates, and Type 2 calibrated to match yellow session betting rates.

	Session Type: (Color, Treatment)			
Agent-type	Blue, Control	Blue, Test	Yellow, Control	Yellow, Test
RL (Type 1)	0.93	0.97	0.92	0.97
RL (Type 2)	0.1	0.21	0.15	0.05
Base	0.97	0.97	0	0
Pavlovian	0.97	0.97	0	0.05
Lab Data	0.97	0.97	0.05	0.1

Note: Our Pavlovian model calibrations specifically target the excess yellow betting observed in Test versus Control, rather than absolute betting rates. As such, the model fails to predict the baseline 5% yellow betting rate in the Control Treatment, which likely stems from the non-Pavlovian factors hypothesized in Fig.2. By incorporating one such factor—loss tolerance—into the Pavlovian model, we can match both the baseline 5% betting rate in Control and the elevated 10% rate in Test, as shown in Table 8 of Appendix Appendix E.

Participant-level model comparisons We fitted each model to individual participants by generating trial-by-trial action values and choice probabilities, and maximizing the log-likelihood (LL) of observed choices in the training dataset (sequences 1, 3, and 5).[19] We then evaluated model

[19]The base model has three parameters (α, λ, β), to which the Pavlovian variant adds κ_1, κ_2, and θ. The RL model has five: β, η and three initial values $U_0(G)$ (one per training sequence).

performance using out-of-sample LL computed on the remaining sequences as our measure of model goodness of fit.

We opted for this two-step procedure for two reasons. First, it effectively penalizes overfitting (Ho et al., 2007, 2008). Second, it possesses adequate "*recovery*" properties in the sense of Palminteri et al. (2007): a model outperforms its competitors if and only if the data exhibit behavioral patterns that constitute the model's behavioral signature. This property ensures adequate power while minimizing false positives. To verify it, we conducted extensive analyses with artificially generated data. The main findings are as follows (for more details, see Appendix Appendix E.1). When simulating 99 participants per treatment using the Pavlovian model, the latter outperformed the base model for about 90% of the participants exhibiting harmful betting, demonstrating adequate power to detect Pavlovian effects when present. Conversely, when simulating 99 rational agents, the Pavlovian model never outperformed, showing satisfactory "*restrictiveness*" in the sense of Fudenberg et al. (2023). Similar recovery rates emerged in RL versus base model comparisons, confirming that each model exhibits distinct and identifiable behavioral signatures in our task.

Comparing the models' out-of-sample LL revealed two key findings about participant behavior. First, participants exhibited model-based rather than model-free behavior, with the base model outperforming RL (Fig.6A[20]). Second, participants exhibited behavior consistent with our hypothesized Pavlovian bias, with the Pavlovian model outperforming the base model (Fig.6B[21]).

Next, we quantified the relative evidence for each model. For each participant, we computed the posterior probability that they behaved according to each model, inspired by Dal Bó and Fréchette (2011) and Kendall and Oprea (2024):[22]

$$\hat{p}_n(k) \propto e^{L_n(k)} \cdot p_{n,k}, \qquad \sum_{k=1}^{3} \hat{p}_n(k) = 1, \tag{5}$$

where k indexes the three models (RL, Base, Pavlovian) as $k = 1, 2, 3$, $p_{n,k}$ denotes the prior probability that participant n adopted model k, and $L_n(k)$ denotes the cumulative out-of-sample LL of model k for participant n. Following Dal Bó and Fréchette (2011), we also computed each model's "*importance*" across the participant group, here defined as its average posterior probability across participants:

$$\tau(k) = \frac{1}{N} \sum_{n=1}^{N} \hat{p}_n(k), \tag{6}$$

where N is the size of the participant group.

The evidence strongly favors the Pavlovian model both at the individual and group levels: its posterior probability is close to 1 for most maladaptive bettors (Fig.6C), and it dominates at the

[20] Paired Wilcoxon test comparing goodness-of-fits ($N = 77$): $V = 35$, $p < 0.001$, effect size (rank-biserial) $= 0.98$.
[21] Paired Wilcoxon test comparing goodness-of-fits ($N = 77$): $V = 190$, $p < 0.001$, effect size (rank-biserial) $=0.87$.
[22] Dal Bó and Fréchette (2011)'s "*Strategy Frequency Estimation Method*" is designed for games with deterministic strategies, binary actions, and a single free parameter per task, so we adjusted it to our setting.

group level (6D).[23]

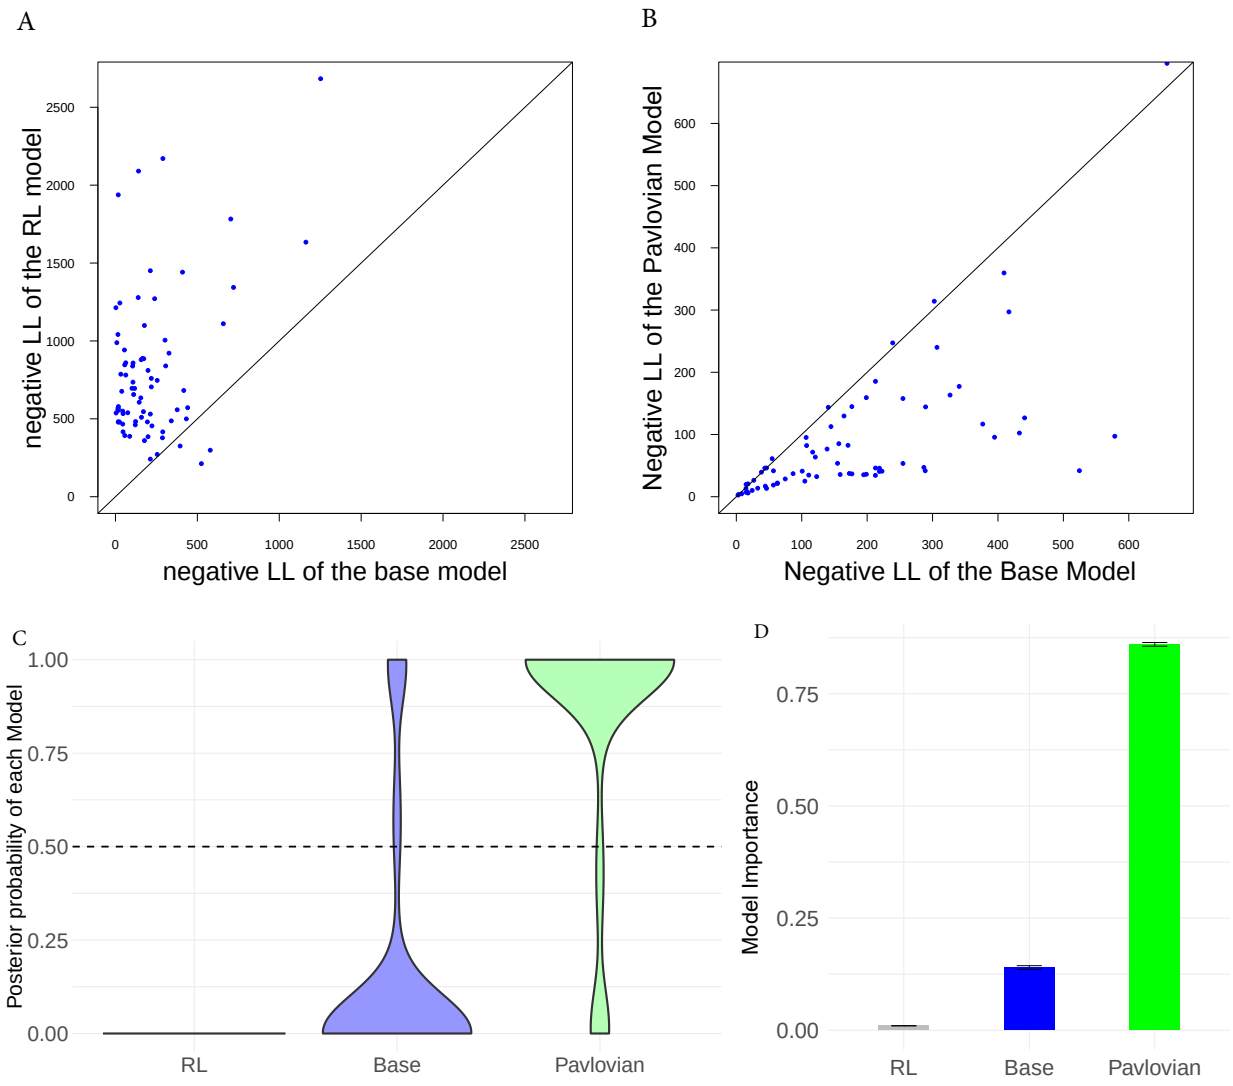
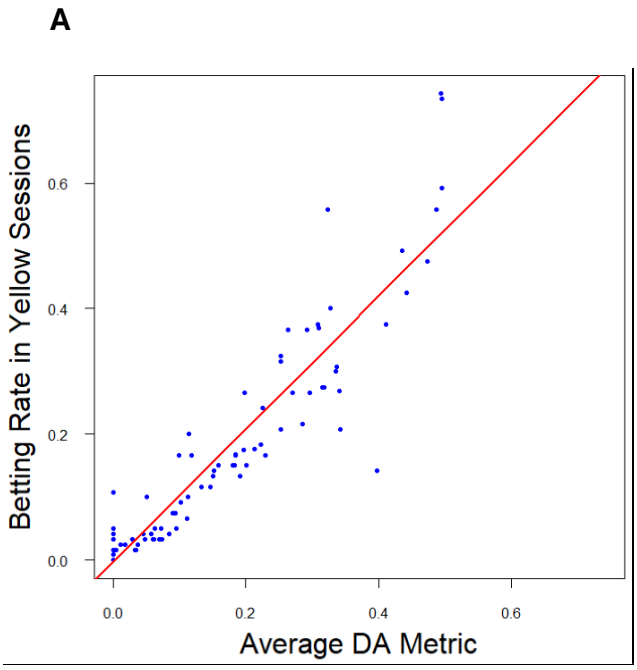


Figure 6. Model fit comparisons for participant behavior. **(A)** Out-of-sample negative LL: base model (x-axis) vs. RL model (y-axis) per maladaptive bettor in the lab ($N = 77$). **(B)** Out-of-sample negative LL: base model (x-axis) vs. Pavlovian model (y-axis). In both **(A)** and **(B)**, points above the 45-degree line indicate a superior fit for the base model. **(C)** Posterior probability that participant behavior followed each model, computed at the participant level using Eq. (5). A probability above 0.5 indicates evidence favoring that model, with values closer to 1 indicating stronger evidence. The violin plot's width represents the density of participants at each probability value. **(D)** Each model's importance among maladaptive bettors. Importance is defined as average posterior probability across participants (Eq. (6)).

The estimated Pavlovian influence predicts maladaptive betting. Individual differences further support the role of Pavlovian influences in maladaptive gambling: Fig.7 shows that par-

[23]The posterior probabilities reported in these figures were computed using an agnostic prior ($p_{n,k} = 1/3$ for all n, k). Similar results emerged in separate analyses with different priors $p_{n,k}$, including conservative ones that placed low weight on the Pavlovian model.

ticipants' estimated Pavlovian attraction predicts their yellow background betting rates, a direct measure of maladaptive gambling in our task.



Regression Output

Variables	Coefficient	Std. Error	p-value
Intercept	-0.0015	0.0037	0.6770
Average DA	1.0539	0.0249	<2e-16
R^2	0.9017	F-statistic	1797
Adjusted R^2	0.9012	Model p-value	< 2.2e-16
Residual Std. Error	0.0447		

Figure 7. Pavlovian attraction to gambling predicts yellow-session (maladaptive) betting. Each point represents one participant ($N = 199$). X-axis: mean DA metric averaged across yellow-session trials, computed using Eq. (3) and participant-fitted values of κ_1, κ_2, and θ. Y-axis: betting rate in yellow sessions. The best-fit line shows an intercept close to 0 and a slope close to 1 (see "Regression Output"), consistent with the theoretical prediction (Eq. (4) and Fig.5).

Robustness checks. The above findings persist across multiple specifications: in both the complete sample and the subset of maladaptive bettors (those betting at least twice across yellow sessions); under various assumptions about risk preferences (allowing for loss tolerance: $\lambda < 1$); and with an RL model augmented for "undirected exploration" (Sutton, 2018) to address the RL model's low blue-session betting, among other robustness checks detailed in Appendix Appendix F.

3.4.2 Nature of the Pavlovian bias.

Drawing on both regression analyses and model comparisons, our findings so far point to the possibility that a Pavlovian bias drove harmful betting in our task, and that this bias partly operated through sign-tracking mechanisms. As stressed earlier, a defining characteristic of sign-tracking lies in its paradoxical relationship with uncertainty: the Pavlovian attraction to a reward cue strengthens as this cue becomes *less* reliable in predicting reward. Therefore, if sign-tracking significantly contributed to the Pavlovian bias in our task, we may observe stronger Pavlovian responses when play availability was uncertain (high-uncertainty sessions) compared to when it was guaranteed.

We found support for this prediction. To test it, we extended the Pavlovian model by introducing two distinct κ_1 parameters, one for high-uncertainty sessions, one for the others. Recall κ_1 is the critical 'gating' parameter governing the Pavlovian attraction to gambling (see "Pavlovian model", p.16). The mean fitted value of κ_1 among maladaptive bettors is higher in the high-uncertainty sessions (Fig.8A[24]). Moreover, 61% of maladaptive bettors showed an unambiguously[25] higher κ_1 in the high-uncertainty sessions, suggesting sign-tracking affected a significant share—likely over half—of the maladaptive bettors.

Motivated by these findings, we refined the Pavlovian model into a more specific sign-tracking variant where κ_1 is no longer a free parameter but a function of the uncertainty associated with gambling:

$$\kappa_1(t) = C_{\kappa_1} \cdot H(X_t), \tag{7}$$

where C_{κ_1} is a free parameter and $H(X_t)$ is the information entropy, quantifying the average level of uncertainty associated with the reward variable X_t at trial t:

$$H(X_t) = -\sum_{x \in S_t} p_t(x) \log p_t(x),$$

where X_t takes values in the outcome set S_t and is distributed according to $p_t(\cdot)$. By design, information entropy is 0.94 for high-uncertainty trials and 0.5 for low-uncertainty trials.

This sign-tracking variant outperforms the fixed-κ_1 Pavlovian model in two ways. First, it successfully predicts the observed uncertainty effect: higher betting rates in high-uncertainty sessions (Table 4). Second, it provides superior fit at the participant level (Fig.8B). The model's importance score corroborates our earlier finding suggesting that sign-tracking affected more than half of the maladaptive bettors (Fig.8C).

Competing explanations. The behavioral literature suggests two alternative explanations to sign-tracking that could explain the uncertainty effect. The first draws on "*frustration theory*" (Amsel and Roussel, 1952; Amsel, 1992): violations of reward expectations generate frustration, leading to enhanced responding when a systematically delivered reward is suddenly withheld. Whether this enhanced responding extends to contexts of reward uncertainty remains debated (Anselme and Güntürkün, 2019). Assuming that it does, and further assuming that gambling denial generated frustration in our participants, one predicts increased betting after gambling denial. The data support this prediction: betting rates are significantly higher following gambling denial (82% versus 50%). However, the uncertainty effect reported in Fig.4 persists even after excluding these post-denial trials (Fig.9A). Thus, while frustration may have contributed to heightened betting in the high-uncertainty sessions of our task, it cannot fully account for the uncertainty effect.

[24] Paired Wilcoxon test comparing goodness-of-fits ($N = 77$): $V = 2067$, $p < 0.001$, effect size (rank-biserial) = 0.66.

[25] We classified participants as displaying the uncertainty effect only if their κ_1 was at least 50% larger in the high-uncertainty trials. This conservative threshold ensures both adequate power and low false positive rates (see Appendix Appendix E.2 for details).

Table 4. Average betting rates in yellow sessions by uncertainty level and agent type. Simulations used 99 agents per treatment across three models: rational choice (base), standard Pavlovian, and the more specific sign-tracking variant of the Pavlovian model. Agents generated trial-by-trial action values, computed choice probabilities using Eq. (2) (base model) or Eq. (4) (Pavlovian models), and made trial-specific decisions. Model parameters were calibrated to match the mean betting rates observed across lab participants. Results shown assume risk-neutrality for simplicity, but the key finding remains robust across different risk attitudes: only the sign-tracking variant of the Pavlovian model can generate the uncertainty effect (higher betting under high uncertainty).

	Yellow Sessions	
Agent-type	Low-uncertainty	High-uncertainty
Base	0	0
Original Pavlovian	0.03	0.03
Sign-tracking variant	0.01	0.04
Lab Data	0.06	0.09

The second competing mechanism builds on a variant of "*sensation seeking theory*" (Zuckerman, 2007; Conlisk, 1993). To challenge our main sign-tracking hypothesis, we explored whether sensation seeking could explain the observed higher gambling in the high-uncertainty sessions. The traditional model by Conlisk (1993) fails to account for this effect: it predicts that gambling appeal increases with outcome variance, but our high-uncertainty sessions have lower variance by design. However, by modifying the theory to focus on entropy rather than variance—inspired by so-called "*suspense theory*" (Lehne and Koelsch, 2015), we could construct a variant of Conlisk (1993)'s model that generates the observed pattern without invoking sign-tracking. Specifically, if utility increases with outcome unpredictability, gambling becomes more appealing in the high-uncertainty sessions:

$$U^{suspense}(G) = U(G) + SRA \cdot H(X_t), \tag{8}$$

where $U(G)$ is the gambling utility of rational agents (Eq. 1), SRA is a constant quantifying suspense-related arousal, and $H(X_t)$ is the information entropy defined in Eq. (7). However, this suspense-augmented model fails as a credible alternative to sign-tracking: since suspense operates equally in both treatments, it cannot explain the excess yellow betting observed in the Test Treatment (see "Main treatment effect", p.9, and Fig.2). Therefore, to give suspense theory its strongest chance in the 'horse race' against the sign-tracking hypothesis, we develop a more comprehensive "Pavlovian + suspense" model, combining suspense-related arousal (Eq. 8) with the Pavlovian influence (Eq. 4). Unlike the simple suspense model, this augmented version can generate both effects we observe in the data: the treatment effect (through its Pavlovian component) and the

uncertainty effect (through its suspense component).

The evidence strongly favors the sign-tracking model over the "Pavlovian + suspense" model at both individual and group levels. At the individual level, a paired Wilcoxon test comparing model fits ($N = 77$) reveals a large effect: $V = 2719$, $p < 0.001$, effect size (rank-biserial) $= 0.81$.[26] The group-level analysis confirms this superiority, with Fig. 9B showing higher importance for the sign-tracking model.[27] This superior performance stems from the alternative model's tendency to predict excessive betting in blue sessions while also suffering from overfitting.

[26] For reference, rank-biserial effect sizes of 0.1, 0.3, and 0.5 are conventionally considered small, medium, and large effects, respectively (Cohen, 2013).

[27] Our analysis used equal prior probabilities for both models, though "*Occam's Razor*" principle would mandate favoring the simpler sign-tracking model, making our results even more compelling.

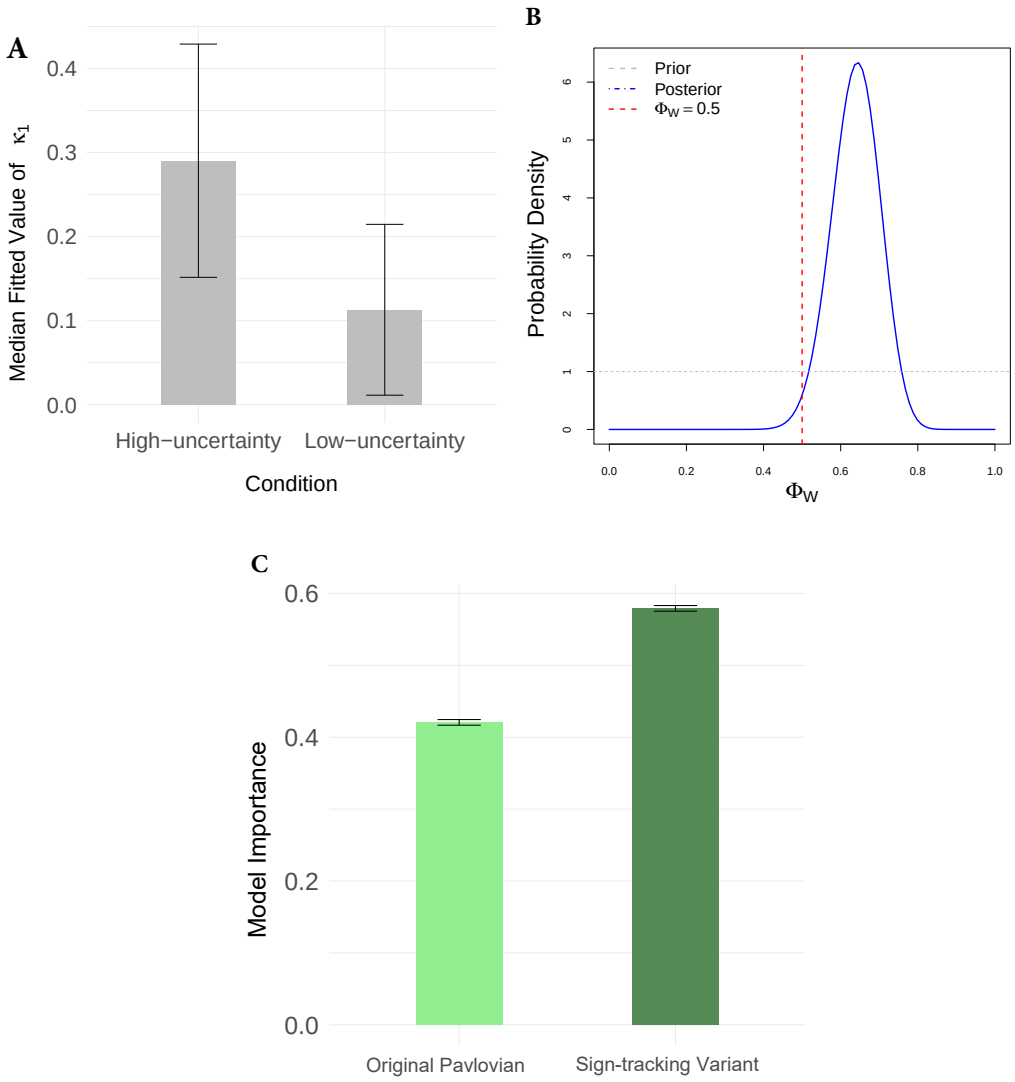


Figure 8. Evidence for sign-tracking. **(A)** Fitted κ_1 values among maladaptive bettors ($N = 77$), show higher medians in high-uncertainty versus low-uncertainty sessions, suggesting uncertainty intensifies the Pavlovian system's appetite for gambling. Due to non-normal distribution, bars display medians with interquartile ranges (25th-75th percentiles). **(B)** Bayesian Wilcoxon test comparing sign-tracking (Eq. (7)) versus original (fixed-κ_1) Pavlovian models' fits across maladaptive bettors. The test statistic $\Phi_W \in [0, 1]$ represents relative model performance based on signed ranks of paired fit differences, with $\Phi_W > 0.5$ favoring the sign-tracking model. Starting from a uniform prior over $[0, 1]$, the posterior (beta distribution: $a = 37, b = 21$) yields a 98% posterior probability of a superior fit by the sign-tracking variant ($P(\Phi_W > 0.5) = 0.98$) with a Bayes factor of 63 (> 30 indicates very strong evidence). We used the Bayesian variant of the Wilcoxon test as high kurtosis made the standard variant unreliable—both with the lab data and the synthetic data generated as part of our validation tests (see Appendix Appendix E.2 for details). **(C)** Model importance scores (computed using Eq. (6) with equal prior probability for each model) comparing the original Pavlovian model (fixed-κ_1) to its sign-tracking variant (Eq. (7)) across maladaptive bettors.

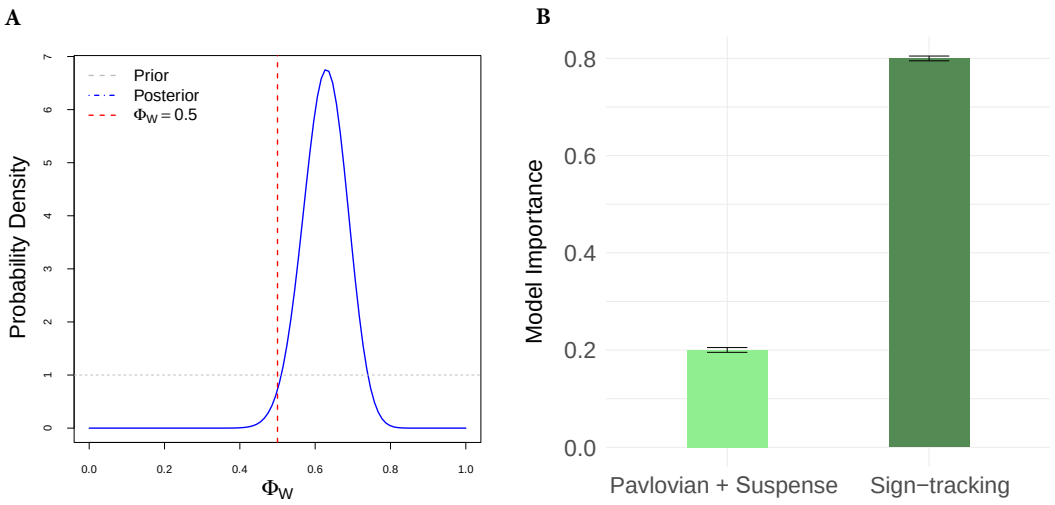


Figure 9. Evidence for sign-tracking against competing explanations. **(A)** Bayesian Wilcoxon test comparing the betting rates between high-uncertainty and low-uncertainty trials during yellow background sessions, at the participant level *after excluding post-denial trials*. The test statistic $\Phi_W \in [0, 1]$ represents relative betting frequency based on signed ranks of paired fit differences, with $\Phi_W > 0.5$ indicating higher betting rates in high-uncertainty trials. Starting from a uniform prior over $[0, 1]$, the posterior (beta distribution: $a = 42, b = 25$) yields a 98% posterior probability of higher betting rates in high-uncertainty trials ($P(\Phi_W > 0.5) = 0.98$) with a Bayes factor of 54 (> 30 indicates very strong evidence). We used the Bayesian variant of the Wilcoxon test as high kurtosis (23) made the standard variant unreliable. **(B)** Model importance scores (computed using Eq. (6) with equal prior probability for each model) comparing the "Pavlovian + suspense" model to the sign-tracking model (Eq. (7)) across maladaptive bettors.

4 Discussion

While no single piece of evidence is decisive, the collection of findings reported in this paper supports the idea that hardwired Pavlovian mechanisms can disrupt decision-making even in sophisticated agents, potentially helping explain well-documented phenomena such as investor inertia and specific manifestations of problem gambling in recreational gamblers, particularly persistent play and absorbed attention to gaming machines.

Understanding the mechanisms behind the hypothesized Pavlovian bias is critical for both market and clinical applications. Therefore, in this study, not only did we test for such a bias, but we took a first step toward examining whether its underlying mechanisms extend beyond seeking monetary rewards to include sign-tracking, an attraction to gambling cues themselves generating an apparent "utility of gambling" (Conlisk, 1993). Our findings suggest that sign-tracking may affect a significant share of problem gamblers, implying a need for dual-focused interventions in financial and gambling domains, whether online or physical: both recalibrating outcome expectations and redesigning environments to minimize the presence and salience of gambling-associated stimuli.

That said, as this study provides the first evidence for sign-tracking in a controlled monetary decision environment, and comes with inherent limitations of laboratory experiments, the results necessarily call for further testing. One such limitation concerns the identification of the underlying mechanism. Section 3.4.2 demonstrates the superiority of sign-tracking over competing explanations such as uncertainty-enhanced arousal, both in terms of predictive power and parsimony. However, as with any laboratory evidence, we cannot definitively rule out alternative mechanisms or experimental artifacts that may have escaped our attention. Future research should aim to replicate these findings and extend them through additional laboratory experiments, combining behavioral and neural measurements to validate and refine this initial evidence.

Appendix A Payoff rule

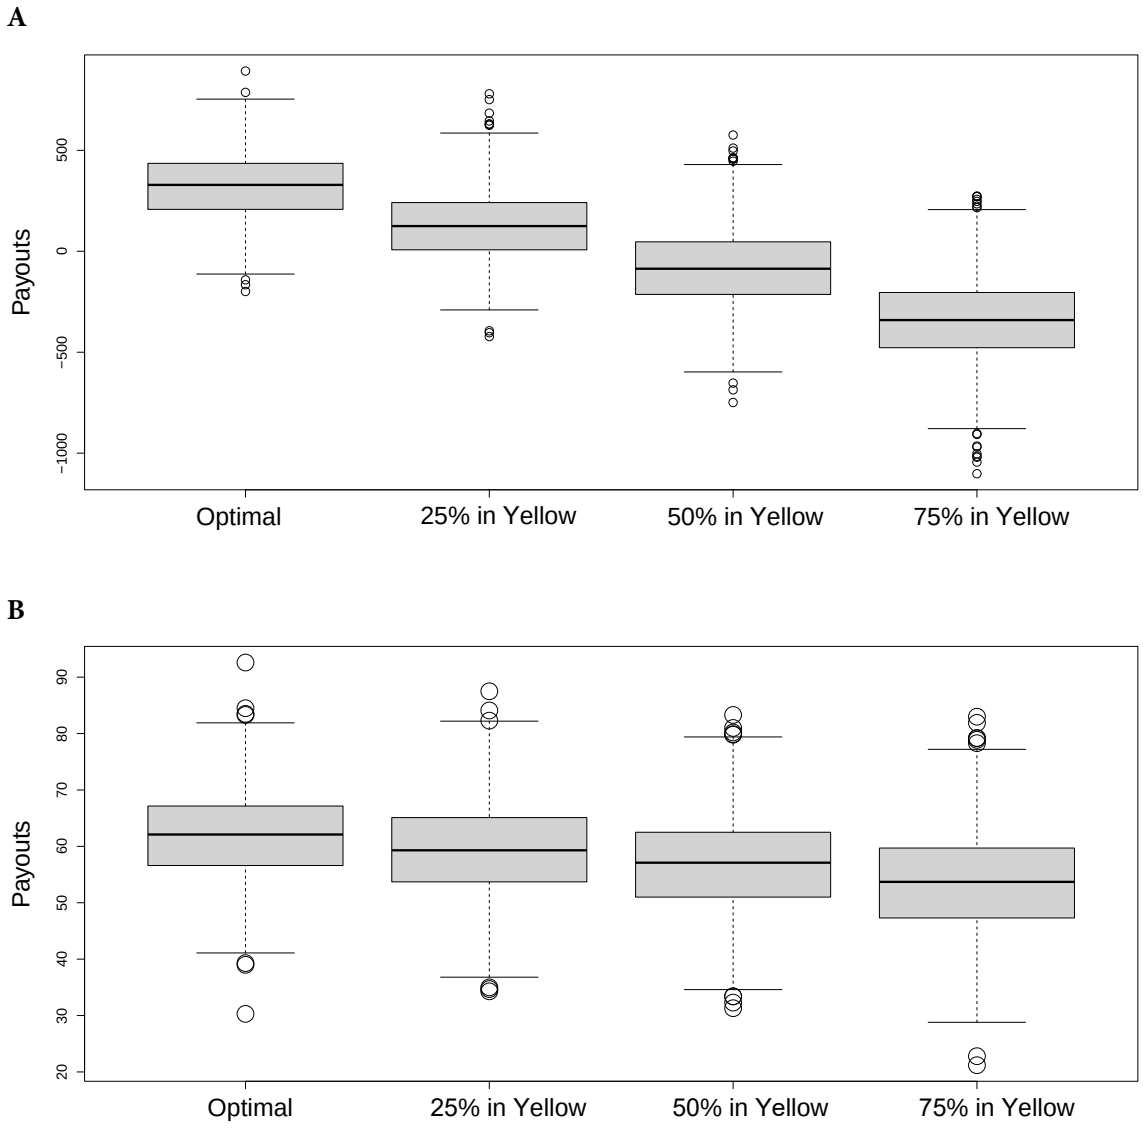


Figure 10. Incentive effects of alternative payment rules. Earnings distributions from 1,000 simulated task runs comparing (A) our experimental payment rule with (B) a "pay one" rule that randomly selects 30% of trials for payment. "Optimal": optimal risk-neutral strategy consists of betting in blue background sessions while passing in yellow background sessions. "25% in Yellow": suboptimal strategy consisting of betting in all blue sessions plus 25% of yellow sessions; "50% in Yellow" and "75% in Yellow" are defined in the same way. Under payment rule **(A)**, the probability of positive earnings declines sharply with suboptimal play: 89% for optimal strategy, falling to 46% when betting in 25% of yellow sessions, 16% when betting in 50%, and merely 1% when betting in 75%. In contrast, under payment rule **(B)**, all strategies achieve 100% positive earnings.

Appendix B Factorial design

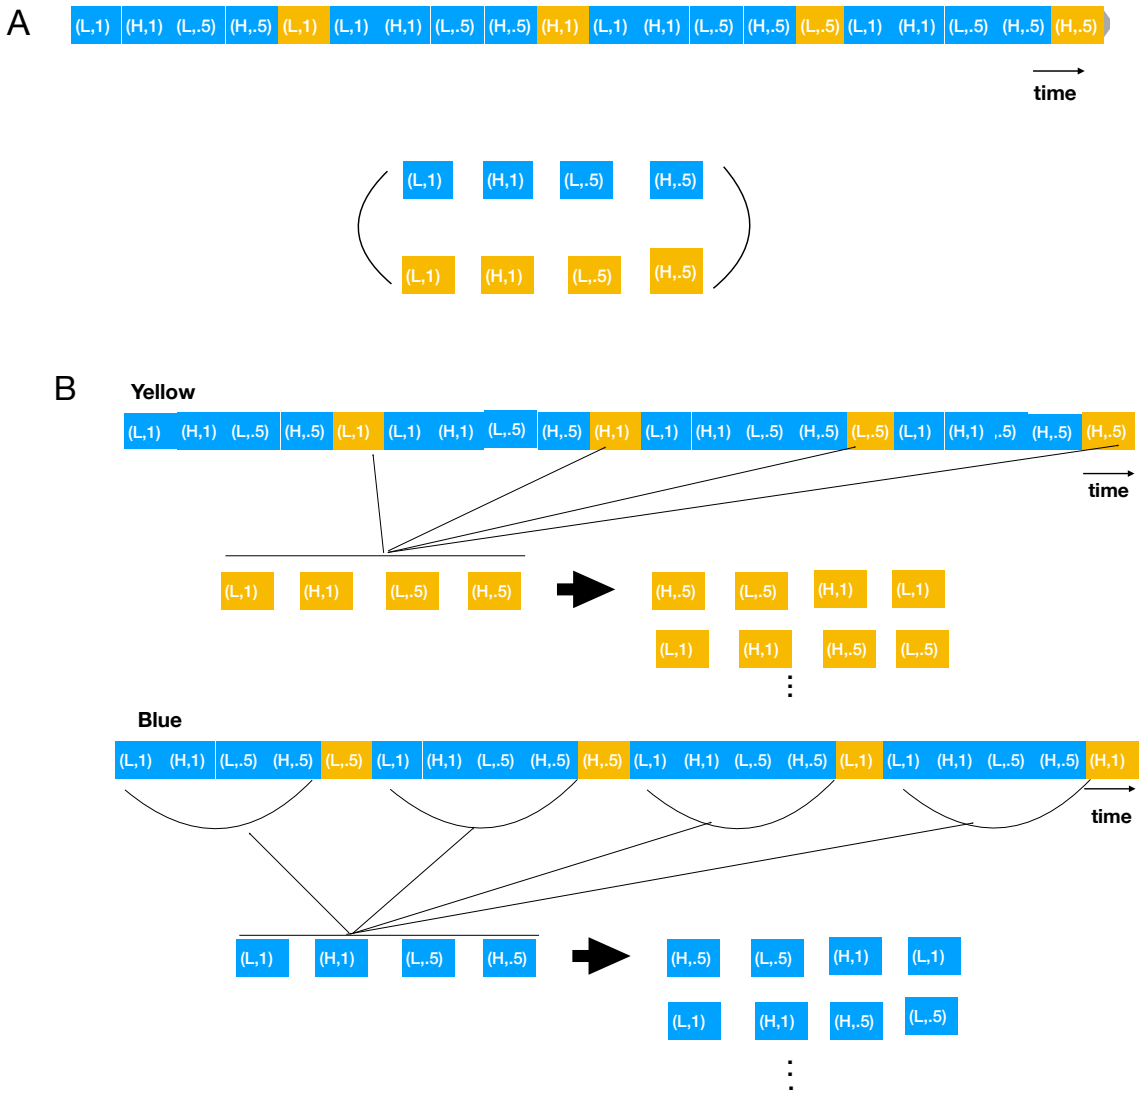


Figure 11. Factorial design. Each square represents one session. Blue: positive expected value for gambling. Yellow: negative expected value for gambling. Session types: (H,0.5): High reward ($2) for winning spin, 50% chance of play. (L,0.5): Low reward ($1) for winning spin, 50% chance of play. (H,1): High reward ($2) for winning spin, 100% chance of play. (L,1): Low reward ($1) for winning spin, 100% chance of play. **(A)** The factorial design ensures each of the 4 session types appears equally often in both blue and yellow sessions, as shown by the grouped boxes. **(B)** To prevent systematic order effects, session order is randomized within each color (yellow and blue), while maintaining the balanced representation of session types. The arrows show example randomizations.

Appendix C Matching population characteristics across the two treatments

Table 5 presents participant demographics for gender, age, and academic discipline, across the treatment groups. The table shows balanced proportions for each characteristic, ensuring comparability between groups.

Table 5. Summary of Participant Demographics by Treatment Group.

	Gender			Age			Discipline		
	F	M	Other	<20	20-25	>25	Business	Stem	Other
Control	39	60	1	8	88	4	68	31	1
Test	39	60	0	7	89	3	60	36	3
N	78	120	1	15	177	7	128	67	4

Appendix D Stepwise process of model building

Tables 6 and 7 present the incremental development of the regression models reported in Table 2, showing coefficient stability for our key variables across successive model specifications.

Table 6. Stepwise building of model from Table 2A. * < 0.05, ** < 0.01, * < 0.001. Standard errors in parentheses. See Table 2 for the legend. Regression (5) corresponds to the model reported in Table 2A.**

	Dependent Variable: choice				
	(1)	(2)	(3)	(4)	(5)
Intercept	4.616 (.309) ***	4.624 (.313) ***	4.632 (.308) ***	2.900 (.269) ***	2.911 (.268) ***
Uncertainty	.156 (.034) ***	.156 (.034) ***	.167 (.034) ***	.153 (.037) ***	.154 (.037) ***
Treatment	-.885 (.263) ***	-.882 (.265) ***	-.911 (.261) ***	-.823 (.227) ***	-.831 (.227) ***
Session colour	-9.991 (.113) ***	-9.991 (.114) ***	-9.983 (.114) ***	-8.504 (.113) ***	-8.527 (.117) ***
Reward value	1.820 (.041) ***	1.819 (.041) ***	1.825 (.041) ***	1.560 (.042) ***	1.561 (.042) ***
Treatment × Colour	1.881 (.130) ***	1.877 (.131) ***	1.859 (.131) ***	1.466 (.130) ***	1.495 (.135) ***
Sequence number	-	-.051 (.017) **	-.013 (.018)	-.029 (.019)	-.030 (.019)
Reaction Time	-	-	.180 (.019) ***	.086 (.019) ***	.085 (.019) ***
Previous choice	-	-	-	1.900 (.039) ***	1.901 (.039) ***
Reward exposure	-	-	-	-	-.017 (.024)
Nakagawa R^2	.840	.840	.841	.837	.837
N	119,096	119,096	119,096	119,096	119,096

Table 7. Stepwise building of model from Table 2B. * < 0.05, ** < 0.01, * < 0.001. Standard errors in parentheses. See Table 2 for the legend. Regression (4) corresponds to the model reported in Table 2B.**

	(1)	(2)	(3)	(4)
Intercept	-3.924 (.351) ***	-3.950 (.339) ***	-4.143 (.306) ***	-4.203 (.310) ***
Uncertainty	1.334 (.107) ***	1.403 (.109) ***	1.281 (.110) ***	1.285 (.110) ***
Treatment	.965 (.323) **	.909 (.313) **	.727 (.281) **	.822 (.292) **
Uncertainty × Treatment	-.910 (.132) ***	-.940 (.135) ***	-.875 (.137) ***	-.879 (.137) ***
Sequence number	-.352 (.032) ***	-.198 (.033) ***	-.208 (.034) ***	-.208 (.034) ***
Reward value	1.595 (.068) ***	1.608 (.070) ***	1.432 (.071) ***	1.436 (.072) ***
Reaction time	-	.525 (.034) ***	.428 (.035) ***	.428 (.035) ***
Previous choice	-	-	1.173 (.067) ***	1.172 (.067) ***
Reward exposure	-	-	-	-.065 (.052)
Nakagawa R^2	.459	.477	.472	.472
N	9224	9224	9224	9224

Appendix E Power and restrictiveness of conducted tests

We conducted an extensive set of validation analyses to verify that our model comparison tests (Section 3.4) are sufficiently powered while controlling for false positives.

Appendix E.1 Validation of model comparisons testing for Pavlovian bias (Section 3.4.1 of main text)

First, we verified that the model comparison analyses run in Section 3.4.1 possess adequate recovery properties in the sense of Palminteri et al. (2007): each behavioral model (base, Pavlovian, and RL) best fits its own simulated behavior.

To assess this, we simulated agents under Pavlovian influence in our experimental task, matching our lab experiment with 99 agents per treatment. The Pavlovian influence parameters were calibrated to reproduce the observed 5 percentage point excess betting rate between treatments (10% Test vs 5% Control).[28]

We then fitted these simulated participants' behavior using the two-step optimization procedure described in the main text. The Pavlovian model outperforms the base model (higher out-of-sample LL) for 87% of participants showing harmful betting. A one-tailed Wilcoxon test comparing out-of-sample LL confirms the Pavlovian model's superiority [$V = 52$, $p < 0.001$, effect size (rank-biserial) $= 0.9$].[29] The Pavlovian model also outperforms the RL model for all participants, demonstrating our procedure's ability to detect Pavlovian features when present.

[28] For each simulated participant, parameters were drawn from truncated normal distributions: θ (mean= 0.95, SD= 0.05, bounds= $[0.9, 1]$), κ_1 (mean= 0.18, SD= 0.05, bounds= $[0, 0.3]$), and κ_2 (mean= 0.8, SD= 0.05, bounds= $[0.7, 0.9]$). The logit parameter was set to 40 (same for all models) to match the low trembling observed in the lab participants.

[29] Standard paired t-tests are inappropriate for this comparison as goodness-of-fit statistics violate normality assumptions, even after Box-Cox transformations.

Next, we verified that the Pavlovian model is "restrictive" enough in the sense of Fudenberg et al. (2023). In a separate simulation batch with 99 agents per treatment devoid of any Pavlovian influence (behavior simulated using the base model), the Pavlovian model never outperforms the base model. Likewise, when simulating 99 RL agents per treatment, the RL model provides the best fit for all participants.[30]

Appendix E.2 Validation of tests characterizing the nature of the Pavlovian bias (Section 3.4.2 of main text)

The foregoing shows that each of the three models (base, Pavlovian, and RL) unequivocally fits its own behavior best. In the same spirit, we evaluated our protocol's ability to detect the uncertainty effect characteristic of sign-tracking while avoiding false positives:

- We simulated 99 Pavlovian agents per treatment, half programmed with the uncertainty effect from Eq. (7), half without it.[31] The sign-tracking version of the Pavlovian model better captures participant behavior than the original Pavlovian model according to a Bayesian Wilcoxon test (posterior probability: 0.9; Bayesian factor: 9). We used the Bayesian variant of the test because the standard test proves unreliable due to high kurtosis (> 10).

- Fitting these simulated participants' behavior to the generalized Pavlovian model [the version of the Pavlovian model that allows κ_1 to vary between high- and low-uncertainty trials, as described in Section 3.4.2, p.22] reveals significantly higher κ_1 values in high-uncertainty trials [one-tailed Wilcoxon test: $V = 4698$, $p < 0.001$, effect size (rank-biserial)=0.7]. We classified participants as displaying the uncertainty effect if their fitted κ_1 was $\geq 50\%$ larger in high-uncertainty trials. This criterion identifies 64 maladaptive bettors exhibiting the uncertainty effect among our 198 simulations, closely matching the true prevalence of 68 maladaptive bettors programmed with the effect. N.B.: For agents behaving rationally (no betting in yellow background sessions), the fitted κ_1 is consistently 0, as expected, indicating no Pavlovian or uncertainty effects.

- To further validate the analysis reported in Section 3.4.2, we repeated the analyses with 99 Pavlovian agents per treatment, all programmed *without* the uncertainty effect. Results show strong evidence *against* its presence (98% posterior probability; Bayesian factor: 52) and no κ_1 differences between uncertainty conditions, confirming that our protocol avoids false positives.

[30] Recovery rates remain robust to variations in learning rate eta (> 0.1) across simulation batches. In all RL simulations, initial Q values were randomly drawn from the outcome range $[-0.7, 1.3]$.

[31] For agents without the uncertainty effect, parameters followed the distributions specified in Footnote 28. For those with the uncertainty effect, distributions remained the same except C_{κ_1} replaced κ_1 (mean= 0.4, SD= 0.05, bounds= $[0.3, 0.4]$).

Appendix F Robustness checks

We conducted an extensive set of robustness checks. One important test was to augment the RL model with undirected exploration (UE), which allows updating of both selected and unselected choice values—a form of "novelty seeking" (Sutton, 2018):

$$U_t(c_t) = U_{t-1}(c_t) + \eta \cdot (x_t - U_{t-1}(c_t)),$$
$$U_t(u_t) = U_{t-1}(u_t) + e,$$

where c_t and u_t represent the selected and unselected choice in trial t, respectively, and $e \geq 0$ is an exploration parameter.

This extension aims to better capture observed betting rates (Table 3). However, as Table 8 shows, it fails to improve upon the original RL performance: while it increases betting in blue sessions as intended, it also raises betting in yellow sessions to implausible levels. The base model's superior fit persists at the participant level, even when compared against this augmented RL model with fitted exploration parameter.

We also tested model performance when allowing for loss tolerance in comparing base versus Pavlovian models. The Pavlovian model maintains its clear superiority (Bayesian Wilcoxon test: 98% posterior probability of superior fit; Bayesian factor: 61). Importantly, after incorporating loss tolerance, the calibrated Pavlovian model fully reproduces the observed betting rates (Table 8).

The main conclusions from our model comparisons are robust to several additional tests. They hold when analyzing the complete participant sample rather than just maladaptive bettors, and when relaxing the choice randomness constraint (allowing the logit parameter to be below 40, unlike in our main specification—see Footnote 28).

Table 8. Average betting rates by treatment, color, and agent type (extended). We simulated 99 agents per treatment for each model: rational choice (base), Pavlovian, and Reinforcement Learning (RL). For RL, we considered both Type 1 and Type 2 agents (as defined in Table 3), with Type 2 further augmented by undirected exploration (UE) to increase blue session betting rates. The base and Pavlovian models include both risk-neutral agents (RN: $\lambda = 1, \alpha = 1$) and loss-tolerant agents (LT: $\lambda < 1, \alpha = 1$). All agents generated trial-by-trial action values and made trial-specific decisions using choice probabilities computed from Eq. (2) (base and RL models) or Eq. (4) (Pavlovian model). Model parameters were calibrated to match the mean betting rates observed across lab participants.

	Session Type: (Color, Treatment)			
Agent-type	Blue, Control	Blue, Test	Yellow, Control	Yellow, Test
RL (Type 1)	0.93	0.97	0.92	0.97
RL (Type 2)	0.1	0.21	0.15	0.05
RL (Type 2) + UE	0.58	0.58	0.37	0.39
Base (RN)	0.97	0.97	0	0
Pavlovian (RN)	0.97	0.97	0	0.05
Base (LT)	0.99	0.99	0.05	0.05
Pavlovian (LT)	0.99	0.99	0.05	0.1
Lab Data	0.97	0.97	0.05	0.1

References

Alekseev, A., Charness, G. and Gneezy, U., 2017. Experimental methods: When and why contextual instructions are important. Journal of Economic Behavior & Organization, 134, pp.48-59.

Amsel, A., 1992. Frustration theory: An Analysis of Dispositional Learning and Memory (No. 11). Cambridge University Press.

Amsel, A. and Roussel, J., 1952. Motivational properties of frustration: I. Effect on a running response of the addition of frustration to the motivational complex. Journal of Experimental Psychology, 43(5), p.363.

Anselme, P. and Güntürkün, O., 2019. How foraging works: uncertainty magnifies food-seeking motivation. Behavioral and Brain Sciences, 42, p.e35.

Anselme, P., Robinson, M.J. and Berridge, K.C., 2013. Reward uncertainty enhances incentive salience attribution as sign-tracking. Behavioural brain research, 238, pp.53-61.

Barberis, N.C. and Jin, L.J., 2023. Model-free and model-based learning as joint drivers of investor behavior (No. w31081). National Bureau of Economic Research.

Benjamin, D.J., 2019. Errors in probabilistic reasoning and judgment biases. Handbook of Behavioral Economics: Applications and Foundations 1, 2, pp.69-186.

Berridge, K.C. and O'Doherty, J.P., 2014. From experienced utility to decision utility. In Neuroeconomics (pp. 335-351). Academic Press.

Berridge, K.C. and Robinson, T.E., 2003. Parsing reward. Trends in neurosciences, 26(9), pp.507-513.

Breland, K. and Breland, M., 1961. The misbehavior of organisms. American Psychologist, 16(11), p.681.

Brown, A.L., Imai, T., Vieider, F.M. and Camerer, C.F., 2024. Meta-analysis of empirical estimates of loss aversion. Journal of Economic Literature, 62(2), pp.485-516.

Brown, P.L. and Jenkins, H.M., 1968. AUTO-SHAPING OF THE PIGEON'S KEY-PECK 1. Journal of the experimental analysis of behavior, 11(1), pp.1-8.

Bushong, B., King, L.M., Camerer, C.F. and Rangel, A., 2010. Pavlovian processes in consumer choice: The physical presence of a good increases willingness-to-pay. American Economic Review, 100(4), pp.1556-1571.

Cavanagh, J.F., Eisenberg, I., Guitart-Masip, M., Huys, Q. and Frank, M.J., 2013. Frontal theta overrides pavlovian learning biases. Journal of Neuroscience, 33(19), pp.8541-8548.

Chapman, J., Snowberg, E., Wang, S.W. and Camerer, C., 2024. Looming large or seeming small? Attitudes towards losses in a representative sample. Review of Economic Studies, p.rdae093.

Charles, C., Frydman, C. and Kilic, M., 2024. Insensitive investors. The Journal of Finance, 79(4), pp.2473-2503.

Charness, G., Gneezy, U. and Halladay, B., 2016. Experimental methods: Pay one or pay all. Journal of Economic Behavior & Organization, 131, pp.141-150.

Chen, H., Cohen, P. and Chen, S., 2010. How big is a big odds ratio? Interpreting the magnitudes of odds ratios in epidemiological studies. Communications in Statistics—simulation and Computation®, 39(4), pp.860-864.

Cohen, J., 2013. Statistical Power Analysis for the Behavioral Sciences. Routledge.

Conlisk, J., 1993. The utility of gambling. Journal of Risk and Uncertainty, 6, pp.255-275.

Collins, L., Young, D.B., Davies, K. and Pearce, J.M., 1983. The influence of partial reinforcement on serial autoshaping with pigeons. The Quarterly Journal of Experimental Psychology, 35(4), pp.275-290.

Dal Bó, P. and Fréchette, G.R., 2011. The evolution of cooperation in infinitely repeated games: Experimental evidence. American Economic Review, 101(1), pp.411-429.

Daw, N.D., 2014. Advanced reinforcement learning. Neuroeconomics, pp.299-320.

Dayan, P. and Balleine, B.W., 2002. Reward, motivation, and reinforcement learning. Neuron, 36(2), pp.285-298.

Dayan, P., Niv, Y., Seymour, B. and Daw, N.D., 2006. The misbehavior of value and the discipline of the will. Neural networks, 19(8), pp.1153-1160.

Dean, M., Kıbrıs, Ö. and Masatlioglu, Y., 2017. Limited attention and status quo bias. Journal of Economic Theory, 169, pp.93-127.

Dorfman, H.M. and Gershman, S.J., 2019. Controllability governs the balance between Pavlovian and instrumental action selection. Nature communications, 10(1), p.5826.

Firth, D., 1993. Bias reduction of maximum likelihood estimates. Biometrika, 80(1), pp.27-38.

Fudenberg, D., Gao, W. and Liang, A., 2023. How flexible is that functional form? Quantifying the restrictiveness of theories. Review of Economics and Statistics, pp.1-50.

Gabaix, X., 2019. Behavioral inattention. In Handbook of Behavioral Economics: Applications and Foundations 1 (Vol. 2, pp. 261-343). North-Holland.

Gabaix, X. and Koijen, R.S., 2021. In search of the origins of financial fluctuations: The inelastic markets hypothesis (No. w28967). National Bureau of Economic Research.

Gershman, S.J., Guitart-Masip, M. and Cavanagh, J.F., 2021. Neural signatures of arbitration between Pavlovian and instrumental action selection. PLoS computational biology, 17(2), p.e1008553.

Giglio, S., Maggiori, M., Stroebel, J. and Utkus, S., 2021. Five facts about beliefs and portfolios. American Economic Review, 111(5), pp.1481-1522.

Hearst, E. and Jenkins, H.M., 1974. Sign-tracking: The stimulus-reinforcer relation and directed action. (No Title).

Hertwig, R., 2017. When to consider boosting: some rules for policy-makers. Behavioural Public Policy, 1(2), pp.143-161.

Ho, T.H., Camerer, C.F. and Chong, J.K., 2007. Self-tuning experience weighted attraction learning in games. Journal of Economic Theory, 133(1), pp.177-198.

Ho, T.H., Wang, X. and Camerer, C.F., 2008. Individual differences in EWA learning with partial payoff information. The Economic Journal, 118(525), pp.37-59.

Hoffrage, U., Lindsey, S., Hertwig, R. and Gigerenzer, G., 2000. Communicating statistical information. Science, 290(5500), pp.2261-2262.

Kendall, C. and Oprea, R., 2024. On the complexity of forming mental models. Quantitative Economics, 15(1), pp.175-211.

Lehne, M. and Koelsch, S., 2015. Toward a general psychological model of tension and suspense. Frontiers in psychology, 6, p.79.

Lo, A., 2017. Adaptive Markets: Financial Evolution at the Speed of Thought. Princeton University Press.

Lusardi, A. and Mitchell, O.S., 2007. Baby boomer retirement security: The roles of planning, financial literacy, and housing wealth. Journal of Monetary Economics, 54(1), pp.205-224.

Martínez-Marquina, A., Niederle, M. and Vespa, E., 2019. Failures in contingent reasoning: The role of uncertainty. American Economic Review, 109(10), pp.3437-3474.

McLean, R.D. and Pontiff, J., 2016. Does academic research destroy stock return predictability?. The Journal of Finance, 71(1), pp.5-32.

Nakagawa, S. and Schielzeth, H., 2013. A general and simple method for obtaining R2 from generalized linear mixed-effects models. Methods in ecology and evolution, 4(2), pp.133-142.

Ortmann, A., 2020. Deception. In: A. Schram and A. Ule, eds. Handbook of Research Methods and Applications in Experimental Economics. Cheltenham: Elgar, pp. 28-39.

Palminteri, S., Wyart, V. and Koechlin, E., 2017. The importance of falsification in computational cognitive modeling. Trends in cognitive sciences, 21(6), pp.425-433.

Payzan-LeNestour, E., 2018. Can people learn about 'black swans'? Experimental evidence. The Review of Financial Studies, 31(12), pp.4815-4862.

Payzan-LeNestour, E. and Bossaerts, P., 2015. Learning about unstable, publicly unobservable payoffs. The Review of Financial Studies, 28(7), pp.1874-1913.

Payzan-LeNestour, E. and Doran, J., 2024. Craving money? Evidence from the laboratory and the field. Science Advances, 10(2), p.eadi5034.

Payzan-LeNestour, E., Dunne, S., Bossaerts, P. and O'Doherty, J.P., 2013. The neural representation of unexpected uncertainty during value-based decision making. Neuron, 79(1), pp.191-201.

Redish, A.D., 2013. The Mind Within the Brain: How We Make Decisions and How Those Decisions Go Wrong. Oxford University Press.

Robinson, M.J., Anselme, P., Fischer, A.M. and Berridge, K.C., 2014. Initial uncertainty in Pavlovian reward prediction persistently elevates incentive salience and extends sign-tracking to normally unattractive cues. Behavioural brain research, 266, pp.119-130.

Russell, T.I. and Robinson, M.J., 2019. Effects of nicotine exposure and anxiety on motivation for reward and gambling-like cues under reward uncertainty. Behavioral neuroscience, 133(4), p.361.

Samuelson, W. and Zeckhauser, R., 1988. Status quo bias in decision making. Journal of Risk and Uncertainty, 1, pp.7-59.

Sanborn, A.N. and Chater, N., 2016. Bayesian brains without probabilities. Trends in cognitive sciences, 20(12), pp.883-893.

Schüll, N.D., 2012. Addiction by design: Machine gambling in Las Vegas. In Addiction by Design. Princeton University Press.

Shmaya, E. and Yariv, L., 2016. Experiments on decisions under uncertainty: A theoretical framework. American Economic Review, 106(7), pp.1775-1801.

Sutton, R.S., 2018. Reinforcement learning: An introduction. A Bradford Book.

Szollosi, A. and Newell, B.R., 2020. People as intuitive scientists: Reconsidering statistical explanations of decision making. Trends in Cognitive Sciences, 24(12), pp.1008-1018.

Szollosi, A., Wang-Ly, N. and Newell, B., 2023. Nudges for People who Think. Working Paper. PsyArXiv, June 8. Available at: https://doi.org/10.31234/osf.io/ravsp.

Winkielman, P. and Berridge, K., 2003. Irrational wanting and subrational liking: How rudimentary motivational and affective processes shape preferences and choices. Political Psychology, 24(4), pp.657-680.

Zhang, J., Berridge, K.C., Tindell, A.J., Smith, K.S. and Aldridge, J.W., 2009. A neural computational model of incentive salience. PLoS computational biology, 5(7), p.e1000437.

Zuckerman, M., 2007. Sensation seeking and risky behavior (Vol. 72). Washington, DC: American Psychological Association.