

SISSA

Scuola
Internazionale
Superiore di
Studi Avanzati

Neuroscience Area - PhD course in
Cognitive Neuroscience

Perceptual Decision Making of Nonequilibrium Fluctuations

Advisor:

Candidate:

Prof. Domenica Bueti

Ayşe Aybüke Durmaz

Co-advisor:

Prof. Mathew Ernest Diamond

Academic Year 2025-26



General Abstract

How does the brain recover a weak signal that is submerged in intense stochastic fluctuations to make fast yet accurate choices? We recast perceptual decision-making as the inference of a nonzero drift (v) in the presence of large diffusivity (D): the observer must determine the direction of motion when trajectories are dominated by diffusion. A concrete analogy is reading wind while hunting: turbulent gusts scramble moment-to-moment cues, yet a subtle, consistent drift in air motion carries actionable information about wind direction. In this perspective, the core problem is signal discovery under diffusion—extracting the sign and magnitude of a weak drift from time-resolved fluctuations.

This framing contrasts with the traditional random-dot motion (RDM) paradigm, where “coherence” (the percentage of dots moving consistently) is an effective but unitless control of difficulty. Although 50% coherence is intuitively “more signal” than 10%, its relationship to the quantity of evidence is ambiguous: is it five times more, or something else? Because coherence does not specify the physical statistics generating samples, an ideal-observer analysis (e.g., a Sequential Probability Ratio Test, SPRT) cannot be uniquely grounded in first principles. By instead specifying a physics-based stimulus with known parameters, this thesis makes the evidence quantified and the ideal-observer well-posed.

We generated visual motion stimuli from a drifted Brownian process. To index the distance from equilibrium, we employed an interpretable nonequilibrium measure (entropy production Σ) proportional to drift–diffusion contrast, which increases as directional drive overwhelms diffusivity. In this generative setting, the momentary log-likelihood ratio (LLR) for direction decisions and the optimal stopping boundaries of the SPRT are derived analytically from (v, D) , providing a physics-grounded benchmark for ideal performance.

Across three behavioral experiments, we asked: i) whether human observers ($N = 67$) detect and exploit graded nonequilibrium dynamics; ii) how closely their choices approach an ideal-observer benchmark; iii) how evidence integration adapts as Σ varies; and iv) whether such adaptation

depends on task structure and the spatiotemporal layout of the stimuli. Results showed that stimulus dynamics (Σ , v , D) robustly shaped decision metrics, demonstrating that observers are indeed sensitive to graded changes in Σ .

An analytical SPRT captured these effects and quantified deviations from ideal performance. Complementarily, an Evidence Integration Model (EIM) fitted to the data revealed a systematic adjustment of the temporal integration window with Σ : in each trial, observers assigned greater weight to the most recent portion of the trajectory (a recency effect) whose strength scaled with Σ . Observers were also sensitive to salient changes in trajectory directionality, consistent with adaptive weighting under nonstationary drift.

Crucially, these effects were stronger in a blocked design—where Σ was held constant within blocks—than in an intermixed design, where Σ varied from trial to trial, indicating that stable nonequilibrium statistics facilitate calibration of integration timescales. Finally, sensitivity to the nonequilibrium structure was modulated not only by the physical parameters (v , D) but also by the spatial and temporal layout of the stimuli. Overall, by embedding perceptual evidence in a physics-based process that specifies its quantity, this work refines the characterization of variables that govern perceptual decisions and clarifies the temporal dynamics underlying efficient sensory evidence integration. It shows that when evidence is measured—rather than merely manipulated—ideal-observer analyses become principled, enabling precise tests of how the brain detects weak directional signals under high diffusion.

Contents

1	General Introduction	1
2	Visual Perceptual Decisions Are Tuned to Thermodynamic Structure of Brownian Motion with Drift*	6
2.1	Abstract	6
2.2	Introduction	7
2.3	Results	9
2.4	Discussion	28
2.5	Materials and Methods	30
3	A Closer Look into Stimulus Dynamics Influencing Decision Under Nonequilibrium Fluctuations	44
3.1	Abstract	44
3.2	Introduction	45
3.3	Results	47
3.4	Discussion	56
3.5	Materials and Methods	59
4	Effects of Stimulus Format on Decision Making under Nonequilibrium Dynamics	66
4.1	Abstract	66
4.2	Introduction	67
4.3	Results	69
4.4	Discussion	87
4.5	Materials and Methods	89
5	General Discussion and Conclusion	101

Chapter 1

General Introduction

Perceptual decision-making requires the brain to extract reliable information from inherently noisy sensory signals in order to make fast yet accurate choices. Classic frameworks—including signal detection theory and sequential sampling models—typically treat sensory noise as a stationary perturbation imposed on an otherwise stable stimulus [1]–[3]. Within this view, the brain’s primary computational challenge is to accumulate enough evidence to overcome this noise. However, the assumption that sensory noise is fixed and featureless is increasingly recognized as inadequate [4]–[9]. In natural environments, sensory “noise” is not simply a uniform limit on performance but a high dimensional, time dependent structure. A noisy stimulus can itself evolve stochastically, providing the observer with not merely an average uncertainty level but a continuously fluctuating temporal signal.

A growing body of work shows that human observers are sensitive to this fine-grained temporal structure. Specific spatiotemporal realizations of a stimulus (not just its overall noise level) systematically influence perceptual choices. Park et al. demonstrated that incorporating the full stimulus realization on each trial allows a stimulus-informed model to account for substantial trial by trial variability in behavior beyond what average noise statistics explain [10]. Similarly, Fard et al. repeatedly presented identical 0% coherence random dot motion stimuli (thus, with no expected net directional signal) with participants assuming there were true underlying signals; the experiments found that certain noise instantiations elicited reproducible choice biases — consistent across repetitions and even across observers — revealing that subtle moment-to-moment fluctuations can systematically shape decisions [11]. These findings reframe noise not merely as a limiting factor, but as a meaningful component of the sensory input that observers implicitly track.

A complementary line of research highlights that natural sensory environments are rarely stationary:

the reliability, direction, and temporal statistics of the evidence often change within a single decision episode, and humans adapt accordingly. When evidence evolves continuously, observers can improve performance by adjusting the timescale of leaky integration to match stimulus dynamics [12]. Neural evidence further shows tight links between normative adaptive accumulation and recurrent cortical decision dynamics when the underlying stimulus state changes [7]. Together, these results emphasize that perceptual decisions depend not only on the amount of sensory evidence available, but also on how that evidence unfolds in time.

Taken together, these findings motivate a key challenge: to identify a physics-based descriptor of stimulus dynamics that captures how sensory evidence unfolds over time, is experimentally controllable, and remains sensitive to the fine-grained spatiotemporal realization of the input. The perspective we adopt here comes from thermodynamics — and, in particular, from the statistical physics of nonequilibrium fluctuations. In simple terms, a system is in equilibrium when its fluctuations are indistinguishable regardless of whether viewed forward or backward in time, such as random pixel flicker (“TV snow”) with no systematic drift. By contrast, nonequilibrium dynamics arise when there is some degree of directionality or “push” in the system’s evolution — much like a river flowing downstream or the molecules in a gas that is expanding when a valve has been opened. Such processes are temporally irreversible: if their evolution is played backward, it appears unnatural or physically implausible. Importantly, nonequilibrium is not a binary property but a graded one. Some processes are only weakly directional, whereas others are strongly driven and exhibit a pronounced arrow of time.

A central quantity characterizing this temporal irreversibility is entropy production, which provides a principled and interpretable measure of how far a process deviates from equilibrium — how strongly its fluctuations express temporal irreversibility [13], [14]. Low entropy production corresponds to nearly reversible dynamics; high entropy production signals strongly driven or directional, asymmetric fluctuations. Entropy production thus captures not only the magnitude of noise but also its temporal organization. Although pervasive in natural environments, this dimension of stimulus structure has received limited attention in perceptual neuroscience. Yet it offers something

uniquely powerful: a principled way to connect the statistics of the sensory evidence to the statistics of observers' choices. In particular, entropy production provides a quantitative measure of the temporal irreversibility present in the stimulus — an explicit description of how much directional, time-asymmetric structure is available to the observer at each moment of the trajectory. By manipulating entropy production, we can therefore parametrically control the informational content of the evidence stream in well-defined statistical-mechanical terms.

This creates a direct mapping between the external input and the computations performed by the perceptual system. Rather than describing the stimulus only through aggregate features such as coherence or signal-to-noise ratio — as in classical random-dot motion paradigms — we can characterize the actual temporal structure of the evidence that the observer receives. Entropy production thus allows us to move from broad notions of “noisy” versus “structured” stimuli to a formal, mechanistic account of how far the stimulus is driven from equilibrium and how much time-asymmetric information it contains. This, in turn, enables us to ask whether — and to what extent — human perceptual decisions track the nonequilibrium structure present in the sensory input.

Ideas rooted in thermodynamics have been fruitfully applied to biological sensing, cellular adaptation, and information processing in fluctuating environments [15]–[17], where many systems must expend energy to maintain robust performance. More directly relevant to perception, theoretical analyses have begun to explore how nonequilibrium structure constrains decision-making. Roldán and colleagues examined environments whose hidden states themselves evolve with nonequilibrium dynamics and showed that the amount of information an ideal observer can extract from an evidence stream is fundamentally bounded by the entropy production of the underlying process [13]. These results suggest that temporal irreversibility in sensory inputs may carry behaviorally relevant information.

Empirical evidence converges on the idea that humans are sensitive to temporal irreversibility: several studies have contrasted forward and time reversed versions of the same stimulus and found that reversing temporal order systematically alters perceptual and neural responses [18]–[20]. How-

ever, these studies rely on a binary manipulation — forward versus backward playback — applied to complex stimuli with unknown generative dynamics. As such, they cannot reveal how perception responds to graded levels of nonequilibrium structure or which specific dynamical features the perceptual system tracks.

Taken together, although natural sensory environments contain rich nonequilibrium fluctuations, and although humans are sensitive to temporal patterns in noisy inputs, no prior work has parametrically varied the entropy production properties of stimuli to test how perceptual decision-making depends on the level of nonequilibrium fluctuations. Doing so requires stimuli that are mathematically tractable, experimentally controllable, and capable of expressing different degrees of temporal irreversibility.

This dissertation addresses this gap by introducing visual motion stimuli generated from a well defined stochastic process — drifted Brownian motion governed by a Langevin equation — and parametrically manipulating their nonequilibrium fluctuations. Because the entropy production rate of this process can be computed analytically and controlled experimentally, it enables the creation of stimulus families that vary systematically in temporal irreversibility. Using these stimuli, we asked whether, and how, human perceptual decision-making depends on the level of nonequilibrium fluctuations in the motion stream. Across experiments, we examined how accuracy, reaction time, and inferred decision thresholds vary with nonequilibrium structure. Behavior was modeled using both normative and descriptive frameworks, including the Sequential Probability Ratio Test [13], [21] and an Evidence Integration Model incorporating memory timescales and leaky accumulation. We also analyzed which features of the stimulus trajectory exert the strongest influence under nonequilibrium conditions—specifically, how observers weight different moments of the trajectory and which moment-to-moment fluctuations bias their choices.

The dissertation is organized as follows. Chapter 2 presents Experiments A and B, which used drifted Brownian motion stimuli with varying nonequilibrium characteristics and compared human behavior with predictions of the Sequential Probability Ratio Test and the Evidence Integration Model. Chapter 3 provides complementary analyses, examining participants’ ability to track differ-

ent temporal moments and exploit salient changes in the stimulus trajectory. Chapter 4 presents Experiment C, in which the same stochastic dynamics were embedded in three distinct spatiotemporal layouts. We again evaluated behavior using both model based and data driven approaches. Finally, Chapter 5 concludes with a discussion of limitations, theoretical implications, and future directions for understanding perceptual decision-making under nonequilibrium fluctuations.

Chapter 2

Visual Perceptual Decisions Are Tuned to Thermodynamic Structure of Brownian Motion with Drift*

Aybüke Durmaz¹, Yonathan Sarmiento^{1,2}, Gianfranco Fortunato¹, Debraj Das², Mathew E. Diamond¹,
Domenica Bueti¹, Edgar Roldán²

1 International School for Advanced Studies (SISSA), Trieste, Italy

2 The Abdus Salam International Center for Theoretical Physics (ICTP), Trieste, Italy

2.1 Abstract

How are rapid and reliable perceptual judgments made when sensory inputs are noisy? Neural recordings from the lateral intraparietal area in non-human primates performing perceptual decision-making tasks have revealed mechanisms of evidence accumulation, often modeled as drift-diffusion processes. In many such experiments, participants judge the net direction of random dot motion stimuli; this form of stimulus, while computationally tractable, is not tied to any physical phenomenon. To better characterize the statistical processes underlying human decision-making in response to realistic physical dynamics, we asked participants to assess fluctuations of physical nonequilibrium stationary states and analyzed their responses within the framework of stochastic thermodynamics. Forty-five participants viewed short movies of a particle undergoing drifted Brownian motion and judged its direction of drift. Performance exhibited limits consis-

*Parts of the analyses in this chapter are also available in Yonathan Sarmiento's thesis.

tent with thermodynamic uncertainty relations linking speed, accuracy, and energetic dissipation: lower entropy production led to longer decision times. Participants also required more time than predicted by Wald’s optimal Sequential Probability Ratio Test, indicating suboptimal evidence integration. To capture this behavior quantitatively, we developed a model with non-Markovian evidence accumulation and a finite memory time constant. Variation in this time constant with entropy production suggests that humans adapt their integration window to the rate of dissipation, favoring temporally extended memory over instantaneous evidence to support effective decisions in nonequilibrium contexts. These findings identify, for decision-making based on stochastic evidence accumulation, both the nature of the evidence and the algorithm by which it is integrated.

2.2 Introduction

The brain is endowed with mechanisms for processing noisy sensory data to make fast and accurate decisions [22]. For a binary decision between two competing hypotheses, H_1 and H_2 , Wald’s Sequential Probability Ratio Test (SPRT) [23], rooted in probability theory [24], provides a universal model of optimal decision-making. SPRT specifies how information should be integrated over time to maximize decision accuracy; its applications include statistics [21], biophysics [25]–[27], and neuroscience [1]–[3]. In the SPRT framework, ongoing evidence is accumulated by computing the log-likelihood ratio between H_1 and H_2 until the weight of evidence for one hypothesis exceeds a predefined threshold. In neuroscience, the SPRT has been implemented through Drift Diffusion Models (DDM) [2] where a decision is made once the accumulated log-likelihood ratio—subject to instantaneous fluctuations but exhibiting a drift over longer timescales—reaches one of two thresholds, or boundaries. Combining the SPRT and DDM paradigms offers a solution wherein the first boundary reached (H_1 or H_2), determines the decision, and the time to reach it corresponds to the decision time [28].

A range of studies in rhesus monkeys [29]–[31], rats [32], and humans [2] have shown decision-making performance close to that predicted by SPRT, lending support to sequential hypothesis-testing models [22], [33]. One widely used paradigm is the random dot motion (RDM) task, in which

subjects judge the net motion within a field of moving dots. In these studies, rhesus monkeys are trained to make a saccade toward the screen side corresponding to the perceived net motion (e.g., rightward saccade for rightward motion) [30]. Across trials, the coherence of motion—i.e., the proportion of dots moving coherently in the target direction—is manipulated to vary task difficulty. Neural recordings show that the activity of spatially selective neurons in the lateral intraparietal area (LIP) correlates with the monkeys’ decisions. Specifically, activity in neurons coding for the chosen location gradually increases and peaks around saccade onset. Moreover, the slope of this rise scales with motion coherence: the higher the coherence, the steeper the increase [22], [29]–[31], [33]. These findings reveal neuronal operations consistent with the accumulation of evidence toward a decision threshold.

Recent research highlights how complex spatiotemporal fluctuations in stimuli influence decision outcomes [10], [11], [33]–[39]. In addition to average noise levels, participants consistently use detailed spatiotemporal noise patterns and stochastic input trajectories within stimuli [11], [34], [40]. Yet, despite the knowledge gained from studies using RDM tasks, the statistical structure of the evidence is not well understood. Beyond the experimenter-defined parameter of motion coherence, the thermodynamic properties of the sensory evidence accessible to the observer, such as the presence of time-irreversible fluctuations or energy dissipation, are difficult to characterize.

The nature of evidence accumulation cannot be understood independently of the properties of the evidence itself. This gap in understanding motivates a shift toward modeling sensory evidence as a structured, physically grounded, nonequilibrium process. To investigate how human observers acquire and evaluate information in such systems, one must consider stimuli whose fluctuations are not purely random but instead carry signatures of nonequilibrium dynamics—fluctuations that break detailed balance, generate net probability currents, and entail energy dissipation [17].

In this work, we conducted visual perception experiments (Experiments A, B, and a control experiment) in which participants judged movies of a driven Brownian particle described by a Langevin equation with a few nonequilibrium degrees of freedom. Our analyses, grounded in the framework of stochastic thermodynamics, establish a robust link between participants’ behavioral responses

and the input parameters of the underlying Brownian process. Experiment A tests a theoretical model rooted in SPRT[13] against behavioral data, quantifying the extent to which human decisions deviate from the model’s definition of optimality. Experiment B examines how non-stationary environmental properties affect optimality, while the control experiment assesses variance due to non-decisional factors. Our approach integrates behavioral psychophysics, computational decision-making, and stochastic thermodynamics to refine decision-making models beyond the classical SPRT and to uncover fundamental thermodynamic constraints [13] that shape judgments of nonequilibrium stationary stimuli.

2.3 Results

We designed a decision-making task that mimics the dynamics of a real physical system: a fluctuating particle undergoing drifted Brownian motion. On each trial, the particle drifts either left or right while subject to stochastic fluctuations, and participants must infer its drift direction based on noisy observations (see Fig. 2.1A-B, and Supplementary Fig. S1). Unlike conventional random dot motion tasks, which use abstract visual noise, our approach embeds evidence accumulation in a thermodynamically-anchored process, where stimulus statistics reflect physical quantities such as entropy production and Péclet number (See Section 2.5.4.2).

To implement this, we constructed stimulus dynamics using principles from stochastic thermodynamics, systematically manipulating drift and diffusion and thereby controlling entropy production. This allowed us to probe how nonequilibrium fluctuations influence evidence integration. We employed two complementary models: i) a normative benchmark based on the Sequential Probability Ratio Test (SPRT), which defines optimal decision-making under the given stimulus statistics; and ii) an Evidence Integration Model (EIM) incorporating a memory time constant to capture the suboptimal decision making observed in human participants. Together, these models allowed us to quantify deviations from optimality and infer how decision strategies adapt to changes in the structure of the environment.

Below, we present behavioral results for each experiment, compare them to normative predictions,

and show how the EIM provides a robust account for how participants make judgments.

2.3.1 Experiment A

This experiment tested four conditions in which the entropy production rate, given by $\Sigma = v^2/D$, increased monotonically from conditions I to IV while the drift-to-diffusion ratio, v/D , was held constant (parameter values are provided in Supplementary Table S1). Trials from the same condition were grouped into blocks, and the order of these blocks was randomized for each participant.

Figs. 2.1C–F show data from a representative participant across the four conditions. In each condition, the middle panel (ordinate X_{dec}) displays the horizontal disk position at the decision time, T_{dec} . Black and red points indicate correct and incorrect decisions, respectively. The corresponding histograms of T_{dec} (top and bottom) and X_{dec} (right) are shown separately for correct and incorrect trials. Decision-time distributions were similar for correct and incorrect responses, and the values of X_{dec} for incorrect trials appear approximately reflected across the midline. These patterns suggest that errors primarily arose from the accumulation of “misleading” evidence—an outcome expected under stochastic stimulus dynamics—rather than from the use of an inappropriate decision strategy.

To evaluate the effects of entropy production rate Σ on decision-making at the group level, we analyzed the mean decision threshold (Fig. 2.2A), mean accuracy (Fig. 2.2B), and mean decision time (Fig. 2.2C), averaging across all trials and participants for the four experimental conditions. These results are shown in the colored plots in Fig. 2.2. Effects of Σ were assessed using linear mixed-effects (LME) models (see Section 2.5.4.1). For participants’ mean decision thresholds, $\langle X_{\text{dec}} \rangle$, a Type III ANOVA on the model estimates revealed no significant main effect of entropy production rate ($F_{3,60} = 1.17$, $p = 0.33$; see Fig. 2.2A). Thus, participants did not systematically adjust their mean decision thresholds across the entropy production conditions. Full ANOVA details are provided in Supplementary Table S2.

Next, we examined participants’ mean accuracy across the different entropy production rate conditions. A Type III ANOVA on the model estimates revealed a statistically significant effect of the

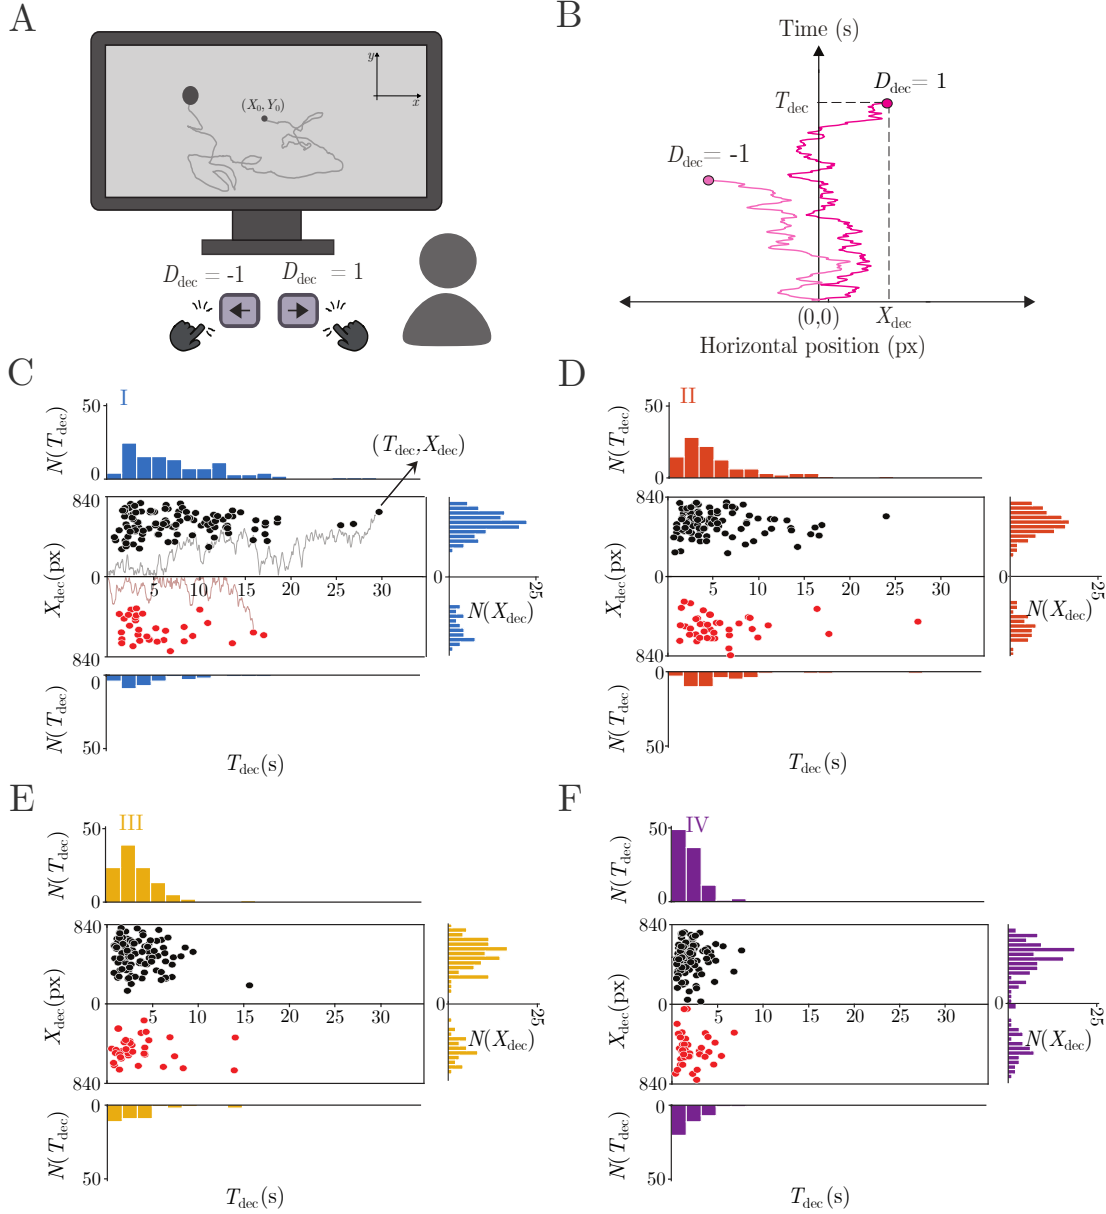


Figure 2.1: **Sketch of the experimental setup and data collection.** (A) Participants judged the net motion direction (left vs. right) of a moving disk displayed on a computer screen by pressing the corresponding arrow key on a keyboard. The disk's motion snapshots were generated from the stochastic simulations described in the main text. (B) Example traces of the disk's x -position over time are shown in fuchsia. For each trial, we recorded the decision time (T_{dec} , in seconds), the disk's horizontal and vertical positions at the time of decision (X_{dec} and Y_{dec} , in pixels; Y_{dec} not illustrated), and the decision outcome ($D_{\text{dec}} = 1$ for right, $D_{\text{dec}} = -1$ for left). We also stored the full trajectory of the disk up to T_{dec} (gray line in panel A). (C–F) Experimental data from one participant across the four conditions of Experiment A, each defined by a selected pair of (v, D) . In each condition, the central panel shows X_{dec} as a function of T_{dec} , with black and red points indicating correct and incorrect decisions, respectively. Histograms of T_{dec} (top: correct; bottom: incorrect) and of X_{dec} (right: correct; left: incorrect) are shown for each condition. In panels C–F, we pooled trials with positive and negative drift for both correct and incorrect decisions.

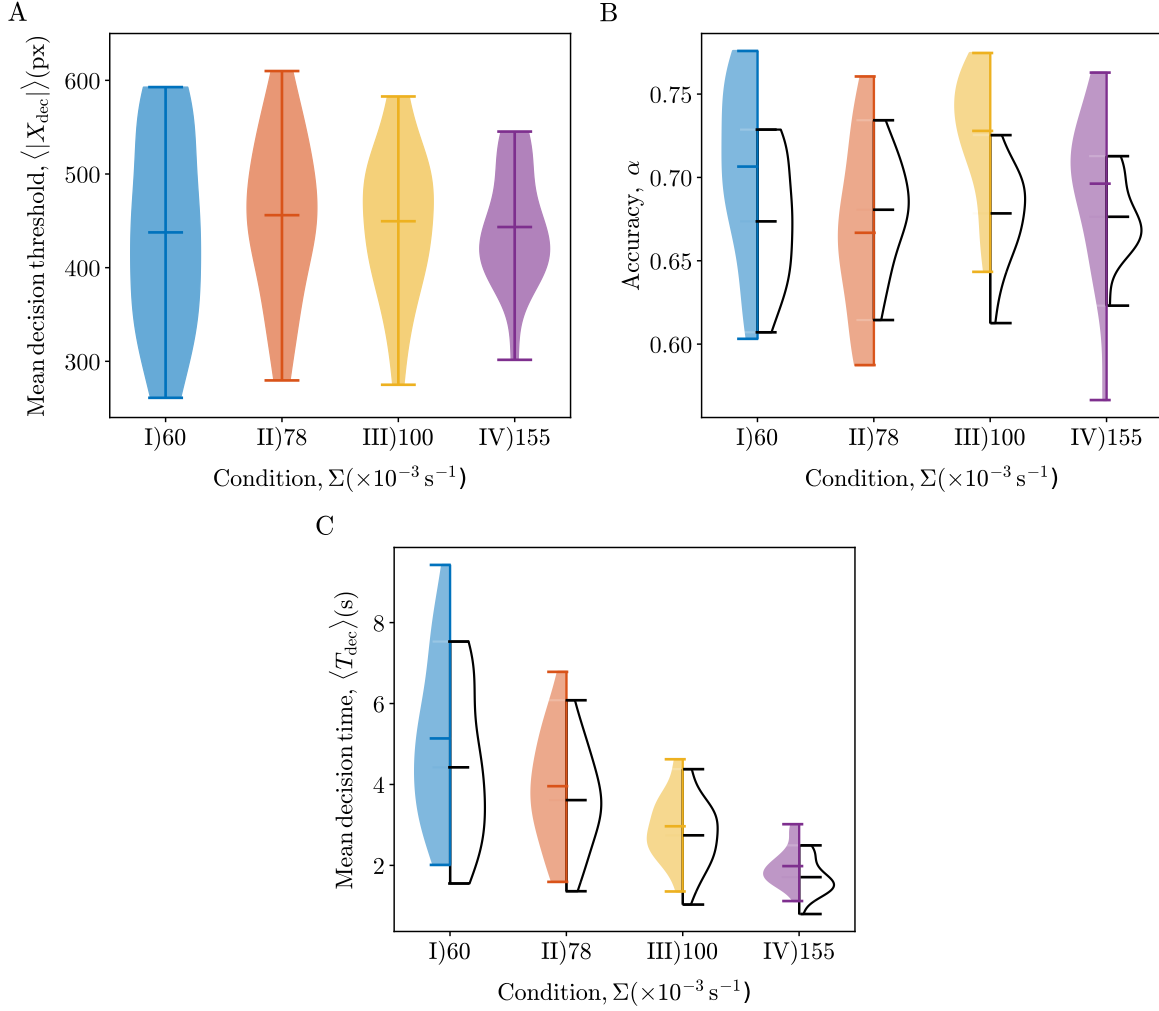


Figure 2.2: **Entropy main effects in Experiment A.** Participants' mean decision-making measures are shown as a function of the entropy production rate Σ . Panel (A) displays the mean decision threshold $\langle |X_{\text{dec}}| \rangle$, panel (B) shows the accuracy $P(D_{\text{dec}} = H)$, and panel (C) presents the mean decision time $\langle T_{\text{dec}} \rangle$. Colored violin plots represent the distributions across the $N = 21$ participants, each value obtained by averaging over all trials for that participant. In panels (B) and (C), the empty violin plots show the corresponding SPRT predictions, computed using thresholds set at $\pm \langle |X_{\text{dec}}| \rangle$ for each participant. Central horizontal lines indicate the means. SPRT predictions of accuracy and mean decision time are obtained by inserting the Péclet number and entropy production rate into Eqs. (2.9) and (2.10), respectively.

entropy production rate on accuracy ($F_{3,60} = 25.94$, $p < 0.001$; see Section 2.5.4.1, and Fig. 2.2B). Post-hoc multiple comparisons showed that all pairwise contrasts were significantly different (all $|t_{60}| > 2.83$, all $p < 0.038$), with the exception of the comparison between conditions I and IV. However, despite these significant differences, accuracy did not exhibit a clear monotonic trend as a function of entropy (see Supplementary Tables S3 and S4 for details of the ANOVA results and the Bonferroni-corrected multiple comparisons).

Finally, we consider participants' mean decision time $\langle T_{\text{dec}} \rangle$ across entropy production rate conditions. A Type III ANOVA on model estimates revealed a main effect of entropy production rate on T_{dec} ($F_{3,60} = 170.26$, $p < 0.001$; see Section 2.5.4.1, and Fig. 2.2C). Post-hoc multiple comparisons revealed that participants responded faster as entropy production rate increased from condition I to IV (all $t_{60} > 3.89$, $p < 0.0015$). Full ANOVA results and the Bonferroni-corrected multiple comparisons are reported in Supplementary Tables S5 and S6. The observed contraction in decision times across conditions aligns with prior theoretical work [13], which predicted that it would take less time to infer the direction of time's arrow when the underlying process is far from equilibrium, as evidence accumulates more rapidly in such cases.

Having established the effect of entropy production rate on participants' behavior, we next compared observed decision-making to the Sequential Probability Ratio Test (SPRT) [21], [23], which describes how an ideal observer accumulates evidence to discriminate between two competing hypotheses.

To calculate the SPRT predictions of accuracy (Eq. (2.10)) and decision times (Eq. 2.9), we considered symmetric thresholds located at $\pm L$, with $L = \langle |X_{\text{dec}}| \rangle$ set at the observed ensemble average decision thresholds; this follows from the fact that the Péclet number is defined as $\text{Pe} = vL/D$. To test the robustness of the decision thresholds in predicting the decision times, we performed a cross-validation analysis in Experiment A for each participant in each condition. We estimated the mean threshold from a training set and used it to predict mean T_{dec} in the remaining test set. We found that the experimental mean decision time of the training data accurately matches the predicted mean decision time of the test data (see Section 2.5.4.5; and Supplementary Fig. S2A).

This suggests that knowledge of decision thresholds of a particular data set may be used to predict decision times of unseen data.

The predictions of SPRT are shown as the black-outline plots in Figs. 2.2B-C, opposite the corresponding colored plots of empirical data (see also Supplementary Fig. S3A and C for the observed accuracy and reaction time data of each participant in each condition against the theoretical prediction of SPRT).

First, to assess how closely the SPRT predictions matched the observed accuracies across entropy production rate, we began by measuring the logit absolute error (AE) between the observed and the predicted accuracies using a LME (see Section 2.5.4.6). Type III ANOVA on model estimates revealed a significant main effect of entropy production rate ($F_{3,63} = 10.30$, $p < 0.001$; Fig. 2.2B). Post-hoc multiple comparisons revealed that condition III was significantly different from II and IV (all $|t_{66.15}| > 4.23$, all $p < 0.0004$). The rest of the comparisons were not significant (all $|t_{66.15}| < 2.71$, all $p > 0.05$). This result shows that the discrepancy between the observed accuracy and SPRT predictions were significant, suggesting suboptimality in participants' performance. See Supplementary Tables S7, and S8 for the details on ANOVA results and Bonferroni-corrected multiple comparisons.

Second, to evaluate how well the SPRT predicted the observed decision times across entropy production rates (Fig. 2.2C), we computed the relative absolute error between observed and predicted decision times using a LME (see Section 2.5.4.6). A Type III ANOVA on the model estimates revealed a significant main effect of entropy production rate ($F_{3,63} = 18.27$, $p < 0.001$). However, the discrepancies between observed and predicted decision times did not show a systematic trend with entropy: condition I did not differ from condition IV, and condition II did not differ from condition III. All other pairwise comparisons were significant (all $|t_{66.15}| > 3.97$, all $p < 0.001$). See Supplementary Tables S9, and S10 for the details on ANOVA results and Bonferroni-corrected multiple comparisons.

These results indicate that participants performed suboptimally relative to an ideal observer implementing the SPRT, with no consistent pattern of deviations across entropy conditions.

In summary, Experiment A shows that decision performance varied with the entropy production rate of the stimuli, leading to faster choices at higher entropy levels. In contrast, mean decision thresholds and accuracy did not differ across experimental conditions. Relative to SPRT predictions, participants exhibited longer decision times than those of an optimal decision maker. In the next section, we investigate the sources of this suboptimality.

2.3.2 Experiment B

The first aim of Experiment B was to examine how uncertainty in stimulus dynamics affects participants' choices. To achieve this, Experiment B replaced the blocked structure of Experiment A—where participants could anticipate the upcoming trial's stimulus statistics and prepare corresponding decision strategies—with an intermixed design in which trials from different conditions were randomly shuffled. We anticipated that this increased uncertainty about the upcoming values of v and D would require participants to adjust their strategy within each trial, leading to slower responses and reduced accuracy [41]. The second goal was to separate the individual contributions of v and D . Experiment B used a 3×3 grid of parameter combinations, allowing us to test the effects of each parameter independently (full parameter grid values are provided in Supplementary Fig. S4). This was not possible in Experiment A, where v and D were yoked so that their ratio (v/D) remained fixed.

2.3.2.1 Effect of predictability of stimuli statistics on optimality

To assess whether the entropy production rate influences participants' choice metrics differently in Experiments A and B, we again applied the LME framework to decision thresholds, accuracy, and decision times (see Section 2.5.4.7). First, we focused on decision thresholds. A Type III ANOVA on the model estimates revealed a significant main effect of experiment (A vs. B; $F_{1,43.073} = 4.35$, $p = 0.043$). Decision thresholds were substantially larger in Experiment B — by approximately 88 pixels — relative to Experiment A ($t_{43} = 3.97$, $p = 0.0003$; see Fig. 2.3A). Neither entropy nor the interaction between entropy and experiment reached significance. Full ANOVA results and Bonferroni-corrected multiple comparison tests are reported in Supplementary Tables S11,

and S12.

We next examined participants' mean accuracy. A Type III ANOVA on the model estimates revealed no significant main effect of experiment (A vs. B), no effect of entropy production rate, and no interaction between the two factors. These results indicate that participants' accuracy did not differ between the experiments (see Fig. 2.3B). Full ANOVA results are reported in Supplementary Tables S13.

Finally, we analyzed decision times. A Type III ANOVA revealed a significant main effect of entropy production rate in Experiment B ($F_{1,45.103} = 681.79, p < 0.001$), reflecting shorter decision times at higher entropy levels, consistent with the pattern observed in Experiment A. The main effect of experiment (A vs. B) was also significant ($F_{1,43.287} = 12.64, p = 0.0009$), whereas the interaction was not. Full ANOVA results are provided in Supplementary Table S14. Post-hoc comparisons showed that participants responded significantly more slowly in Experiment B, with decision times approximately 1.42 seconds longer on average ($t_{43} = 5.628, p < 0.0001$; Fig. 2.3C). Bonferroni-corrected multiple comparison tests are reported in Supplementary Table S15.

Overall, these results indicate that participants in Experiment B used wider decision thresholds and took longer to reach a decision than those in Experiment A. Yet despite accumulating more evidence, their accuracy did not increase significantly, pointing to a reduction in optimality (see also Supplementary Fig. S3B and D for the observed accuracy and reaction time data of each participant in each condition against the theoretical prediction of SPRT). Because the physical stimuli were identical in the two experiments, we attribute this decrease in optimality to the intermixed design.

2.3.2.2 Exploring the parameter space

To explore the parameter space more fully, Experiment B sampled a 3×3 grid of v and D combinations, unlike Experiment A where these parameters were yoked through a fixed v/D ratio and therefore could not be dissociated. We first examined how v and D affected the mean decision thresholds (Fig 2.4A; Section 2.5.4.1). The diagonal conditions in Experiment B, which share the

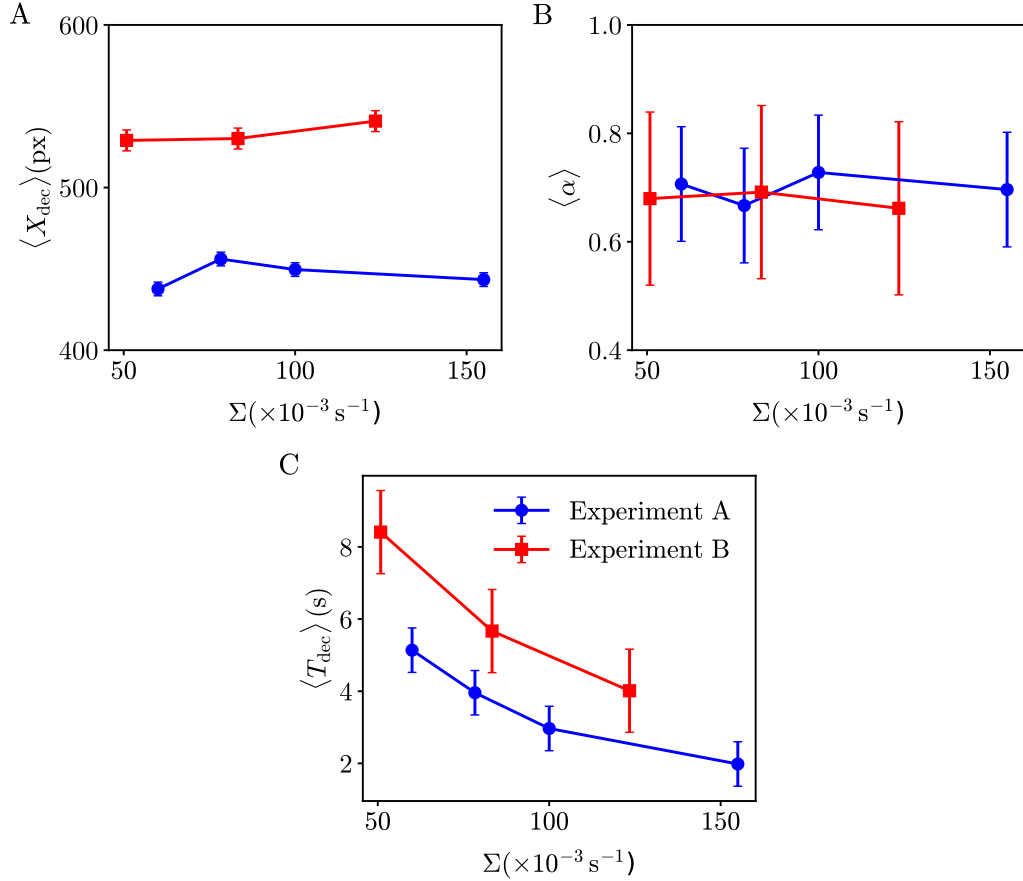


Figure 2.3: **Decision-making metrics as a function of entropy in Experiments A and B.** Decision-making metrics averaged across all participants, are shown as a function of the entropy production rate Σ for Experiment A (blue) and Experiment B (red). The conditions from the two experiments are comparable because they share similar values of Σ and a fixed Pe . (A) Mean decision thresholds $\langle |X_{\text{dec}}| \rangle$ were consistently larger in Experiment B than in Experiment A across all values of Σ , indicating that participants accumulated more evidence in the intermixed design. (B) Mean accuracy as a function of Σ . Statistical analyses revealed no significant differences in mean accuracy between the two experiments. (C) Mean decision times $\langle T_{\text{dec}} \rangle$ as a function of Σ . Both experiments exhibited the inverse relationship between decision time and Σ predicted by the SPRT, although overall decision times were slower in Experiment B. Error bars indicate the standard error of the mean (SEM). The corresponding (v, D) parameter values can be found in Supplementary S1 for Experiment A and in S4 (diagonal entries) for Experiment B.

same velocity-to-diffusivity ratio, correspond to the conditions used in Experiment A. A Type III ANOVA revealed significant main effects of both v ($F_{2,184} = 14.79$, $p < 0.001$) and D ($F_{2,184} = 5.14$, $p < 0.01$), but no interaction. Post-hoc multiple comparisons showed that thresholds in low- v conditions were significantly narrower than in both medium- v ($t_{184} = -4.38$, $p = 0.0001$) and high v ($t_{184} = -4.98$, $p < 0.0001$) conditions. For D , thresholds in low-diffusivity conditions were significantly wider than those in medium- D conditions ($t_{184} = 3.21$, $p = 0.0048$). Full ANOVA results and Bonferroni-corrected multiple comparison tests are reported in Supplementary Tables S16, S17, and S18.

Overall, these findings indicate that decision thresholds did not vary in a simple monotonic way with either velocity or diffusivity. Threshold setting therefore cannot be described as a straightforward function of v or D , suggesting that participants' threshold strategies are shaped by more complex influences than either stimulus parameter alone.

Second, we examined how v and D influenced accuracy (Fig. 2.4B; Section 2.5.4.1). A Type III ANOVA revealed significant main effects of v ($F_{2,184} = 564.85$, $p < 0.001$) and D ($F_{2,184} = 707.54$, $p < 0.001$), as well as a significant interaction between them ($F_{4,184} = 40.98$, $p < 0.001$). Higher v values (all $t_{184} > 3.190$, $p < 0.005$ except one pair) and lower D values (all $t_{184} > 3.644$, $p < 0.001$ except one pair) were associated with higher accuracy. The significant $v \times D$ interaction matches the prediction of the SPRT in Eq. 2.10. Full ANOVA results and Bonferroni-corrected multiple comparison tests are reported in Supplementary Tables S19, S20, and S21.

In our next analysis, we examined how the parameters v and D influenced decision times (Fig. 2.4C; Section 2.5.4.1). Both variables had strong, statistically significant effects, and they also interacted with one another. Type III ANOVA revealed main effects of v ($F_{2,184} = 22.10$, $p < 0.001$), and D ($F_{2,184} = 1551.72$, $p < 0.001$), and their interaction ($F_{4,184} = 33.89$, $p < 0.001$). Post-hoc comparisons showed that for any given value of v , every pair of D conditions differed significantly (all $|t_{184}| > 7.874$, $p < 0.001$). In all cases, higher D produced shorter decision times. This indicates that D was the primary factor shaping how long decisions took, since greater diffusivity causes the underlying evidence signal to fluctuate more rapidly. In contrast, changes in v within the

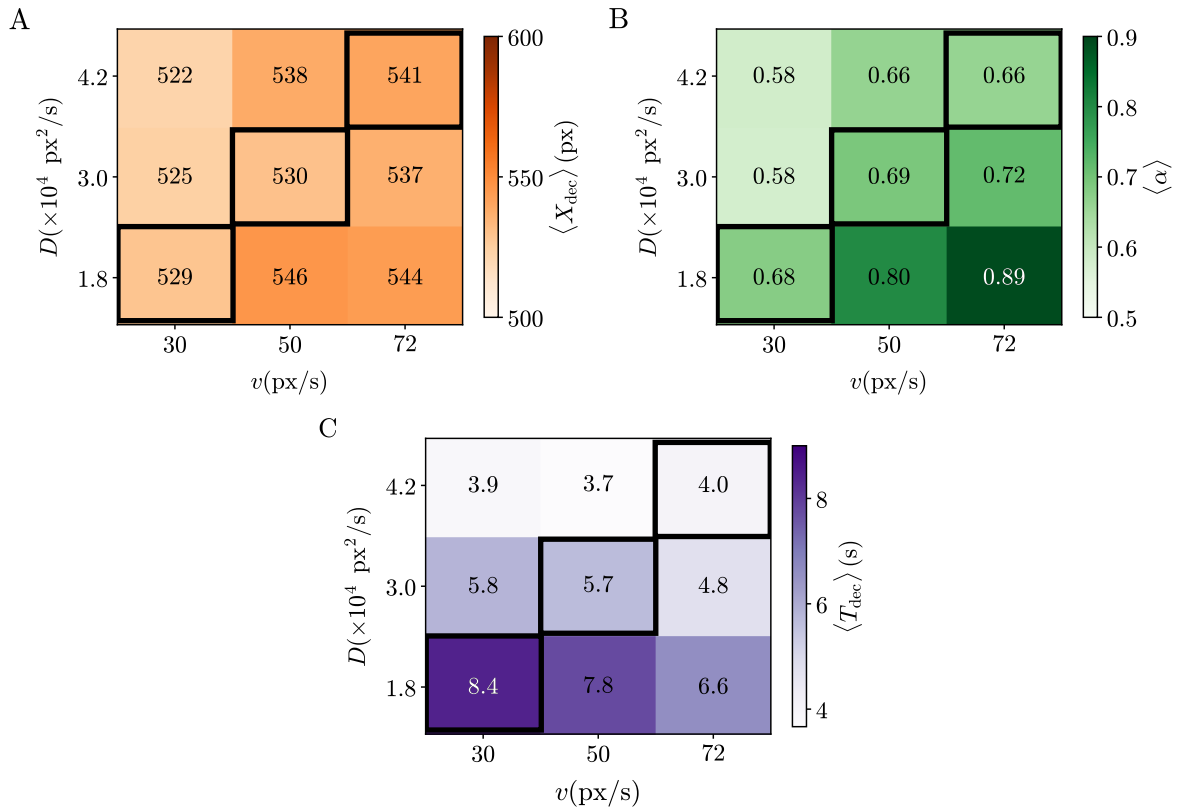


Figure 2.4: **Effects of v and D in Experiment B.** Each panel shows a 3×3 grid spanning the parameter space, with v on the x -axis and D on the y -axis, and each cell representing a specific combination of the two. (A) Mean decision thresholds, $\langle X_{\text{dec}} \rangle$. (B) Mean accuracy, $\langle \alpha \rangle$. (C) Mean decision times, $\langle T_{\text{dec}} \rangle$. Color bars indicate the numerical values associated with each decision-making measure.

tested range had a much smaller impact. Full ANOVA results and Bonferroni-corrected multiple comparison tests are reported in Supplementary Tables S22, S23, and S24.

Considering all behavioral measures together, separating v and D made their distinct roles clear. The decision threshold showed only weak and inconsistent dependence on either parameter. Accuracy, however, varied systematically with both v and D , following the expected ratio between them and displaying a strong interaction. Decision times were influenced most heavily by diffusivity, consistent with the idea that higher D leads to faster evidence dynamics and thus quicker decisions.

2.3.3 Evidence integration model to incorporate memory limits

Although the SPRT offers a strong normative benchmark for perceptual decision-making, our experimental data showed clear departures from its predictions. In particular, participants' decision times were consistently longer than the SPRT would predict for the same accuracy levels, indicating that human observers rely on strategies that are systematically suboptimal relative to the ideal model.

To understand these discrepancies, we examined where the assumptions of the SPRT diverged from actual participant behavior. Two main issues emerged. First, the SPRT assumes that decision thresholds are fixed, yet in our experiments the inferred thresholds varied widely across trials and conditions. The threshold histograms (Fig. 2.1C–F) show broad distributions even within individual participants. This variability cannot be explained solely by stochastic motor-response noise. A control experiment in which the threshold was visually displayed on-screen produced noticeably lower variability (see Section 2.5.3, and Supplementary Figs. S5, S6), ruling out the possibility that reaction-time fluctuations alone account for the observed spread. Instead, it appears that participants adjust their decision criteria dynamically, likely tuning them to the physical properties of each trial—such as entropy, velocity, and diffusivity—as we observed in Experiment B. An additional analysis supports this idea. Since all participants viewed the same set of stimulus trajectories but in a randomized order; when we aligned trials by stimulus identity, the inferred

decision thresholds showed reliable trial-by-trial consistency across participants (Supplementary Fig. S7; Section 2.5.4.9). These findings suggest that a meaningful portion of threshold variance was driven by properties of the specific trajectory realization on that trial.

The second issue is that SPRT assumes perfect, lossless integration of the entire stimulus history. Under this assumption, the accumulated evidence always reflects the full trajectory of the stimulus. Biological systems, however, are subject to memory constraints and leaky integration, especially when processing time-varying and noisy input. These limitations provide a further source of deviation from the idealized SPRT predictions.

These two limitations highlighted the need for a more flexible model—one capable of capturing both dynamic threshold adjustments and imperfect memory. To meet this need, we developed the Evidence Integration Model (EIM).

The EIM preserves the core idea of sequential evidence accumulation but incorporates a biologically plausible constraint: leaky integration. Instead of assuming perfect retention of all past sensory input, the model posits that accumulated evidence decays exponentially over time. Formally, we define the accumulated evidence as:

$$W_t \equiv \frac{1}{\tau_m} \int_0^t ds e^{-s/\tau_m} X_{t-s}, \quad (2.12)$$

where X_t denotes the momentary sensory evidence (e.g., the horizontal position of the moving disk), τ_m is the temporal integration window — a free parameter representing the timescale over which evidence is retained — and W_t is the exponentially weighted accumulation of past evidence up to time t . This formulation assigns greater weight to more recent observations, consistent with the idea that human observers may discount older information when processing dynamic, time-varying stimuli. Fig. 2.5A illustrates trajectories of the accumulator W_t generated from a single instance of the stochastic stimulus X_t for different values of τ_m .

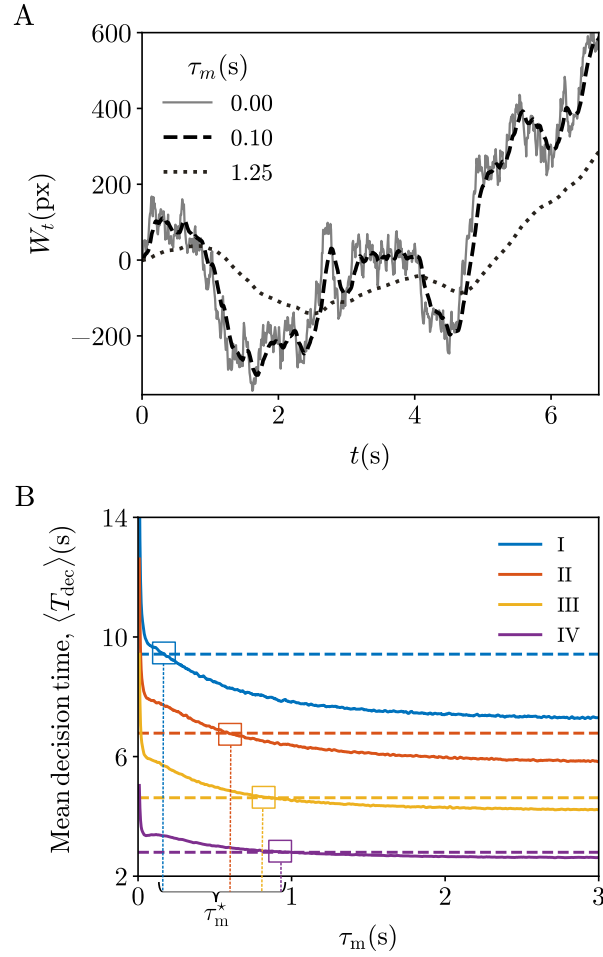


Figure 2.5: **Evidence Integration Model (EIM)**. (A) Sample trajectories of the decision accumulator W_t corresponding to a given position trajectory X_t for different integration windows τ_m . When $\tau_m = 0$, the accumulator simply follows the stimulus trajectory, so $W_t = X_t$ (gray solid line). (B) Estimation of the optimal integration window τ_m^* for one participant across the four conditions. Horizontal dashed lines show the participant's mean experimental decision time, while the solid curves show the mean escape time of W_t from the interval $(-L_W, L_W)$ obtained from numerical simulations as a function of τ_m . Here $L_W = \langle |W_{T_{\text{dec}}}| \rangle$ is computed from the experimental trajectories $X_{[0, T_{\text{dec}}]}$ (see Fig. 2.1). For each condition and participant, τ_m^* is the value of τ_m that minimizes the difference between the experimental and simulated mean decision times, indicated by the crossing point within the colored boxes.

Decisions are modeled as occurring when the accumulated evidence W_t first exits the symmetric interval $(-L_W, L_W)$, where L_W is the decision threshold. If W_t reaches $+L_W$, the model predicts a rightward choice; if it reaches $-L_W$, a leftward choice. This first-exit rule specifies the model’s stopping criterion.

Together, τ_m and L_W specify the decision policy under memory-limited evidence accumulation. We inferred these parameters separately for each participant and condition, and — as we show below — they vary systematically across conditions in ways that reflect the physical properties of the stimulus, such as entropy. This pattern suggests that perceptual decisions are influenced not only by the strength of the sensory evidence but also by how the brain integrates that evidence over time under cognitive constraints.

Following standard practice in cognitive neuroscience, we treated both the temporal integration window τ_m and decision threshold L_W as parameters that vary across participants and conditions. We estimated these parameters by comparing simulated and empirical decision times across a grid of (τ_m, L_W) values generated from the experimental trajectories (see Section 2.5.4.3). To test the robustness of the decision thresholds in predicting the decision times, we performed a cross-validation analysis in Experiment A for each participant in each condition. We estimated the mean threshold from a training set and used it to predict mean T_{dec} in the remaining test set. We found that the experimental mean decision time of the training data accurately matches the predicted mean decision time of the test data (see Section 2.5.4.5; and Supplementary Fig. S2B). This suggests that knowledge of decision thresholds of a particular data set may be used to predict decision times of unseen data.

Fig. 2.5B illustrates this fitting procedure for a representative participant. The intrinsic temporal integration window τ_m^* is defined as the mean exit time from the simulations (solid curves) that most closely matches the participant’s experimentally observed mean decision time (horizontal dashed lines). The points at which these values align are highlighted by the colored box on each curve. The corresponding decision threshold L_W^* is taken from the same parameter pair at this best-matching location. See Supplementary Fig. S8 for the histograms of the probability density of trial-by-trial

EIM decision thresholds W_{dec} obtained from a single participant using Eq.2.14.

First, to assess the robustness of the assumption that each participant exhibits a condition-specific intrinsic temporal integration window, we used a split-half cross-validation analysis (see Section 2.5.4.3). For each participant and condition, we randomly divided trials into two equally sized halves and estimated the integration window parameter separately for each half. The finding that the two sets of temporal integration window values that derived from separate sets of trials lie along a diagonal (see Supplementary Fig. S9) supports the interpretation that τ_m reflects a robust, participant-specific temporal integration window within each condition.

We then validated the EIM by comparing its trial-by-trial predictions with the experimental data (see Fig. 2.6 for Experiment A and Supplementary Fig. S10 for Experiment B). Figs. 2.6A–B show joint histograms of predicted versus observed decision times pooled across all trials, participants, and conditions, using the SPRT (Fig. 2.6A) and the EIM (Fig. 2.6B). We performed the same comparison for decision thresholds in Figs. 2.6C–D, contrasting experimental values with predictions from the SPRT (Fig. 2.6C) and the EIM (Fig. 2.6D). Notably, the distribution of X_{dec} under the SPRT exhibited considerable horizontal spread, despite the use of fixed decision boundaries.

To quantitatively assess model fit, we used the Kolmogorov–Smirnov (KS) test to compare the distributions of predicted and observed decision times and thresholds (see Section 2.5.4.8). For decision times, the KS test showed a significant mismatch between SPRT predictions and the data ($D_{11963} = 0.186$, $p = 0.012$), whereas no significant difference emerged for the EIM ($D_{11963} = 0.103$, $p = 0.421$). The discrepancy for the SPRT stemmed from its consistent tendency to predict faster decisions than those observed experimentally (Fig. 2.6A). For decision thresholds, the EIM again matched the empirical distribution more closely than the SPRT ($D_{11963} = 0.118$, $p = 0.273$ vs. $D_{11963} = 0.262$, $p < 0.0001$, respectively). Together, these results show that the EIM captures inter-trial variability more accurately than the SPRT in Experiment A, and suggest that the EIM is sensitive not only to the average, expected stimulus statistic based on nonequilibrium parameters, but also to individual stimulus trajectories.

In Experiment B, however, we did not observe a comparable level of model performance (see Sup-

plementary Fig. S10). Both the EIM and the SPRT showed significant discrepancies between their predicted and observed decision times ($D_{30018} = 0.038$, $p < 0.0001$ vs. $D_{30018} = 0.06$, $p < 0.0001$, respectively). Likewise, predictions of decision thresholds from both models differed significantly from the empirical distributions ($D_{30018} = 0.147$, $p < 0.0001$ vs. $D_{30018} = 0.312$, $p < 0.0001$, respectively). These findings indicate that the EIM’s performance declined somewhat in Experiment B.

To evaluate the relationship between the integration-window parameter τ_m^* and entropy production, we analyzed the retrieved τ_m^* values from Experiment A using a LME model (see Section 2.5.4.3). A Type III ANOVA on the model estimates revealed a significant main effect of entropy production ($F_{3,60} = 12.45$, $p < 0.0001$). Post-hoc multiple comparisons further showed that all condition pairs differed significantly (all $|t_{60}| > 2.83$, $p < 0.001$), with the exception of the comparison between condition II and condition III. Together, these results indicate that τ_m^* increases with the entropy production rate Σ (see Fig. 2.7A; see also Supplementary Tables S25, and S26 for full ANOVA results and Bonferroni-corrected multiple comparisons).

This pattern suggests that participants tended to extend their integration window when confronted with stimuli exhibiting higher dissipation. Such behavior implies a strategic shift toward greater reliance on memory when stimulus dynamics become more complex and less predictable. Under these conditions, integrating evidence over a longer temporal window can mitigate the influence of local fluctuations and promote more stable, cautious decisions. Conversely, for stimuli closer to equilibrium (i.e., lower dissipation), participants relied more heavily on momentary evidence, consistent with a reduced benefit of temporal integration. Overall, the findings indicate that participants adaptively tuned their integration timescale to the thermodynamic structure of the stimulus, increasing reliance on memory as dissipation rises.

In contrast, the diagonal conditions of Experiment B did not show the same increase in τ_m^* with increasing Σ (Fig. 2.7B). To evaluate whether the intermixed design of Experiment B affected the integration window parameter, we analyzed the recovered τ_m^* values using a LME model (see Section 2.5.4.3). A Type III ANOVA on the model estimates revealed a main effect of entropy

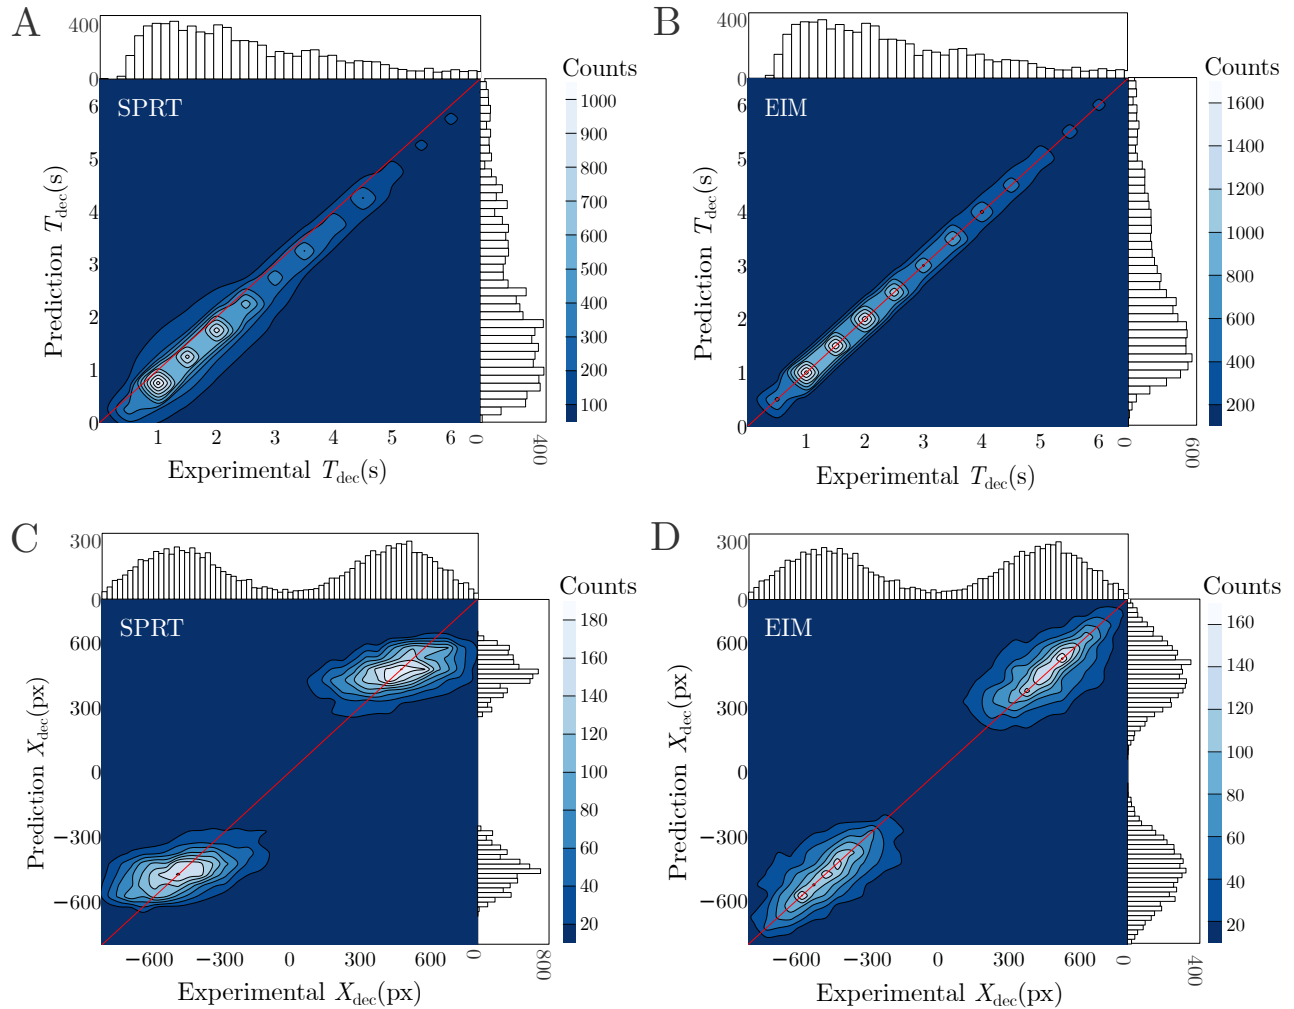


Figure 2.6: **Comparison of SPRT and EIM models to experimental data.** (A–B) Joint histograms comparing the observed trial-by-trial decision times T_{dec} for all participants and conditions with the corresponding predictions from the SPRT (A) and the EIM (B). (C–D) Joint histograms in the same format as panels (A–B), but for the observed decision thresholds X_{dec} . In all panels, the red lines have slope one and indicate perfect agreement between model predictions and experimental data.

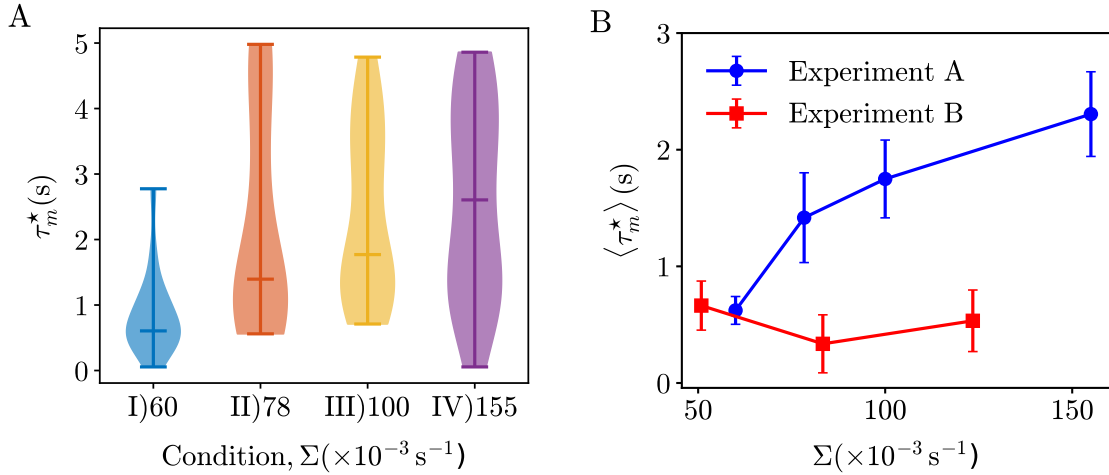


Figure 2.7: **Intrinsic integration window τ_m^* inferred from the EIM fit to human participants' decisions.** (A) Distribution of the inferred τ_m^* values for all participants in each condition of Experiment A. The central horizontal lines indicate the medians. (B) Comparison of the inferred τ_m^* as a function of the conditions of Experiment A (blue) and the three corresponding conditions from Experiment B (red). Error bars represent SEM across participants.

production rate ($F_{2,46} = 6.74$, $p < 0.001$). Post-hoc comparisons indicated that τ_m^* was shortest in the medium-entropy condition, differing significantly from both low and high entropy (all $|t_{46}| > 3.08$, $p < 0.01$). By contrast, low- and high-entropy conditions did not differ from each other. These results indicate that τ_m^* did not scale monotonically with Σ in Experiment B. Full ANOVA results and Bonferroni-corrected multiple comparisons are provided in Supplementary Tables S27, and S28.

These findings suggest that in the blocked structure of Experiment A, participants may adapt their optimal temporal integration window to the Σ governing the current set of trials, effectively carrying forward an entropy-specific τ_m^* from one trial to the next. In this scenario, each trial begins with a prediction about the upcoming stimulus statistics and a corresponding internal τ_m^* . In contrast, the intermixed design of Experiment B likely prevented participants from initiating trials with such a preset integration window. The temporal integration window did not appear to adjust “bottom-up” based on stimulus statistics in real time, and calibrating τ_m^* may require more time than the duration of a single trial permits. We therefore speculate that participants in the intermixed design operated with a relatively short τ_m^* , a pragmatic compromise that supports acceptable performance across unpredictable levels of Σ .

Finally, we evaluated an alternative evidence-integration framework based on an Ornstein–Uhlenbeck (OU) process. In this model, the position term X_{t-s} in the kernel of Eq. (2.12) is replaced by the velocity \dot{X}_{t-s} , yielding $\tau_o \dot{W}_t = -W_t + Hv + \sqrt{2D}\xi_t$, a Markovian stochastic accumulator with temporal integration window τ_o .

As shown in Supplementary Fig. S11, we were unable to identify values of τ_o that reliably reproduced the behavioral data across most experimental conditions. In contrast, the EIM provided accurate fits and robust predictions under the same conditions. A likely explanation is that, in this task, the velocity is substantially less informative about the latent stimulus structure than the position itself. OU-type models are naturally suited to cases where the underlying evidence is unobservable and must be tracked as a hidden state, whereas here the evidence carries essential history-dependent structure. This result underscores the importance of non-Markovian components in the evidence-accumulation process: models that rely solely on Markovian updates fail to capture the temporal dependencies that appear necessary to explain the observed decisions.

2.4 Discussion

On the basis of statistics obtained from human participants performing a perceptual decision-making task, this work demonstrates that judgments of nonequilibrium phenomena are constrained by fundamental thermodynamic trade-offs among decision time, accuracy, and entropy production. To our knowledge, this study represents the first exploration within cognitive neuroscience of recently proposed thermodynamic uncertainty relations for Markovian systems [13], [42]–[44], and thus introduces the potential utility of such theoretical tools as a means to probe decision-making processes in the brain.

We summarize the main findings of our study as follows. i) The moving-disk paradigm provides explicit access to the evidence available to participants, enabling the application of stochastic thermodynamics and a precise computation of the ideal observer’s performance. ii) In Experiment A, ideal SPRT performance exceeds human performance in a small but systematic manner. iii) This suboptimality is reproduced by incorporating memory constraints via the EIM. iv) Participants

adjust the EIM integration window τ_m^* according to the entropy production rate of the stimulus. v) In Experiment B, intermixing stimuli with different entropy production rates reduces performance, seemingly because participants apply a single integration window across trials; this parameter does not reset on a trial-by-trial basis.

Human perceptual decision-making is typically modeled by fitting the accumulation of an unobservable evidence variable using parameters such as drift rate and decision threshold. In contrast, our approach exploits a paradigm in which the evidence variable — the instantaneous position of a moving stimulus — is directly observable. Having explicit access to the quantity of available sensory evidence enables us to apply principles from stochastic thermodynamics to empirical data, and to quantify deviations of human decisions from the optimality defined by the Sequential Probability Ratio Test (SPRT). The observed deviations in both decision times and thresholds were small but systematic, motivating the search for a mechanistic explanation of the brain’s suboptimality.

Because the nervous system, unlike the SPRT, is subject to limitations in information storage [45], we introduced a “forgetting” mechanism, formalized as the Evidence Integration Model (EIM). The EIM accurately predicts participants’ choices on a trial-by-trial basis. Crucially, alternative Markovian models such as the Ornstein–Uhlenbeck (OU) process failed to account for our findings, validating the non-Markovian nature of evidence accumulation in our paradigm. Moreover, the dependence of the EIM’s integration window parameter on the stimuli’s entropy production rate suggests that participants adjusted their temporal integration window according to the dissipation properties of the observed dynamics. In Experiment A, for example, participants adopted a longer integration window when integrating evidence from stimuli that are farther from equilibrium. Adjusting the temporal integration window has been proposed as an adaptive strategy for mitigating stimulus volatility by averaging out noise and preventing false alarms [4], [5], [12]. It is tempting to speculate that the integration window parameter identified here may relate to those inferred in previous studies where participants’ estimates of stimulus intensity or duration are modeled through leaky accumulation processes [46]–[52].

Our findings can be situated within a broader literature on optimality in decision making. Prior

work shows that deviations from optimality can be reduced through extensive training [53]–[57]. In the present study, participants were neither over-trained nor informed about the structure of the task. Thus, any learning about stimulus parameters across conditions had to occur implicitly. Furthermore, because sessions contained a fixed number of trials and rewards were not contingent on response speed, participants had no incentive to favor faster decisions. This contrasts with paradigms using fixed-duration sessions, where maximizing reward rate naturally encourages speed-weighted strategies.

Exploring the neural dynamics governing perceptual decision-making during non-equilibrium fluctuations offers a promising venue for future research. NMDA-mediated synaptic mechanisms, in particular, have been proposed as modulators of the temporal integration window, and may offer a biologically plausible implementation of the timescale adjustments observed here [58]–[60]. Beyond neurobiological mechanisms, macroscale brain dynamics also provide promising avenues. Neuroimaging studies have shown that human and primate brain activity exhibits speed–accuracy trade-offs [17] and measurable levels of temporal irreversibility [61], with these features varying across cognitive tasks [61], states of consciousness [62], [63], and clinical conditions [64], [65]. These perspectives not only deepen our understanding of human cognition, but also shed light on the general principles governing complex systems that must make swift and accurate decisions in changing environments.

2.5 Materials and Methods

2.5.1 Participants

All the participants were tested in the Cognitive Neuroscience laboratory of SISSA. Prior to the experiment, the written consent form, approved by the Ethics Committee of the SISSA was obtained from the participants. At the end of the experiments, all participants were provided with a debriefing session. They received 10 euros per hour for taking part.

In Experiment A, twenty-one right-handed participants (12F, 9M, mean age=30.76, SD=8.8, all

right-handed) with normal or corrected-to-normal vision were recruited. The participants were tested for approximately 1 to 1.5 hours depending on their own pace.

Among the participants of Experiment A, nine individuals also participated in the control experiment.

In Experiment B, twenty-four participants (11F, 13M, mean age=27.17, SD=3.48, 1 left-handed) with normal or corrected-to-normal vision are recruited. None of them participated in Experiment A. The participants were tested for approximately 4 hours, separated into 2 sessions conducted on different days.

2.5.2 Stimuli and apparatus

At the beginning of each trial, a dark gray disk (diameter of 2° visual angle) was displayed at the center of the screen against a light gray background. From there, at time 0, the disk began to move in a trajectory defined by the nonequilibrium dynamics of an overdamped Brownian particle in two dimensions. It is important to note that such simulations mimic the fluctuating motion of a floating particle, like dust seen in a beam of light. In physical terms, each disk trajectory was generated by performing an Euler numerical integration of the Langevin equation $\dot{X}_t = Hv + \sqrt{2D}\xi_t$ and $\dot{Y}_t = \sqrt{2D}\eta_t$. Here, $v > 0$ is the drift velocity along the x -axis, and $D > 0$ is the diffusion coefficient, which are equal along the x and y axes. The corresponding (v, D) parameter values can be found in Supplementary Table S1 for Experiment A and in Supplementary Fig. S4 for Experiment B. The left/right direction of the drift in each trial, defined by H , was set to -1 or 1 with equal probability. The quantities ξ_t and η_t are independent Gaussian white noises, each with zero mean and satisfying correlations $\langle \xi_t \xi_s \rangle = \langle \eta_t \eta_s \rangle = \delta(t - s)$, $\langle \xi_t \eta_s \rangle = 0$ for all t, s . this is an overdamped Brownian particle dragged at a constant net speed v in a Newtonian fluid with viscosity $\gamma = k_B T / D$, where k_B is Boltzmann's constant and T is the temperature of the fluid.

The trial finished either when the subject responded or when the disk reached one of the screen edges along x -axis (right/left). If the top or bottom border was reached the disk was reflected back to its previous vertical position. The discrete time step in numerical simulations was the inverse of

the screen’s refresh rate.

In Experiment A, stimuli were displayed on a Samsung 2233RZ series business monitor (resolution = 1680×1050 pixels, refresh rate = 120 Hz). In Experiment B, stimuli were displayed on a BenQ ZOWIE XL2731 monitor (resolution = 1920×1080 pixels, refresh rate = 144 Hz). In the control experiment, stimuli were shown on a BenQ ZOWIE XL2731 monitor (resolution = 1920×1080 pixels, refresh rate = 144 Hz). The temporal resolution in the motor implementation experiment was set the same as in Experiment A. A chin rest was used to restrict head movements and comfortably fix head position at a distance of 57 cm from the computer screen.

2.5.3 Task and experimental design

Both Experiments A and B employed a two-alternative forced-choice (2AFC) design. On each trial, participants judged whether the disk’s net motion was rightward or leftward by pressing the corresponding arrow key before the disk reached either screen edge (Fig. 2.1A). No information regarding the underlying generative dynamics of the stimulus was provided.

For each trial, we recorded: i) the true drift direction $H = \{1, -1\}$; ii) the decision outcome $D_{\text{dec}} = \{1, -1\}$; iii) the decision time T_{dec} ; iv) the horizontal position of the disk at decision time, $X_{\text{dec}} = X(t = T_{\text{dec}})$, corresponding to the single-trial decision threshold; and v) the full trajectory $X_{[0, T_{\text{dec}}]}$ up to the moment of decision. The vertical position at decision time, Y_{dec} , was also recorded but proved not to be critical to any analysis.

After completing 20 training trials, participants proceeded to the testing session. In both training and testing, participants received immediate feedback: correct responses yielded +35 and +40 points in Experiments A and B, respectively; incorrect responses yielded -20 points; and missed responses (failure to respond before the disk reached an edge) resulted in -50 points. Feedback was color-coded (green for positive and red for negative outcomes), accompanied by a brief tone, and displayed with the cumulative point total. When ready, participants initiated each new trial by pressing the space bar (see Supplementary Fig. S1).

Experiment A. Experiment A comprised four conditions, defined by combinations of drift velocity and diffusivity. Each condition contained 3 blocks of 50 trials, for a total of 12 blocks (600 trials). Importantly, each block contained only trials from a single condition (block-design).

Experiment B. Experiment B included nine conditions, formed by the combination of three drift velocities and three diffusivities. Each condition was presented in 144 trials, arranged in 36 blocks of 36 trials, yielding a total of 1296 trials. Unlike in Experiment A, the trials from all nine conditions were intermixed in pseudorandom order, such that each 36-trial block contained four trials from each condition.

Control experiment. A separate control experiment was conducted to characterize response behavior when the decision boundary was externally visible rather than internally computed from accumulated evidence. It consisted of three conditions with different stimulus parameters, presented in randomized order across participants. Each condition included 160 trials divided into four blocks. In all cases, participants were instructed to press the corresponding arrow key as soon as the moving disk reached a prominent red line on either the left or right side of the display.

The "no diffusion" condition used stimuli with pure drift ($v = 480$ px/s). In the "Y-axis diffusion" condition, the disk had the same drift but included diffusivity only in the vertical dimension (21600 px²/s). The "original parameter" condition reproduced the velocity and diffusivity used in Experiment A. No trial-by-trial feedback was provided. In this control task, participants were not required to judge motion direction; they simply reacted when the stimulus reached one of the visible boundaries.

Note that, in all experiments, to assess decision consistency across participants for identical stimuli, each participant viewed the same set of stochastic trajectories, but the order of trial presentation was randomized across individuals.

All stimulus presentation and data collection were implemented in MATLAB using Psychophysics Toolbox Version 3 (PTB-3) [66].

2.5.4 Data analysis

2.5.4.1 Behavioral performance analysis

The behavioral analysis aimed to assess the effects of the experimental manipulations (entropy production rate condition, velocity, and diffusivity) on the decision metrics: decision threshold, accuracy, and decision time. In both Experiments A and B, these metrics were evaluated using linear mixed-effects (LME) models.

Prior to analysis, we removed all "missed trials," defined as cases in which participants failed to make a decision before the disk reached either screen edge ($\sim 5\%$ of trials in Experiment A and $\sim 3\%$ in Experiment B). We also excluded "false start" trials, in which $T_{\text{dec}} < 150$ ms. Such trials occurred only in Experiment B ($\sim 0.06\%$) and were absent in Experiment A.

Decision threshold, X_{dec} . Trial-by-trial X_{dec} data of both experiments were transformed as $\text{sqrt}(\text{constant-abs}(X_{\text{dec}}))$ due to the left-skewed distribution before taking the average for each condition and participant. Then, transformed decision thresholds were analyzed with two LME models with Eq. 2.1 for Experiment A (marginal R^2 : 0.007, conditional R^2 : 0.83), and Eq. 2.2 for Experiment B (marginal R^2 : 0.014, conditional R^2 : 0.93), where id refers to the participant identity:

$$\langle X_{\text{dec}} \rangle \sim \text{condition} + (1 \mid \text{id}). \quad (2.1)$$

$$\langle X_{\text{dec}} \rangle \sim v * D + (1 \mid \text{id}). \quad (2.2)$$

The fixed effects used to model the behavioral metrics in Experiment A was entropy production rate (referred as condition), while in Experiment B both velocity v and diffusivity D , as well as

their interaction were used as fixed effects. In all experiments participant id was taken as a random effect.

Accuracy, α . Accuracy data were analyzed with LME models with Eq. 2.3 in Experiment A (marginal R^2 : 0.25, conditional R^2 : 0.74), and Eq. 2.4 in Experiment B (marginal R^2 : 0.84, conditional R^2 : 0.93). The same fixed and random effects were used as before:

$$\alpha \sim \text{condition} + (1 \mid \text{id}). \quad (2.3)$$

$$\alpha \sim v^*D + (1 \mid \text{id}). \quad (2.4)$$

Decision time, T_{dec} . Prior to the analysis, T_{dec} data were transformed as $(-\text{constant}/T_{\text{dec}})$ in Experiment A, and log-transformed in Experiment B. Then, transformed T_{dec} data were analyzed with LME models with Eq. 2.5 in Experiment A (marginal R^2 : 0.52, conditional R^2 : 0.92), and Eq. 2.6 in Experiment B (marginal R^2 : 0.59, conditional R^2 : 0.96). The same fixed and random effects were used as before:

$$\langle T_{\text{dec}} \rangle \sim \text{condition} + (1 \mid \text{id}). \quad (2.5)$$

$$\langle T_{\text{dec}} \rangle \sim v^*D + (1 \mid \text{id}). \quad (2.6)$$

All LME models were fitted with the lme4 package [67]. LME model variance explained was computed using the MuMIn package [68]. The Satterwaite’s method [69] implemented in the lmerTest package[70] was used to estimate the degrees of freedom for the LME model ANOVA, and estimated marginal means were compared using the emmeans package[71] using Bonferroni adjustment.

2.5.4.2 Sequential probability ratio test (SPRT)

To compare observed human decision-making with that of an optimal observer, we used the Sequential Probability Ratio Test (SPRT) [21], [23], which prescribes how evidence should ideally be accumulated to discriminate between two competing hypotheses. In our setting, the task is to infer the direction of the disk’s drift (left vs. right), corresponding to hypotheses H_1 and H_2 .

In the SPRT, the accumulated evidence is expressed as the log-likelihood ratio

$$\mathcal{L}_t = \ln \left[\frac{\mathcal{P}(X_{[0,t]} | H_1)}{\mathcal{P}(X_{[0,t]} | H_2)} \right], \quad (2.7)$$

where $\mathcal{P}(X_{[0,t]} | H_i)$ denotes the probability of observing trajectory $X_{[0,t]}$ under hypothesis H_i .

For a one-dimensional drift–diffusion process (1DDM), where X_t is driven by drift v and diffusivity D , this expression simplifies to

$$\mathcal{L}_t = \left(\frac{v}{D} \right) X_t, \quad (2.8)$$

a quantity that also corresponds to the stochastic entropy production accumulated along the trajectory [13], [72]. This establishes a direct connection between statistical inference and thermodynamic irreversibility.

The SPRT defines the decision time T_{dec} as the first moment at which \mathcal{L}_t exits the symmetric interval $(-\lambda, \lambda)$, where λ is the decision threshold selected by the observer to achieve a desired accuracy. Under this model, the mean decision time depends on both the entropy production rate and the Péclet number:

$$\langle T_{\text{dec}} \rangle_{\text{SPRT}} = \frac{\text{Pe}}{\Sigma} \tanh \left(\frac{\text{Pe}}{2} \right), \quad (2.9)$$

where $\Sigma = v^2/D$ is the entropy production rate and $\text{Pe} = vL/D$ is the Péclet number, with L denoting the decision threshold in position space. This relation reflects an inverse dependence of mean decision time on dissipation and selected decision threshold.

The corresponding prediction for the decision accuracy is

$$\alpha_{\text{SPRT}} = \frac{1}{1 + \exp(-\text{Pe})}, \quad (2.10)$$

a sigmoidal function of the Péclet number. This expression is consistent with the fact that once the decision threshold (and thus the accuracy target) is fixed via λ , the combination of stimulus parameters (v, D) determines the resulting Péclet number and hence the predicted accuracy.

Together, Eqs. (2.9) and (2.10) provide a compact normative link between entropy production and the expected speed and accuracy of decisions.

Finally, comparing any real agent with the optimal SPRT yields a thermodynamic uncertainty-type bound:

$$\Sigma \langle T_{\text{dec}} \rangle \geq \text{Pe} \tanh\left(\frac{\text{Pe}}{2}\right), \quad (2.11)$$

which places a fundamental lower limit on the mean decision time given i) the irreversibility of the stimulus, and ii) the Péclet number. Inequality Eq. (2.11) represents a novel application of thermodynamic uncertainty relations (TURs) [42]–[44], [73] to perceptual decision-making. This type of bound was originally conjectured for Markovian processes [13] and later proven for generic nonequilibrium steady states [74]. Further details and derivations are provided in the Supplementary Section A.

2.5.4.3 Evidence integration model (EIM)

Continuous time interpretation. To capture memory dynamics that may constrain how the brain integrates evidence, we introduced an Evidence Integration Model (EIM) that incorporates a parameter — absent in the SPRT — that reflects loss of stored information over time. In the EIM, evidence accumulated up to time t is represented by an accumulator W_t , defined as a functional of the trajectory $X_{[0,t]}$:

$$W_t \equiv \frac{1}{\tau_m} \int_0^t ds e^{-s/\tau_m} X_{t-s}. \quad (2.12)$$

Here, τ_m is the temporal integration window. Thus, W_t corresponds to an exponentially weighted moving average of past stimulus positions, where more recent values of X_{t-s} contribute more strongly than older ones. In other words, the integral in Eq. (2.12) reads information backward in time, discounting past evidence according to the factor e^{-s/τ_m} relative to the evidence available at time t . The accumulation process terminates once W_t reaches the decision threshold L_W .

Differentiating Eq. (2.12) yields the relation $\tau_m \dot{W}_t = -W_t + X_t$. When the stimulus evolves according to the stochastic process $\dot{X}_t = Hv + \sqrt{2D} \xi_t$, the accumulator W_t obeys the underdamped Langevin equation

$$\tau_m \ddot{W}_t = -\dot{W}_t + Hv + \sqrt{2D} \xi_t, \quad (2.13)$$

where τ_m acts as an effective inertial relaxation timescale.

Discrete time interpretation. Eq. (2.12) can also be written in discrete time with time step Δt :

$$W_{t+\Delta t} = e^{-\Delta t/\tau_m} W_t + \frac{\Delta t}{\tau_m} X_{t+\Delta t}. \quad (2.14)$$

The first term on the right-hand side represents the evidence W_t accumulated up to time t , scaled by the exponential factor $e^{-\Delta t/\tau_m}$. As before, increasing the temporal integration window τ_m increases memory retention by reducing the rate at which older evidence is discounted.

The second term in Eq. (2.14) contributes the most recent observation, $X_{t+\Delta t}$, as instantaneous evidence added to the accumulator. A practical advantage of this discrete formulation is that Δt can be aligned with the refresh interval of the visual stimulus.

Fitting procedure. We fitted the EIM to the T_{dec} data for each participant and each experimental condition separately. To recover the optimal temporal integration window τ_m^* for a participant in a given condition, we evaluated a range of candidate τ_m values. For each candidate τ_m , we computed W_t using Eq. (2.12) along the empirical trajectories $X_{[0, T_{\text{dec}}]}$.

For each trajectory, we obtained the decision threshold $L_W = \langle |W_{\text{dec}}| \rangle = \langle |W_{T_{\text{dec}}}| \rangle$ defined as the

mean absolute value of the accumulator at decision time. This yielded a collection of parameter pairs $\{(\tau_m, L_W)\}$.

For each pair (τ_m, L_W) , we simulated the EIM to compute the mean exit time of W_t from the interval $(-L_W, L_W)$. The fitted temporal integration window τ_m^* for a given participant was defined as the value whose simulated mean exit time most closely matched that participant’s observed mean decision time.

Cross-validation of τ_m^* . Here, we checked the robustness of our assumption that each participant builds for each condition an intrinsic τ_m^* . To do this, we randomly split trials into two equally sized halves for each participant and in each condition. We then fitted the EIM separately to the reaction-time data from each half, using the same fitting procedure as above, obtaining two independent estimates of the integration window, denoted by τ_m^\bullet and τ_m° . Robustness was evaluated by comparing τ_m^\bullet and τ_m° across splits.

Analysis of temporal integration window parameter, τ_m^* . We analyzed the τ_m parameters (square-root-transformed) retrieved from the EIM fits for each participant in Experiment A using the following LME with formula (2.15), with condition as the fixed effect and participant (id) as the random effect (marginal $R^2 = 0.20$, conditional $R^2 = 0.57$):

$$\tau_{m,A}^* \sim \text{condition} + (1 \mid \text{id}). \tag{2.15}$$

We applied the same analysis to the square-root-transformed τ_m parameters retrieved from the EIM fits for each participant in the three comparable conditions of Experiment B. The data were modeled using the following LME with formula (2.16), again with condition as a fixed effect and participant (id) as a random effect (marginal $R^2 = 0.06$, conditional $R^2 = 0.67$):

$$\tau_{m,B}^* \sim \text{condition} + (1 \mid \text{id}). \tag{2.16}$$

We applied the same methodological procedures described in Section 2.5.4.1.

2.5.4.4 Ornstein–Uhlenbeck model

In order to compare with the EIM, we also used Ornstein–Uhlenbeck model. If we replace X_{t-s} with \dot{X}_{t-s} in defining the evidence accumulator W_t (see Eq. (2.12)), then W_t becomes

$$W_t \equiv \frac{1}{\tau_o} \int_0^t ds e^{-s/\tau_o} \dot{X}_{t-s} \quad (2.17)$$

from which one may write $\tau_o \dot{W}_t = -W_t + \dot{X}_t$. For $\dot{X}_t = Hv + \sqrt{2D}\xi_t$, the accumulator W_t obeys an Ornstein–Uhlenbeck process

$$\tau_o \dot{W}_t = -W_t + Hv + \sqrt{2D}\xi_t. \quad (2.18)$$

We did the analysis to infer the optimal temporal integration window τ_o^* for each individual for the case of the Ornstein–Uhlenbeck (OU) model, accumulating evidence through W_t in Eq. (2.18). We followed an analogous procedure to that used for EIM to the OU model to estimate τ_o^* of each participant under each experimental condition.

2.5.4.5 Cross-validation of mean decision times

In order to test the robustness of the decision thresholds in predicting the decision times (in both SPRT and EIM), we performed a cross-validation analysis in Experiment A for each participant in each condition. We divided the trial-by-trial data into two equal-sized halves at random (training and test) and use the mean decision thresholds of the training data to predict T_{dec} for the test data. Note that for the case of EIM to obtain the decision threshold, we use τ_m^* in the training set trials.

2.5.4.6 Comparison between observed data and SPRT predictions

We computed two metrics capturing the discrepancy between the observed data and the SPRT-predicted values for each participant and condition in Experiment A: i) deviations in accuracy predictions, quantified as the logit absolute error of accuracy (AE_α); and ii) deviations in decision time predictions, quantified as the relative absolute error of decision times ($AE_{T_{\text{dec}}}$).

Each metric was analyzed using a linear mixed-effects (LME) model with a fixed effect of entropy production rate (condition) and a random effect of participant (id).

The logit absolute error of accuracy was modeled using Eq. 2.19 (marginal $R^2 = 0.26$, conditional $R^2 = 0.30$), and the relative absolute error of decision times was modeled using Eq. 2.20 (marginal $R^2 = 0.27$, conditional $R^2 = 0.60$):

$$AE_\alpha \sim \text{condition} + (1 \mid \text{id}). \quad (2.19)$$

$$AE_{T_{\text{dec}}} \sim \text{condition} + (1 \mid \text{id}). \quad (2.20)$$

We applied the same methodological procedures described in Section 2.5.4.1.

2.5.4.7 Comparing behavioral metrics between Experiments A and B

To evaluate how the intermixed design of Experiment B affected decision behavior relative to the block design of Experiment A, we compared decision thresholds X_{dec} , accuracies α , and decision times T_{dec} across the three comparable conditions shared between the two experiments (see Supplementary Fig. S4). In all analyses, the entropy production rate Σ was treated as a continuous covariate.

Decision thresholds, X_{dec} . Decision thresholds X_{dec} were analyzed using the linear mixed-effects (LME) model in Eq. 2.21 (marginal $R^2 = 0.05$, conditional $R^2 = 0.22$):

$$X_{\text{dec}} \sim \text{experiment} * \Sigma + (\Sigma | \text{id}). \quad (2.21)$$

The fixed-effects structure included: i) experiment (A vs. B), and ii) the entropy production rate Σ , as well as their interaction. Participant-specific random slopes for Σ were included as random effects.

Accuracies, α . We aggregated accuracy values α by participant and condition, and analyzed them using the LME model in Eq. 2.22 (marginal $R^2 = 0.05$, conditional $R^2 = 0.54$), with the same fixed and random effects structure as the model for X_{dec} :

$$\alpha \sim \text{experiment} * \Sigma + (\Sigma | \text{id}). \quad (2.22)$$

Decision times, T_{dec} . Log-transformed decision times were analyzed using the LME model in Eq. 2.23 (marginal $R^2 = 0.22$, conditional $R^2 = 0.32$), again with the same fixed and random effects structure:

$$\log T_{\text{dec}} \sim \text{experiment} * \Sigma + (\Sigma | \text{id}). \quad (2.23)$$

We applied the same methodological procedures described in Section 2.5.4.1.

2.5.4.8 Comparison between SPRT and EIM

To evaluate how well the model simulations captured the empirical distributions of participants' behavior, we compared trial-by-trial predictions of decision times and decision thresholds from both the SPRT and the EIM against the observed data. Specifically, we conducted two-sample Kolmogorov–Smirnov (KS) tests between the model-predicted and empirically observed distributions.

The KS test quantifies the maximum distance between the cumulative distribution functions (CDFs) of two samples, providing a non-parametric measure of their distributional dissimilarity. Tests were implemented in Python using the SciPy function `ks_2samp`, which returns both the KS statistic and the associated p -value. Lower KS statistics indicate greater similarity between the predicted and observed distributions. Separate comparisons were performed for decision times and for decision thresholds.

2.5.4.9 Consistency between participants' decision thresholds

We explored the variance of the decision thresholds further by analyzing consistency among participants' decision thresholds for the same trials. To proceed, let us denote the observed decision threshold of participant p in trial i for condition c by $X_{\text{dec}}^{p,c,i}$. For the participant p and condition c , we computed in each trial i the relative difference between the decision threshold $X_{\text{dec}}^{p,c,i}$ and its average $\langle |X_{\text{dec}}^{p,c,j}| \rangle_j$ over all trials j :

$$\Delta X_{\text{dec}}^{p,c,i} \equiv \frac{\langle |X_{\text{dec}}^{p,c,j}| \rangle_j - |X_{\text{dec}}^{p,c,i}|}{\langle |X_{\text{dec}}^{p,c,j}| \rangle_j}. \quad (2.24)$$

Then, we took an average of $\Delta X_{\text{dec}}^{p,c,i}$ over the participants, i.e., $\langle \Delta X_{\text{dec}}^{p,c,i} \rangle_p$, as a function of trials (i). We then calculated the null distribution of $\Delta X_{\text{dec}}^{p,c,i}$ by shuffling trials in order to find the 95% confidence intervals.

Chapter 3

A Closer Look into Stimulus Dynamics Influencing Decision Under Nonequilibrium Fluctuations

Aybüke Durmaz¹, Gianfranco Fortunato¹, Edgar Roldán², Mathew E. Diamond¹, Domenica Buetti¹

1 International School for Advanced Studies (SISSA), Trieste, Italy

2 The Abdus Salam International Center for Theoretical Physics (ICTP), Trieste, Italy

3.1 Abstract

This chapter provides an extended analysis of the data presented in Chapter 2, with the goal of identifying which aspects of stimulus dynamics shape perceptual decisions under nonequilibrium fluctuations. In particular, we examine how human observers temporally weight stochastic evidence and which features of the stimulus trajectory exert the strongest influence on their choices.

First, using a decision-locked reverse-correlation analysis, we observed a robust recency effect in both Experiments A and B: choices tended to occur shortly after informative segments of the trajectory—brief, directionally coherent “jumps” (defined below) — while earlier portions of the trajectory were progressively discounted. Importantly, in Experiment A, where entropy-production conditions were presented in a blocked design, the relative weighting of recent evidence compared to earlier segments became increasingly pronounced at higher entropy-production rates. In Experiment B, where entropy-production levels varied on a trial-by-trial basis, the same qualitative recency effect was present but showed much weaker modulation across conditions. This pattern

suggests that observers adjusted their temporal integration strategy in accordance with the broader task structure: when participants knew the upcoming entropy-production level, they exhibited a steeper weighting function for high-entropy stimuli, whereas when they could not anticipate the entropy level, they instead “smoothed” their weighting function, treating each trial as if the expected entropy were low.

Second, using an index quantifying the influence of brief trajectory segments, we found that coherent, salient changes in motion (“jumps”) reliably evoked choices, with later jumps exerting the strongest impact in both experiments. This indicates that observers were sensitive not only to continuous moment-to-moment fluctuations in the motion stream but also to short, high-information events embedded within it. As with the recency effect, sensitivity to these salient jumps was more pronounced in the blocked design of Experiment A than in the intermixed design of Experiment B.

These findings align with predictions from the Evidence Integration Model (EIM), in which a finite memory window (τ_m) naturally produces recency through leaky accumulation and amplifies the impact of brief, low-noise evidence. Taken together, our results indicate that participants rely on a temporally graded, feature-sensitive integration strategy when forming perceptual decisions under nonequilibrium fluctuations.

3.2 Introduction

The previous chapter (Chapter 2) established a link between the thermodynamic structure of stochastic stimuli and human perceptual decisions. Using a visual motion stimulus generated by a Brownian process with controlled drift v and diffusion D , we showed that participants’ mean decision times $\langle T_{\text{dec}} \rangle$ systematically decreased with the entropy production rate—defined as the ratio $\Sigma = v^2/D$ —in line with the normative predictions of the Sequential Probability Ratio Test (SPRT; Eq. 2.11). However, behavior deviated from SPRT in a small but systematic way: decision times were consistently slower than optimal. Introducing the Evidence Integration Model (EIM) accounted for this suboptimality and captured trial to trial variability more effectively than SPRT

in Experiment A. Crucially, the temporal integration window τ_m^* — which reflects the duration of stimulus history retained in memory — increased with the entropy production rate in the blocked design of Experiment A. This result suggests that observers adjusted their temporal integration policy to the dissipation level (i.e., entropy production rate) of the stimulus. Notably, this adjustment emerged only when participants experienced the same dissipation level across many consecutive trials. In contrast, in the intermixed design of Experiment B — where dissipation levels changed on a trial by trial basis — the effect was markedly attenuated. Our interpretation was that, under intermixed conditions, trial to trial changes in dissipation increased participants’ uncertainty about the upcoming level, thereby weakening the influence of entropy production on the integration window τ_m^* .

Overall, these results characterize how well observers approximate a normative bound given the nonequilibrium statistics of the input, and reveal that memory dependent integration of stimulus trajectories depends on both the dissipation level and the task structure. However, they do not clarify how participants sample information within each trial or which aspects of the stimulus most strongly influence decisions. Two key questions therefore remain open. i) Do observers integrate evidence uniformly across the trial, or do they weight early versus late samples differently (primacy vs. recency)? ii) Beyond generic accumulation, are observers sensitive to brief, statistically meaningful cues embedded within the trajectory—such as short lived directional changes (“jumps”)?

To address these questions, this chapter presents two complementary analyses applied to Experiments A and B. First, we estimated temporal weighting profiles using choice time-locked reverse correlation to quantify the contribution of each moment in the stimulus trajectory to the final choice. This well established approach distinguishes between primacy and recency strategies in perceptual decision making and provides a direct test of time dependent versus uniform integration [38], [75]–[79]. It allowed us to uncover the temporal policy of evidence accumulation and to derive an empirical, within trial signature of participants’ memory dynamics.

Second, we examined participants’ sensitivity to brief, salient directional changes within the stimulus trajectories. We quantified whether short, consistent motion segments (“jumps”) were de-

tectable above the background noise inherent to the stochastic process. A related work demonstrated that stimulus consistency (i.e., successive evidence samples tend to support the same choice alternative) increases the weight assigned to evidence consistent with the recent history and improves model predictions [80]. This analysis enabled us to determine how these jumps influenced choices in both experiments and to assess whether participants relied on structured, short-lived events when interpreting nonequilibrium fluctuations.

Together, these two analyses provide a more mechanistic account of how evidence is temporally weighted and how specific trajectory features shape decision making. Combined with the previous chapter’s results, they allow us to understand not only how well participants approximate optimality, but also how they construct decisions from stochastic, nonequilibrium dynamics.

3.3 Results

3.3.1 Temporal weighting profile

To characterize how participants weighted stimulus information over time within each trial, we estimated temporal weighting profiles using a reverse-correlation analysis (see Section 3.5.1.1). For each participant, we computed a reverse-correlation kernel that captured the contribution of each moment in the stimulus trajectory to the final decision. From each kernel, we then derived a recency index (RI), which summarizes whether choices were driven more strongly by early (primacy) or late (recency) portions of the trajectory. To compute this metric, we split each kernel into two equal temporal halves and calculated the difference between their mean weights. Positive RI values indicate recency — a greater influence of later evidence — whereas negative values indicate primacy, reflecting stronger weighting of early evidence. This analysis was performed for both Experiments A and B.

3.3.1.1 Experiment A

Figure 3.1A shows the reverse-correlation kernels averaged across participants and color-coded by entropy-production condition. To assess whether the entropy production rate influenced individual

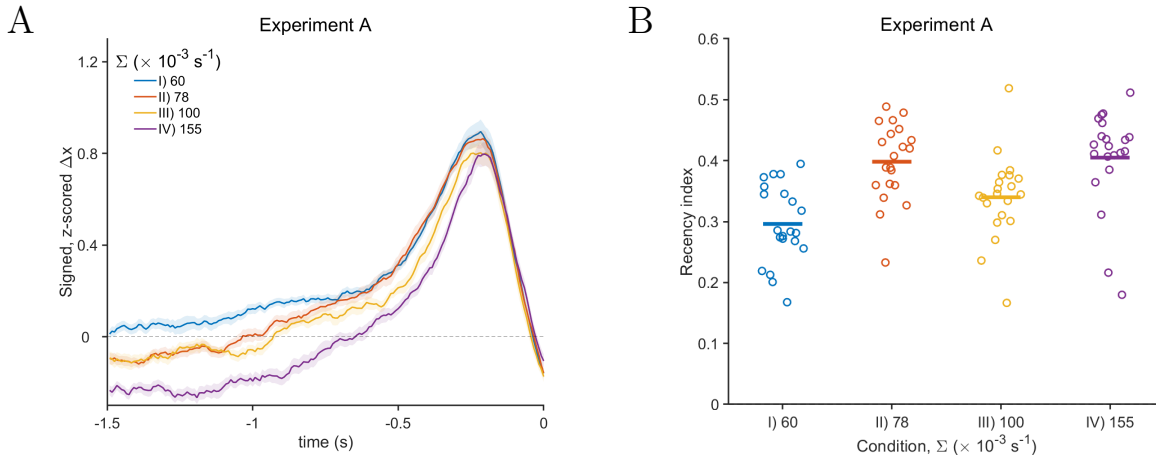


Figure 3.1: **Temporal weighting profiles across entropy-production rates in Experiment A.** (A) Reverse correlation kernels averaged across participants for each entropy production condition. Kernels were computed as decision-signed averages of standardized stimulus displacements across trials. The x -axis corresponds to the final 1.5 seconds of each trial, with time 0 marking the instant of decision. (B) RI values for each participant. Colored horizontal bars show the mean RI for each entropy production condition.

temporal weighting profiles — as captured by the decision-locked recency index (RI_{dec} ; Fig. 3.1B) — we fitted a linear mixed-effects model (LME; see Section 3.5.1.1). A Type III ANOVA on the model estimates revealed a significant main effect of entropy production rate ($F_{3,63} = 20.49$, $p < 0.001$; see Supplementary Table S29 for the details of the ANOVA results).

Moreover, post-hoc multiple comparisons revealed that recency indices increased progressively with increasing entropy production rate, with one exception: condition II showed slightly higher values than condition III ($t_{66.2} = 3.54$, $p = 0.0045$). All other adjacent entropy conditions exhibited a stronger recency effect than the next lower condition (all $|t_{66.2}| > 3.95$, all $p < 0.0012$), except for the comparison between conditions I and III, which only approached significance ($t_{66.2} = 2.65$, $p = 0.06$). Bonferroni-corrected multiple comparisons are reported in Supplementary Table S30.

These findings show that recency increased with entropy production rate, producing a near-graded change in the decision-locked kernels across entropy conditions. Building on these observations, we next examine whether a similar pattern emerged in Experiment B.

3.3.1.2 Experiment B

As discussed in Chapter 2, Experiment B extended our analysis to a broader range of entropy production levels; further, unlike Experiment A, which used a blocked design, Experiment B employed an intermixed design in which the entropy production rate varied unpredictably from trial to trial. This manipulation allowed us to test whether the recency weighting pattern observed in Experiment A would generalize when participants could not anticipate the upcoming entropy level. As reported in Chapter 2, participants in Experiment B showed reduced Wald optimality, shorter temporal integration windows, and — critically — an integration window that no longer scaled systematically with entropy production. Here, we asked how these behavioral changes affected the temporal weighting of the stimulus trajectory.

Figure 3.2A presents the reverse correlation kernels averaged across participants for each entropy condition (color-coded). We first examined whether velocity or diffusivity individually contributed to the systematic relationship between recency and entropy observed in Experiment A.

Figures 3.2C and 3.2D show participant-level recency indices grouped by velocity and diffusivity, respectively. We analyzed these recency indices using a LME (see Section 3.5.1.1). A Type III ANOVA on the model estimates revealed significant main effects of velocity ($F_{2,192} = 14.72$, $p < 0.001$) and diffusivity ($F_{2,192} = 16.46$, $p < 0.001$), as well as a significant interaction between the two factors ($F_{4,192} = 12.89$, $p < 0.001$). Full statistical results are provided in Supplementary Table S31.

We found a larger recency effect in high v (all $|t_{200}| > 3.88$, all $p < 0.004$ except one pair), and high D (all $|t_{200}| > 3.62$, all $p < 0.001$ except one pair) conditions (see Supplementary Tables S32 and S33). However, the interaction between velocity and diffusivity did not reveal a systematic, graded effect of the velocity-diffusivity combination on the recency effect. Bonferroni-corrected multiple comparisons are provided in Supplementary Tables S34 and S35.

In sum, although we observed stronger recency in some high- v and high- D conditions, the interaction between v and D did not yield a systematic, gradual pattern across the full parameter grid.

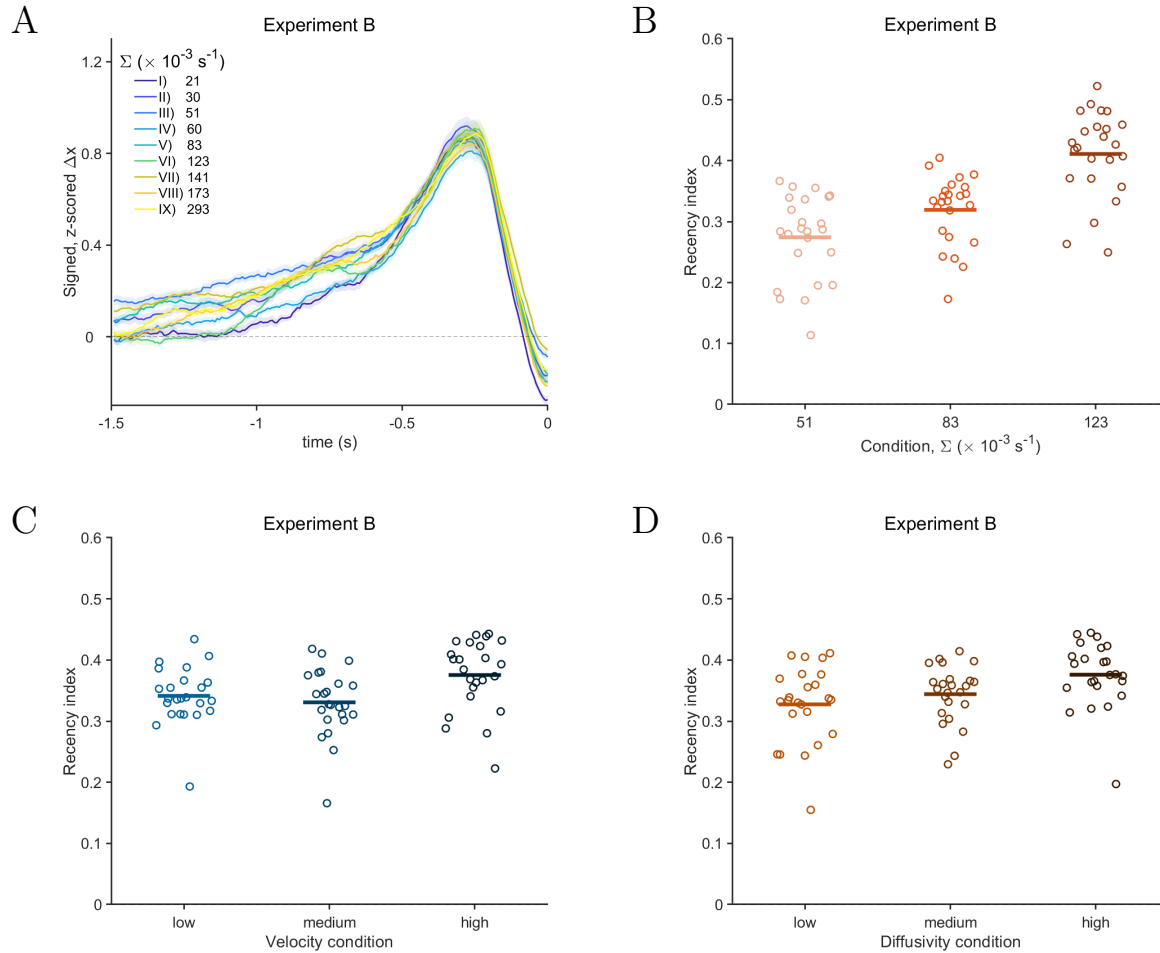


Figure 3.2: **Temporal weighting profiles across entropy production rates in Experiment B.** (A) Reverse correlation kernels averaged over participants for each entropy condition are shown. Kernels were computed as decision-signed averages of the standardized stimuli displacements across trials. X -axis refers to the last 1.5 seconds of the trials with 0 being the decision moment. (B) Recency indices for the three conditions in Experiment B, that are comparable with the Experiment A conditions, are shown. Individual recency indices for each participant were computed by subtracting the second and first half of the kernel means as an indicator of recency. Positive values indicate recency effect; while, negative values indicate primacy. The mean recency indices across participant for each condition is shown with a colored horizontal bar. We also compared recency indices across conditions grouped by velocity (C), and diffusivity (D).

Thus, in Experiment B the relationship between recency and the stimulus parameters appears less smooth than in Experiment A as can be seen in Fig. 3.2A.

Next, to further investigate the contrast in recency indices between Experiments A and B we focused on the three conditions in Experiment B that are comparable with the conditions in Experiment A. Note that these three conditions had a fixed velocity-to-diffusivity ratio with an increasing entropy production rate as in Experiment A. Fig. 3.2 B shows the recency indices of each participant in these three conditions. We tested these recency indices using a LME (see Section 3.5.1.1). Type III ANOVA on model estimates revealed a significant main effect of entropy production rate ($F_{2,48} = 56.05$, $p < 0.001$; see Supplementary Table S36). We found a gradual increase in recency with higher entropy (all $|t_{50.1}| > 3.35$, all $p < 0.005$). Detailed contrasts are provided in Supplementary Table S37.

3.3.2 Salient features in stimulus trajectories (jump analysis)

Next, we asked whether brief, coherent displacements of the stimulus in one direction—“jumps”—served as decision-relevant cues. Fig. 3.3 illustrates the procedure used to identify jumps in the stimulus trajectory. For each participant and each trajectory, we first smoothed the trajectory along the horizontal axis, and then scanned it with a moving segment until a decision was made. We marked a segment as a jump when it showed a sustained movement in one direction that lasted at least ~ 200 ms, and covered more than 150 pixels. For each detected jump segment, we fitted a line to obtain the following jump’s features: duration, distance, slope, and latest position. Finally, depending on its direction we labeled each jump as decision-congruent or incongruent. The congruency of the jump is a key feature for our subsequent analysis (see Section 3.5.1.2). We reasoned that if participants are influenced by these jump segments, then the properties of jumps that are congruent with the participant’s decision should differ systematically from those of incongruent jumps. This analysis complements the temporal weighting analysis by focusing on segments that provide clear directional evidence and testing whether participants leverage them, and whether this effect depends on the task structure (blocked design vs. intermixed design).

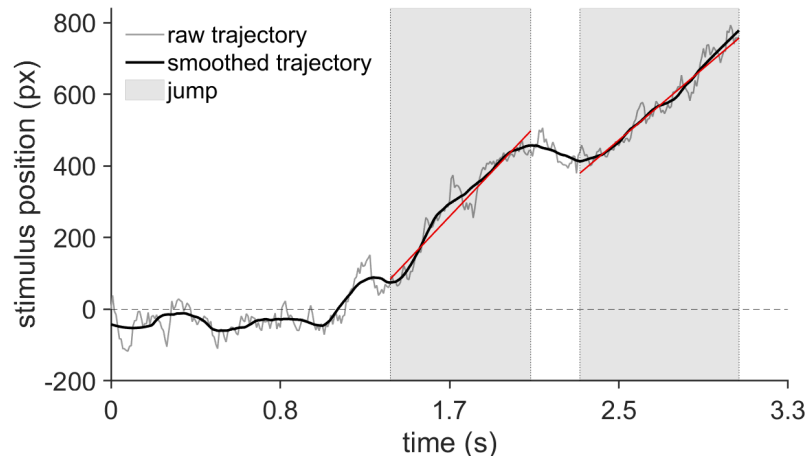


Figure 3.3: **Jump detection illustration of an example trajectory.** Trajectory until decision (gray line) is smoothed (black line), and it is scanned with a moving segment filter until decision time. Each segment that has a consistent derivative with a minimum length and duration is detected as jump. In this trajectory two segments are defined as decision-congruent jumps shown within the gray shaded area. As shown here, the jumps are brief consistent displacements in one direction (either congruent or incongruent with the decision direction). Then, we fitted a line to the detected jumps (red lines) to obtain features of the jumps, such as slope.

3.3.2.1 Experiment A

Fig. 3.4 shows box plots, for each of the identified feature of the jumps (i.e., duration, slope, length, position), the difference between congruent and incongruent segments. As expected, congruent jumps occurred more frequently overall than incongruent jumps (see Supplementary Fig. S12A).

The quantification of the differences between decision-congruent and decision-incongruent jumps across the four features (see Section 3.5.1.2) revealed that congruent jumps exceeded incongruent ones for all features (two-sided Wilcoxon rank-sum tests: all $U_C > 327/441$, all $p_{\text{Holm}} < .0077$; see Supplementary Table S38). Among these, the latest jump position showed the largest effect size (Cohen’s $d = 0.380$), whereas jump slope showed the smallest effect size (Cohen’s $d = 0.040$; see Supplementary Table S39). We confirmed these findings using a paired sign-flip permutation test on the mean difference between congruent and incongruent jumps (all $p_{\text{Holm}} < .0001$; see Supplementary Table S40).

We next wondered whether the influence of the jumps accumulates over a trial. Consistent with the recency effect, we expected later jumps to be more predictive of participants’ choices. To test this,

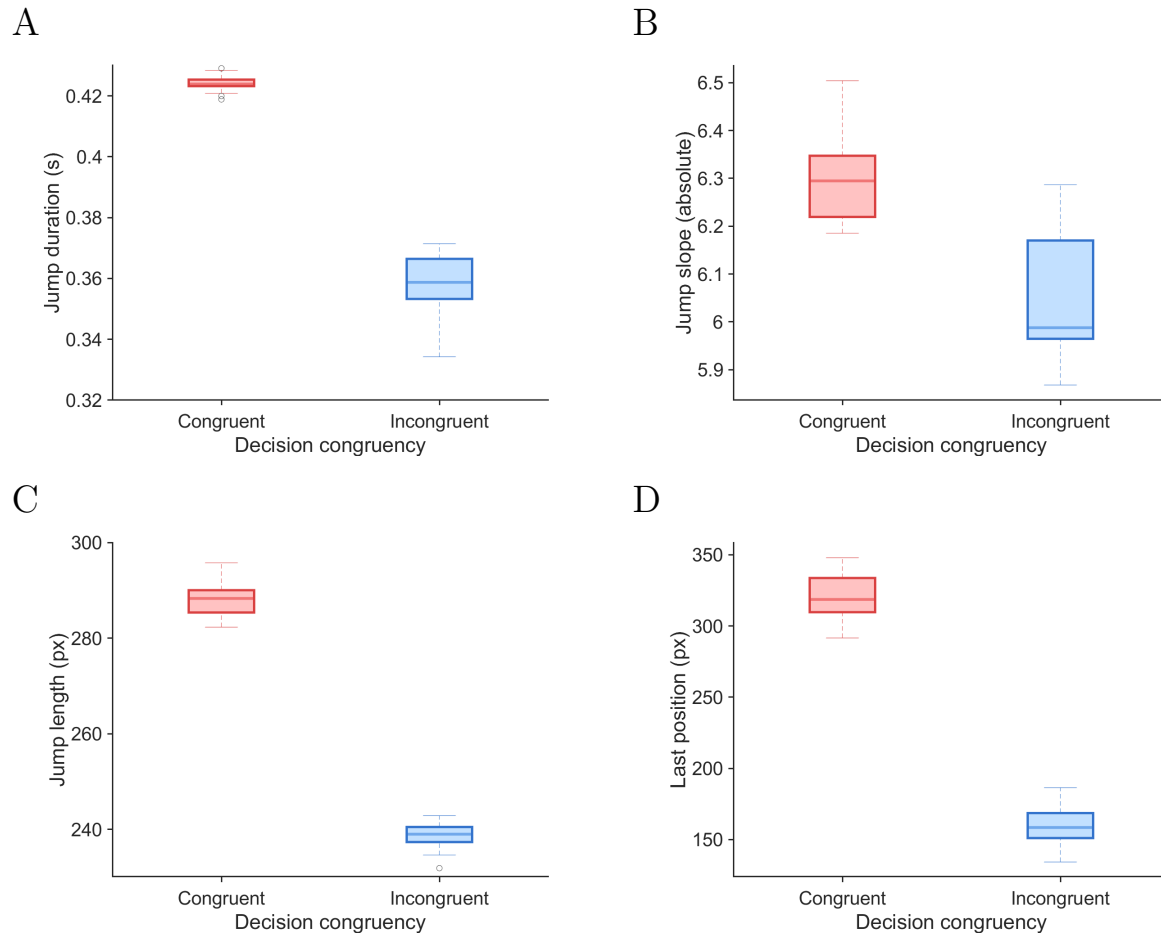


Figure 3.4: **Features of the jumps grouped by decision-congruency in Experiment A.** (A) The duration of the jumps, (B) the slope of the jumps, (C) the spatial length of the jumps, and (D) the last position that the jump arrives (as distance to the starting point of the disk) were divided into whether the jumps' direction is consistent with participants' response (decision-congruent jumps), or not (decision-incongruent jumps). Data points of the box plots were obtained by taking average of the relevant feature within each participant.

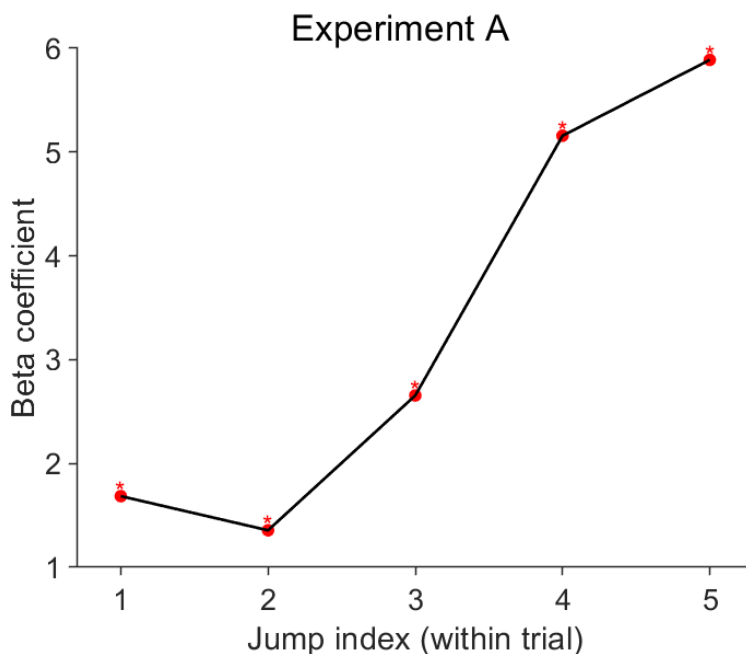


Figure 3.5: **Jumps are integrated in Experiment A.** In order to test whether the influence of jumps on decision accumulates over a trial we fitted a logistic regression that predicts the decision from each of the five jump’s direction. Jump index indicates the order of the jump occurring within a trial. All beta coefficients were positive and significant (shown with red data points and stars) and showed an increasing trend with jump index. This suggests that later jumps predicted choice more strongly, consistent with an integrative use of these cues as the trial progresses.

we focused on trials with the same number of detected jumps ($J = 5$), which allowed us to compare jump order directly (first through fifth). We then fitted a logistic regression to predict participant’s decision from the direction of each of the five jumps (see Section 3.5.1.2). As can be seen in Fig. 3.5, all five coefficients were positive and significant (all $|t_{785}| > 3.33$, all $p < 0.0009$), suggesting that all five jumps reliably biased participants’ choices in their direction (See Supplementary Table S41); the model was also significant overall ($F_{5,785} = 518.7$, $p < 0.0001$; Supplementary Table S42). Multiple comparisons revealed that most of the comparisons were significant (all $|t_{785}| > 3.82$, all $p < 0.0014$), except that the first jump did not differ from the second and third; neither the fourth from the fifth (Bonferroni-corrected multiple comparisons are provided in Supplementary Table S43). These results suggest that the later jumps predict choice more strongly, consistent with the recency effect .

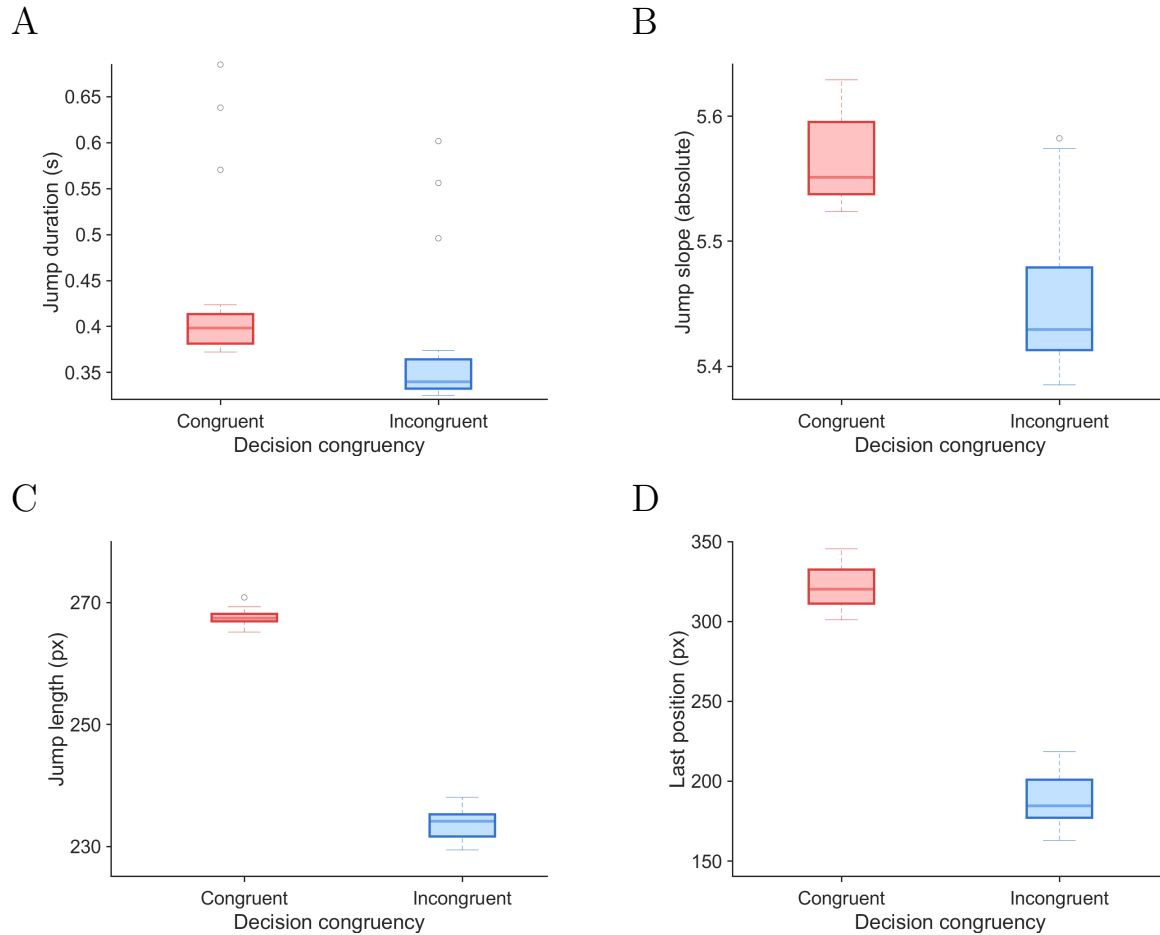


Figure 3.6: **Features of the jumps grouped by decision-congruency in Experiment B.** (A) The duration of the jumps, (B) the slope of the jumps, (C) the spatial length of the jumps, and (D) the last position that the jump arrives (as distance to the starting point of the disk) were divided into whether the jumps' direction is consistent with participants' response (decision-congruent jumps), or not (decision-incongruent jumps). Data points of the box plots were obtained by taking average of the relevant feature within each participant.

3.3.2.2 Experiment B

An identical set of analysis was performed for Experiment B. Fig. 3.6 shows box plots, for each of the identified feature of the jumps (i.e., duration, slope, length, position), the difference between congruent and incongruent segments in Experiment B. Similar to Experiment A, congruent jumps occurred more frequently overall than incongruent jumps (see Supplementary Fig. S12B).

The quantification of the differences between decision-congruent and decision-incongruent jumps across the four features revealed that congruent jumps exceeded incongruent ones in three features

(two-sided Wilcoxon rank-sum tests showing: all three $U_C > 437/576$, all $p_{\text{Holm}} < .0044$) except jump slopes which did not show a reliable difference between congruent and incongruent jumps (see Supplementary Table S44). We confirmed these findings using a paired sign-flip permutation test on the mean difference between congruent and incongruent jumps. This test showed that congruent jumps exceeded incongruent jumps in three features (all $p_{\text{Holm}} < .0001$; see Supplementary Table S45); while jump slopes were slightly lower (mean difference = -0.009) in congruent segments ($p_{\text{Holm}} = .02$). Among these features the latest jump position showed the largest effect size (Cohen's $d = 0.289$; see Supplementary Table S46). Additionally, across features, effect sizes were generally smaller in Experiment B than in Experiment A (see Supplementary Tables S39–S46).

We next wondered whether in Experiment B the influence of the jumps accumulates over a trial in a similar pattern to what we observed in Experiment A. To test this, we focused on trials with the same number of detected jumps ($J = 5$), which allowed us to compare jump order directly (first through fifth). We then fitted a logistic regression to predict participant's decision from the direction of each of the five jumps (see Section 3.5.1.2). As can be seen in Fig. 3.7, all five coefficients were positive and significant (all $|t_{2677}| > 12.5$, all $p < 0.0001$), suggesting that all five jumps reliably biased participants' choices in their direction (See Supplementary Table S47); the model was also significant overall ($F_{5,2677} = 1812.3$, $p < 0.0001$; Supplementary Table S48). Multiple comparisons revealed that many of the comparisons were significant (all $|t_{2677}| > 4.05$, all $p < 0.0005$), except that the first three jumps did not differ (Bonferroni-corrected multiple comparisons are provided in Supplementary Table S49). These results suggest that the later jumps predict choice more strongly, consistent with the recency effect.

3.4 Discussion

The aim of this chapter was to complement the work of Chapter 2 by further examining the aspects of the stimulus dynamics influencing participants' choices. Specifically, we investigated whether participants weighted different portions of the stimulus trajectory differently and whether abrupt changes in the trajectory influenced their decisions. To address these questions, we applied




Figure 3.7: **Jumps are integrated in Experiment B.** In order to test whether the influence of jumps on decision accumulates over a trial we fitted a logistic regression that predicts the decision from each of the five jump’s direction. Jump index indicates the order of the jump occurring within a trial. All beta coefficients were positive and significant (shown with red data points and stars) and show an increasing trend with jump index. This suggests that later jumps predict choice more strongly, consistent with an integrative use of these cues as the trial progresses.

two new analyses to Experiments A and B.

First, we asked whether observers relied more on early or late parts of the motion trajectory when forming their decision. To answer this, we computed temporal weighting profiles using reverse correlation. This analysis revealed a robust recency effect, indicating that late stimulus samples were weighted more heavily than early ones. In Experiment A (blocked design), recency progressively increased with the entropy production rate (dissipation level), whereas in Experiment B (intermixed design), recency remained present but its modulation across entropy production conditions was less pronounced.

Second, we quantified participants’ sensitivity to brief, salient changes in the stimulus trajectory (“jumps”). These jumps biased choices toward their direction, with later jumps exerting the strongest influence in both experiments. This effect was again more pronounced in Experiment A than in Experiment B. Importantly, not only the direction of the jumps mattered: their duration,

length, amplitude, and number also contributed to decision-making.

Taken together, these results show that observers did not integrate stimulus information uniformly across the trial. Instead, participants’ decisions were strongly influenced by the later portions of the motion trajectory and by sudden changes in its direction. Moreover, these effects were modulated by both stimulus dissipation (i.e., entropy production rate) and task structure (blocked vs. intermixed design)

The results of these new analyses can be interpreted in light of the Evidence Integration Model (EIM) introduced in Chapter 2. By positing a finite memory timescale τ_m , the model helps explain i) why recency emerges in both experiments, ii) why salient directional segments (“jumps”) would matter in both experiments, and iii) why these effects differ between experiments A and B. Below, we discuss these points separately.

First, the EIM inherently generates a recency effect: because accumulated evidence decays exponentially with time constant τ_m , more recent samples carry greater weight on the evolving decision variable. Second, the EIM predicts higher influence of salient changes in the stimulus trajectory on decision. When the stimulus trajectory is less noisy across time—as during a “jump”—successive samples accumulate faster within the finite memory window, driving the decision variable W_t toward the decision threshold more efficiently.

Third, to interpret the differences between the results of Experiment A and B, we focused on the task structure. In the blocked design (Experiment A), same entropy production rate repeats across many consecutive trials. This repetition may help participants to learn the stimulus statistics (i.e., the current dissipation level of the stimulus) and gradually tune their temporal integration policy to it (e.g., by setting an appropriate temporal integration window). Related recent works also show that when the environment is stable, observers can integrate information over a timescale that matches the statistics of the task [4], [81]–[83]. This view is consistent with our observation that recency and τ_m^* scale with entropy production rate in the blocked design but less so in the intermixed design. In Experiment B, the intermixed design may discourage or limit such tuning because the observer would need to infer the stimulus statistics quickly and re-adjust the integration policy on

each trial. This would place an extra demand on the observer. In fact, previous works suggest that performance is typically slower and/or less accurate in intermixed blocks, consistent with the higher cost of switching the strategy [41], [84], [85]. One possible consequence is that the perceptual system adopt a more generic compromise policy that is less tuned to the current dissipation level, which would attenuate entropy-dependent modulation of recency, as we observed. The same logic can also explain why jump usage is weaker in the intermixed design: when the task structure reduces the ability to tuning to the stimulus statistics, perceptual expectations and attentional allocation may be less well matched to the trial’s statistics, which can reduce the sensitivity to the stimulus features (i.e., jumps) [86], [87]. Importantly, this account is an interpretation of how the blocked vs intermixed design changes the demands of identifying the stimulus statistics and tuning the decision process on the specific stimulus condition; future work should clarify these points better.

In sum, the EIM not only accounts for the suboptimal performance relative to SPRT but also correctly predicts the recency effect and sensitivity to sudden salient changes in the stimulus trajectory. In the blocked design (Experiment A), the integration window τ_m expands with entropy production rate producing both stronger recency and higher sensitivity to stimulus trajectory changes. In the intermixed design (Experiment B), where the entropy production rate changes on a trial by trial basis, this adjustment diminishes. In sum, human observers construct decisions from stochastic, nonequilibrium dynamics using a temporally graded, feature-sensitive integration that depends on both the task structure and the dissipation level of the stimulus.

3.5 Materials and Methods

The analyses presented in this chapter use the same behavioral dataset obtained in Experiment A and B as described in the previous chapter (Chapter 2). To avoid redundancy, the full details regarding participant sample, stimuli, apparatus, and experimental design are provided in Section 2.5.

3.5.1 Data analysis

In this chapter, we analyzed the data of Experiments A and B using two approaches: i) reverse correlation analysis to identify participants’ temporal weighting profile, and ii) the identification and analysis of ”jumps” i.e., brief salient changes of the stimulus trajectory.

3.5.1.1 Temporal weighting profile

To assess how participants used stimulus information across time within each trial, we estimated temporal weighting profiles using a decision-locked reverse-correlation analysis. This method allowed us to infer the contribution of moment-to-moment fluctuations in the horizontal motion signal to the eventual choice, thereby providing an estimate of the reverse correlation kernel for each experiment.

Reverse correlation kernels. For each trial, we extracted the displacement of the moving stimulus in the decision-relevant (horizontal) axis, defined as Δx_t . We computed these displacements as running differences over a smoothing window of 0.2 s, converted into the corresponding number of frames given the screen refresh rate. To enable comparisons across entropy levels and participants, we standardized the displacement values by z -scoring within each participant \times entropy level:

$$z_t = \frac{\Delta_t - \mu}{\sigma}, \quad (3.1)$$

where μ and σ are the mean and standard deviation of the displacement series across all trials for that participant and entropy level. This procedure removes scale differences in variance introduced by entropy while preserving the temporal structure of fluctuations.

Kernels were computed as decision-signed averages of the standardized displacements: each trial’s z_t series was multiplied by the participant’s decision sign ($D_{\text{dec}} \in \{-1, +1\}$, coding leftward as -1 and rightward as $+1$), and then averaged across N trials:

$$K(\tau) = \frac{1}{N} \sum_{i=1}^N D_{\text{dec}}^{(i)} z_{\tau}^{(i)}, \quad (3.2)$$

where i indexes trials and τ denotes time relative to the decision.

We analyzed kernels aligned to the decision moment, spanning a fixed time window of 1.5 s prior to response. For each trial, we took the last 1.5 s of the standardized displacement series, reversed its time axis so that $\tau = 0$ corresponds to the decision, and averaged across trials. Trials that have a decision time below 1.5 second were skipped, which corresponds to 31.4% of the trials in Experiment A; and 11.1% of the trials in Experiment B.

Recency indices. To quantify the relative influence of late versus early evidence, we computed a recency index (RI) as the difference in mean kernel amplitude between the last and the first halves of the decision-locked window until the peak:

$$RI = \langle K_{\text{late}} \rangle - \langle K_{\text{early}} \rangle. \quad (3.3)$$

Positive RI values indicate recency (greater weight on later evidence), while negative values indicate primacy (greater weight on earlier evidence).

Linear mixed-effect models. Recency indices in Experiment A derived from individual reverse correlation kernels were analyzed using a linear mixed-effect model (LME) with a fixed effect of entropy production rate (Σ), and a random effect of participant (id) using the following formula (marginal R^2 : 0.30, conditional R^2 : 0.59):

$$RI_{\text{dec}} \sim \Sigma + (1 \mid \text{id}). \quad (3.4)$$

Then, to analyze Experiment B data we used two models with different fixed effects. First, we analyzed recency indices using a LME with fixed effects of velocity v and diffusivity D , and a

random effect of participant (id) using the following formula (marginal R^2 : 0.23, conditional R^2 : 0.56):

$$RI_{\text{dec}} \sim v * D + (1 | \text{id}). \quad (3.5)$$

Next, we restricted the Experiment B conditions to the three conditions that are comparable with the Experiment A (same velocity-to-diffusivity ratio), and analyzed these recency indices using a LME with a fixed effect of entropy production rate (Σ), and a random effect of participant (id) using the following formula (marginal R^2 : 0.44, conditional R^2 : 0.72):

$$RI_{\text{dec}} \sim \Sigma + (1 | \text{id}). \quad (3.6)$$

Common analysis procedures. All LME models were fitted with the `lme4` package [67]. LME model variance explained was computed using the `MuMIn` package [68]. The Satterwaite’s method [69] implemented in the `lmerTest` package[70] was used to estimate the degrees of freedom for the LME model ANOVA, and estimated marginal means were compared using the `emmeans` package[71] using Bonferroni adjustment.

3.5.1.2 Salient features in stimulus trajectories (jump analysis)

To test whether participants detect and use brief statistical cues in the stimulus trajectories, we performed a complementary analysis focusing on jumps. Jumps were defined as short temporal segments in which the disk’s velocity shows a consistent displacement in one direction, exceeding the fluctuations expected from background diffusion. We ran this analysis for both Experiment A and B.

Jumps. For each trial, we extracted the trajectory along the decision-relevant axis (Δx_t) until decision. We smoothed this trajectory using MATLAB’s `smooth` function with the Savitzky–Golay

option ('`sgolay`'), applying a smoothing window of 50 frames (corresponding to ~ 0.35 s at 144 Hz). We then scanned the smoothed trajectory with a moving segment filter until the decision time. A segment was classified as a jump if it satisfies three conditions: i) the derivative of the segment has a consistent sign (the disk does not change direction), ii) the segment length is at least 25 frames (≈ 200 ms), iii) the net displacement between start and end points of the segment is larger than 150 pixels. This procedure ensures that detected jumps reflect coherent, directional motion rather than random noise.

Features of jumps. To fit each jump, we used MATLAB's `fit` function with a first-order polynomial model (`poly1`), which returns the slope (p_1) and intercept (p_2):

$$\hat{y}(t) = p_1 t + p_2. \tag{3.7}$$

We retrieved four parameters from these fits as features of the jumps: The duration, spatial distance, slope and spatial position at the segment end (last position) of the jumps. All the features were recorded separately for decision congruent and incongruent segments.

Testing congruent vs. incongruent jump features. Once jumps were identified, we quantified their behavioral impact in the following analyses:

For each participant p , feature f (duration, spatial distance, slope and spatial position), and congruency (incongruent vs. congruent), we computed a participant-level summary by averaging the feature across all jumps within that cell (participant \times feature \times congruency). This produced two paired vectors for each feature: C (congruent) and I (incongruent).

For each feature, we formed paired differences:

$$d_{p,f} = C_{p,f} - I_{p,f}. \tag{3.8}$$

We tested whether the mean difference d was zero with a two-tailed paired sign-flip permutation test that is magnitude-sensitive. Under the null hypothesis, the signs of $d_{p,f}$ are exchangeable.

We generated a null distribution by randomly flipping signs (Monte Carlo with 100.000 iterations) and computed the mean d as statistic, as well as the p -value. We controlled the family-wise error rate across the four features within each experiment with the Holm step-down adjustment. As an effect size in original units, we reported the Hodges-Lehmann estimate with nonparametric 95% confidence intervals from a bootstrap over participants (5.000 resamples with replacement).

We also compared the congruent and incongruent distributions with a two-sided Wilcoxon rank-sum (Mann-Whitney U) test for each feature. Alongside the two-sided p -value, we reported the common-language effect size

$$\text{CL}_f = P(C > I) = \frac{U_C}{n_I n_C}, \quad (3.9)$$

where U_C is the Mann-Whitney statistic for the congruent group C against the incongruent group I (mid-ranks for ties), and n_I, n_C are the group sizes. Here, $P(C > I)$ means the probability that a randomly drawn observation from the congruent group exceeds a randomly drawn observation from the incongruent group. We quantified uncertainty in CL_f with a stratified bootstrap ($R = 5.000$): we resampled participants with replacement within each group, hold n_I, n_C fixed, and recomputed CL_f ; and took the 2.5% and 97.5% percentiles as the confidence interval. We applied the Holm adjustment across the four features to control the family-wise error rate.

For completeness, we also computed Cohen’s d as effect size of each feature on the same participant-level summaries as:

$$d = \frac{\bar{X}_{\text{cong}} - \bar{X}_{\text{incong}}}{s_p}, \quad (3.10)$$

$$s_p = \sqrt{\frac{(n_{\text{cong}} - 1) \text{var}(X_{\text{cong}}) + (n_{\text{incong}} - 1) \text{var}(X_{\text{incong}})}{n_{\text{cong}} + n_{\text{incong}} - 2}}. \quad (3.11)$$

Here \bar{X}_{cong} and \bar{X}_{incong} denote the sample means of the feature values for decision-congruent and decision-incongruent segments, respectively; $\text{var}(X_{\text{cong}})$ and $\text{var}(X_{\text{incong}})$ are the corresponding sample variances; n_{cong} and n_{incong} are the number of observations in each group; and s_p is the pooled standard deviation.

Jump integrator. Next, to test whether participants integrate jumps to guide decisions, we fitted logistic regression models using MATLAB’s `fitglm` function with a binomial distribution and logit link. Trials were restricted to those with exactly five detected jumps ($J = 5$), yielding $N = 791$ trials in Experiment A (6.6% of all trials), and $N = 2683$ trials in Experiment B (8.9% of all trials). For each trial, we constructed a feature vector $X = (x_1, x_2, \dots, x_J)$, where $x_j = 1$ if the j th jump was in the participant’s chosen direction, and $x_j = 0$ otherwise. The binary decision outcome $y \in \{0, 1\}$ was modeled as

$$\Pr(y = 1 \mid X) = \frac{1}{1 + \exp \left[- \left(\beta_0 + \sum_{j=1}^J \beta_j x_j \right) \right]}. \quad (3.12)$$

The regression coefficients β_j quantified the influence of each jump on the decision. A positive and significant β_j indicates that the corresponding jump reliably biases choices in its direction.

The fitted coefficients, their standard errors, and associated p -values were extracted directly from the `fitglm` output. We visualized these coefficients across jump order and highlighted statistically significant effects.

To assess whether the set of jump-index predictors jointly improved fit, we computed an ANOVA-style omnibus statistic on the model estimates and reported the F approximation. Pairwise differences between jump positions were tested as model-based coefficient contrasts using t -tests (estimate/SE with the GLM covariance), reporting $\Delta\beta$, SE, $t(\text{df})$, and Bonferroni-adjusted two-sided p -values across the 10 comparisons.

Chapter 4

Effects of Stimulus Format on Decision Making under Nonequilibrium Dynamics

Aybüke Durmaz¹, Gianfranco Fortunato¹, Edgar Roldán², Mathew E. Diamond¹, Domenica Bueti¹

1 International School for Advanced Studies (SISSA), Trieste, Italy

2 The Abdus Salam International Center for Theoretical Physics (ICTP), Trieste, Italy

4.1 Abstract

In this chapter, we present Experiment C, which examined whether human sensitivity to the thermodynamic arrow of time—the irreversibility embedded in nonequilibrium sensory fluctuations—generalizes across different presentation formats of identical stochastic motion. Using three types of stimulus formats (“linear accumulation”, “radial accumulation”, and “radial reset”), we tested how the spatial layout and the temporal continuity of the stimulus together shape perceptual decision-making under controlled entropy production levels. In linear accumulation, a drifting particle moved toward the left/right screen edges, whereas in the radial formats the same particle trajectory was mapped onto motion around a circle; in radial reset, the particle periodically returned to its starting position within the same trial. Participants judged the drift direction (i.e., left/right or clockwise/anticlockwise) of a Brownian particle whose motion carried identical physical evidence across these three formats. Behavioral analyses revealed that removing externally imposed spatial boundaries (i.e., the horizontal screen edges) in the radial formats increased decision thresholds, while introducing resets decreased performance accuracy and reduced sensitivity to entropy levels. Comparisons with the Sequential Probability Ratio Test (SPRT) indicated that decisions were closest to normative optimality in the linear accumulation format, less so in radial accumulation,

and least in radial reset. Reverse correlation and jump analyses further revealed that temporal weighting (recency vs. primacy) and sensitivity to brief, salient motion segments in the trajectory (“jumps”) depended strongly on stimulus format. In linear accumulation, observers exhibited pronounced recency and showed sensitivity to salient motion segments in the stimulus trajectory. In radial accumulation and radial reset, both of these effects were attenuated: recency was reduced, and jump sensitivity was not reliably observed. Together, these results demonstrate that participants’ sensitivity to the nonequilibrium structure was shaped not only by stimulus physics but also by the spatial and temporal layout of the stimulus evidence, revealing format-dependent constraints on how the brain converts physical irreversibility into perceptual evidence.

4.2 Introduction

Real-world sensory environments are rarely static or reversible. From the gait of a walking person to the swirling of clouds in the atmosphere, the signals we encounter are generated by stochastic processes that unfold irreversibly in time. Such time-asymmetric signatures are formally captured by entropy production in stochastic thermodynamics, which quantifies the arrow of time in fluctuating systems [88]. In these nonequilibrium contexts (i.e., systems that continuously exchange energy with their surroundings rather than settling into time-reversible equilibrium) the temporal statistics of sensory inputs inherently reflect ongoing dissipation. For perception, this raises a fundamental question: do humans process these irreversibility cues when forming judgments about noisy sensory inputs?

Perceptual decision-making research has classically examined how observers integrate noisy evidence using tasks such as the random dot motion paradigm [2], [22]. Sequential sampling models such as the Drift Diffusion Model (DDM) or the Sequential Probability Ratio Test (SPRT) [23], [54] describe this process as a gradual accumulation of evidence toward a decision threshold, explaining key behavioral metrics such as decision times, accuracies, and threshold settings. While powerful, these models typically abstract away the physical structure of the sensory input, using only the average noise of the stimulus and disregard how the noise evolves in time.

Recent work presented in Chapter 2 bridges this gap by linking perceptual decision-making to the physics of nonequilibrium fluctuations. In Experiments A and B, we showed that participants were sensitive to entropy production in stochastic motion trajectories generated by Brownian dynamics with drift. More specifically, higher entropy production led to shorter decision times. Suboptimal evidence integration suggested by SPRT led us to develop a model with non-Markovian evidence accumulation and a finite memory time constant (temporal integration window). Variation in this temporal integration window as a function of entropy production rate suggested that observers adjust their integration window to the rate of dissipation, favoring temporally extended memory over instantaneous evidence to support effective decisions in nonequilibrium contexts. These findings demonstrate that the human perceptual system is sensitive to the thermodynamic arrow of time when evaluating noisy motion [13], [89].

A critical next step is to test the generalizability of this sensitivity. In Experiments A and B, all evidence was presented in a linear accumulation format, where a drifting particle moved horizontally until reaching the screen edges. This setup provided positional cues and an externally imposed bound on decision time, as the likelihood of the true stimulus direction (i.e., right/left) was higher when it is closer to the corresponding screen edge. However, the same physical evidence can be reformatted in ways that change its stimulus display layout without altering the underlying stochastic dynamics. In Experiment C, we therefore compared three presentation formats: i) linear accumulation, where the particle drifts horizontally until it hits the right or left screen edges; ii) radial accumulation, where the same displacements are mapped onto angular motion along a circle, reducing positional cues and removing spatial bounds; and iii) radial reset, which further introduces periodic resets that disrupt temporal continuity and add memory demands. By contrasting these conditions, we asked whether sensitivity to nonequilibrium signatures is invariant across formats, or whether it is shaped by the spatiotemporal layout of the stimulus.

This chapter is organized into four parts. First, we analyzed decision thresholds, accuracies, and decision times to characterize how stimulus format and entropy production jointly shape perceptual judgments. Second, we compared these behaviors to the predictions of the SPRT model, quantifying

deviations as measures of suboptimality. Third, we estimated temporal weighting profiles using reverse correlation, probing whether participants place greater weight on early or late stimulus fluctuations. Finally, we tested whether participants detect and use brief, salient motion segments (i.e., “jumps”) emerging within trajectories, and whether such reliance depends on stimulus format. Together, these analyses provide a comprehensive picture of how nonequilibrium sensitivity interacts with stimulus presentation format to shape perceptual decision-making.

4.3 Results

In Experiment C, we investigated how the format in which evidence is presented shapes perceptual decision-making, beyond the physical properties of the stimulus itself. Participants judged the motion direction of a disk (a driven Brownian particle described by Langevin equation, see Chapter 2) under three stimulus format conditions: linear accumulation, radial accumulation, and radial reset (see Section 4.5.2). The stimulus format manipulation was blocked rather than inter-mixed: participants completed the task over two sessions and within each session they performed consecutive blocks of one format before switching to the next. Format order was counterbalanced across participants, and blocks for each format were divided equally between the two sessions (see Section 4.5.3). Importantly, the underlying stimulus trajectories were identical across all three conditions: the decision-related displacements (i.e., horizontal displacements) in the linear accumulation condition were re-mapped as arc length displacements along a circular path in the radial conditions. Fig. 4.1 shows an example stimulus trajectory for all three stimulus formats.

In the linear accumulation condition, participants decided whether the disk moves leftward or rightward before it reaches the screen edges. This format was the same as in Experiments A and B (see Chapter 2). Because the screen edges impose an external decision boundary, this format tends to encourage narrower thresholds and faster decisions. Additionally, the position of the disk relative to the edges already conveys some directional evidence: near the left or right edge, the likelihood of the direction is higher, giving participants clearer cues. In the radial accumulation condition, participants judged whether the disk rotates clockwise or anticlockwise. Here, no boundary con-

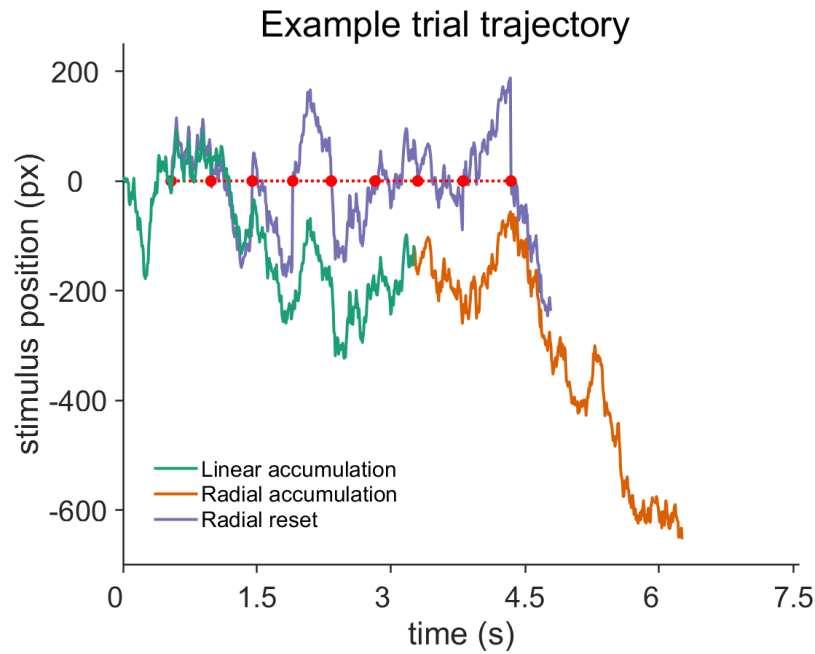


Figure 4.1: **Decision-relevant displacement of an example trajectory as a function of time across all three stimulus formats.** The figure shows an example trial trajectory until decision time. The decision-relevant displacement time series, defined as Δx_t in the linear accumulation condition and Δs_t (arc-length displacement) in the radial conditions were identical across three stimulus format conditions. Positive values in the vertical axis refers to rightward/clockwise motion direction. Due to presence of an external boundary (vertical screen edges), the linear accumulation trajectory ends earlier compared to the radial conditions. The trajectory of the radial reset condition is interrupted by resets, shown as red dots, every 0.4-0.5 seconds.

strains decision thresholds, and at every angular position the disk could plausibly belong to either direction, resulting in weaker spatial cues. In addition to this, in the radial reset condition, the trajectory resets to the initial position repeatedly within a trial, removing the possibility of relying on continuous accumulation and adding an additional working-memory demand.

In addition to the format in which evidence was presented, a central variable of interest in this work is the entropy production rate of the stimulus trajectories. In stochastic thermodynamics, entropy production quantifies the degree of nonequilibrium irreversibility in a dynamical system: the extent to which its fluctuations unfold in a preferred temporal direction. For perception, this means that trajectories with higher entropy production contain stronger directional signatures thereby providing stronger cues about the underlying drift. By contrast, trajectories with lower entropy production are closer to equilibrium, with weaker directional signatures. In Experiments A and B, we demonstrated that participants' decisions were systematically tuned to these nonequilibrium properties, suggesting that human observers are sensitive to nonequilibrium fluctuations when extracting sensory evidence from noisy motion [89]. In Experiment C, we extended this logic to test whether the “format” of evidence presentation interacts with entropy production rate, shaping the thresholds, times, and accuracies of perceptual decisions. In each format, we manipulated three levels of entropy production rate while keeping the velocity-to-diffusivity ratio fixed, using an intermixed design in which entropy levels were presented in random order within each block (see Section 4.5.3).

To this end, we analyzed three decision metrics: decision thresholds (X_{dec}), accuracies (α), and decision times (T_{dec}). We used linear mixed-effects (LME) models to assess how these metrics vary across stimulus formats and entropy production levels.

4.3.1 Decision thresholds: Participants adopted similar threshold levels in radial settings with a more pronounced effect of entropy

First, we focused on participants' mean decision thresholds $\langle X_{\text{dec}} \rangle$ using a LME (see Section 4.5.4.1; and Fig. 4.2). Type III ANOVA on model estimates revealed significant main effects of stimulus

format ($F_{2,176} = 153.04, p < 0.001$); and entropy level ($F_{2,176} = 128.64, p < 0.001$), as well as their interaction ($F_{4,176} = 16.57, p < 0.001$). ANOVA results are provided in Supplementary Table S51,

Post-hoc multiple comparisons revealed that at every entropy level, decision thresholds were significantly narrower in the linear condition than in both radial accumulation and radial reset (all $|t_{184.38}| > 3.12, p < .01$); whereas radial accumulation and radial reset did not differ. Crucially, entropy modulated decision thresholds in the linear accumulation (e.g., a modest increase from the lowest to highest entropy; $t_{184.38} = -2.70, p = .023$), but it produced even stronger monotonic increases in both radial accumulation and radial reset (all $|t_{184.38}| > 5.5, p < .0001$). Notably, the largest entropy-driven increase was observed in radial reset. Bonferroni-corrected multiple comparisons are provided in Supplementary Tables S52, and S53.

Together, these results demonstrate mainly that i) as expected decision thresholds were larger in radial conditions compared to linear condition since there was no boundary pressure in radial conditions; ii) interestingly participants kept similar decision thresholds between radial accumulation and radial reset conditions suggesting a similar decision boundary setting even if there is a higher demand on memory in the reset condition; and iii) the entropy effect on decision thresholds is strongly dependent on the stimulus presentation format: minimal in linear accumulation but pronounced in radial and radial reset.

4.3.2 Accuracy: Stimulus format explains accuracy differences, while resetting reliably lowers accuracy

Next, we focused on decision accuracy (α), to examine whether stimulus format and entropy also affect the overall quality of perceptual judgments, using a LME (see Section 4.5.4.1; and Fig. 4.3). Type III ANOVA on model estimates revealed significant main effects of stimulus format ($F_{2,176} = 25.70, p < 0.001$) and entropy ($F_{2,176} = 3.67, p = 0.027$), as well as their interaction ($F_{4,176} = 3.23, p = 0.014$). The effect size estimates indicate that stimulus format ($\eta^2 = 0.23$) is the dominant determinant of accuracy, whereas entropy ($\eta^2 = 0.04$) has only a weak effect on its own. Full

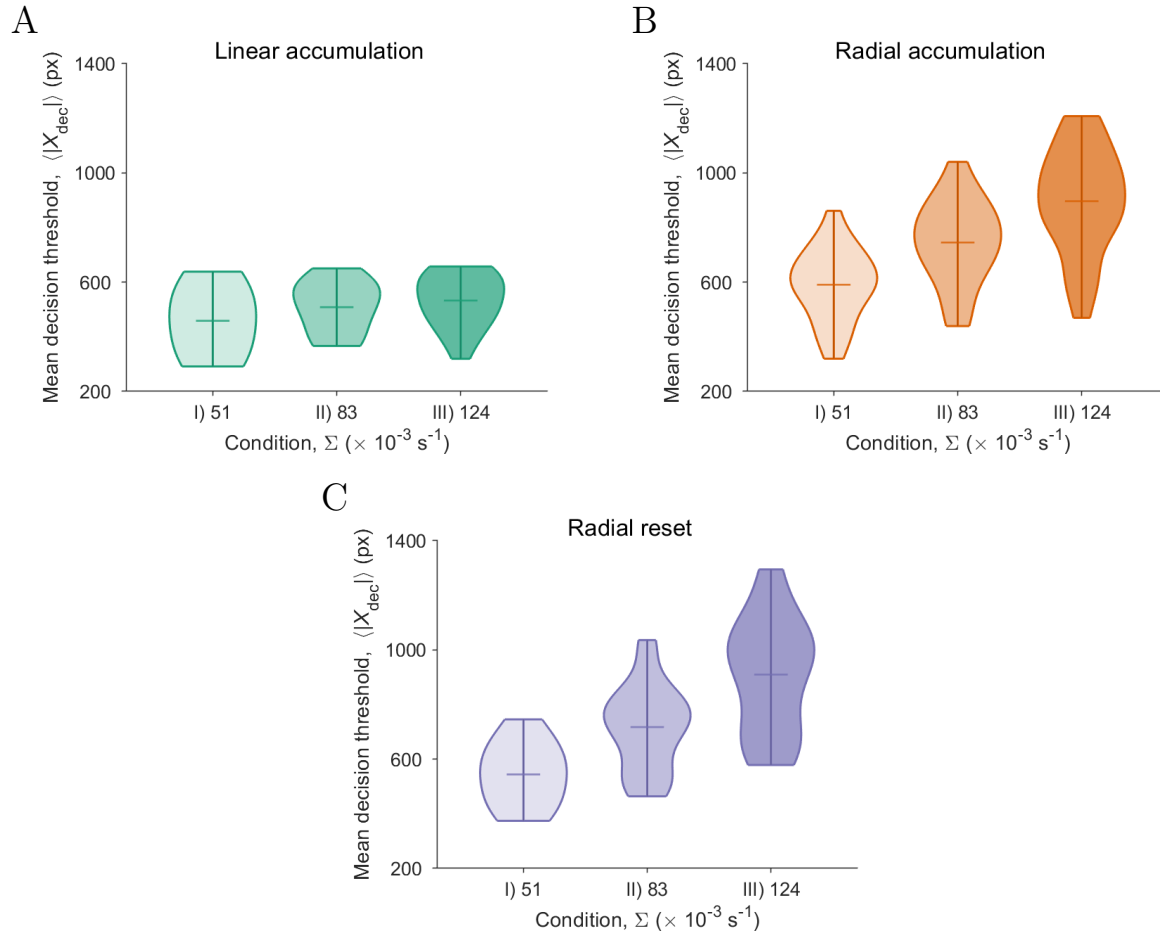


Figure 4.2: **Experiment C, statistics of the participants' mean decision-thresholds $\langle |X_{\text{dec}}| \rangle$ as a function entropy production rate Σ levels.** (A) Participant-by-participant $\langle |X_{\text{dec}}| \rangle$ distributions in linear accumulation condition, (B) the same metric in radial accumulation, and (C) in radial reset condition. Data points in each violin plot were obtained by averaging over all trials for every participant. Central horizontal lines show the respective means.

ANOVA results are provided in Supplementary Table S54.

Post-hoc multiple comparisons revealed no evidence for a systematic relationship between entropy level Σ and accuracy when comparisons were grouped by stimulus format. In particular, none of the entropy contrasts were significant in either radial accumulation or radial reset (all $|t_{184.38}| < 2.16$, $p > .09$). In the linear accumulation condition, accuracy dropped only from low to medium entropy ($\Delta = 0.045$, $t_{184.38} = 3.12$, $p = .006$), with no further reliable changes. Moreover, radial accumulation had higher accuracy than radial reset at all entropy levels (all $|t_{184.38}| > 3.53$, $p < .002$). Bonferroni-corrected multiple comparisons are provided in Supplementary Tables S55, and S56.

Together, these results indicated that i) overall participants' accuracy was robust to the entropy changes; and ii) the accuracy dropped in radial reset condition compared to radial accumulation even if they had similar threshold settings.

4.3.3 Decision Times: Sensitivity to entropy diminishes across stimulus formats

Finally, we focused on participants' mean decision times $\langle T_{\text{dec}} \rangle$ using a LME (see Section 4.5.4.1; and Fig. 4.4). Type III ANOVA on model estimates revealed significant main effects of stimulus format ($F_{2,175} = 251.78$, $p < 0.001$), and entropy ($F_{2,175} = 50.85$, $p < 0.001$), as well as an interaction ($F_{4,175} = 8.45$, $p < 0.001$). Full details on ANOVA is provided in Supplementary Table S57.

Post-hoc multiple comparisons revealed that the difference between linear and the radial conditions was robust at all entropy levels (all $|t_{183.38}| > 7.12$, $p < .0001$), with no reliable differences between radial accumulation and radial reset. Furthermore, in linear accumulation, decision times were faster progressively with higher entropy (all $|t_{183.38}| > 4.6$, $p < .0001$). In the radial accumulation condition, entropy effects were weaker but still significant (all $|t_{183.38}| > 2.77$, $p < .019$), except low vs. medium). In the radial reset condition, entropy effects were minimal with only the low vs. high contrast reaching significance (faster decision times at high entropy: $t_{183.38} = 2.59$, $p = .031$); while

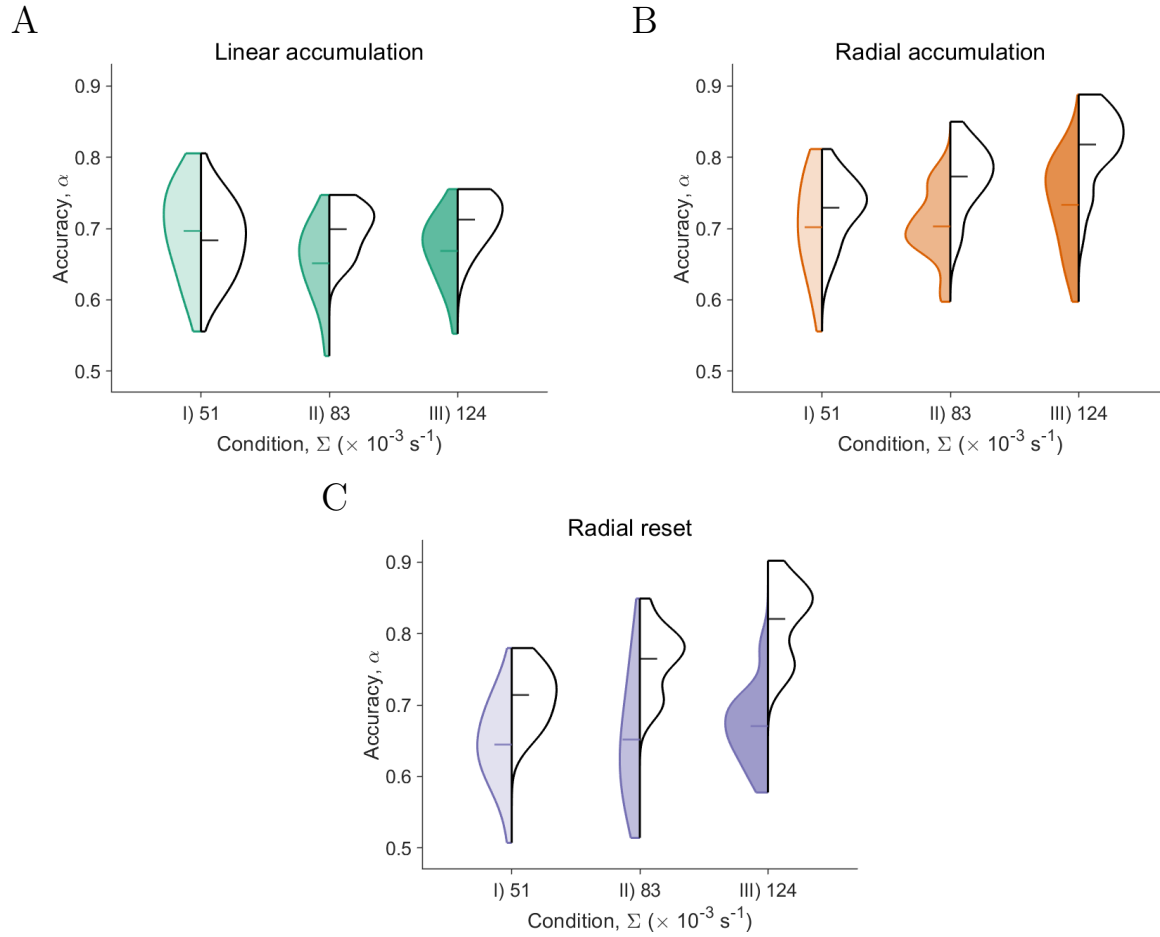


Figure 4.3: **Experiment C, statistics of the participants' accuracy α as a function of entropy production rate Σ levels.** (A) Participant-by-participant accuracy $P(D_{\text{dec}} = H)$ distributions in linear accumulation condition, (B) the same metric in radial accumulation, and (C) in radial reset condition. The colored violin plots refer to the experimental data; whereas, the empty violin plots illustrate the distribution of the theoretical predictions from SPRT for each participant. SPRT predictions of accuracy were obtained by plugging Péclet number and entropy production rate in Eq. (2.10). Central horizontal lines show the respective means.

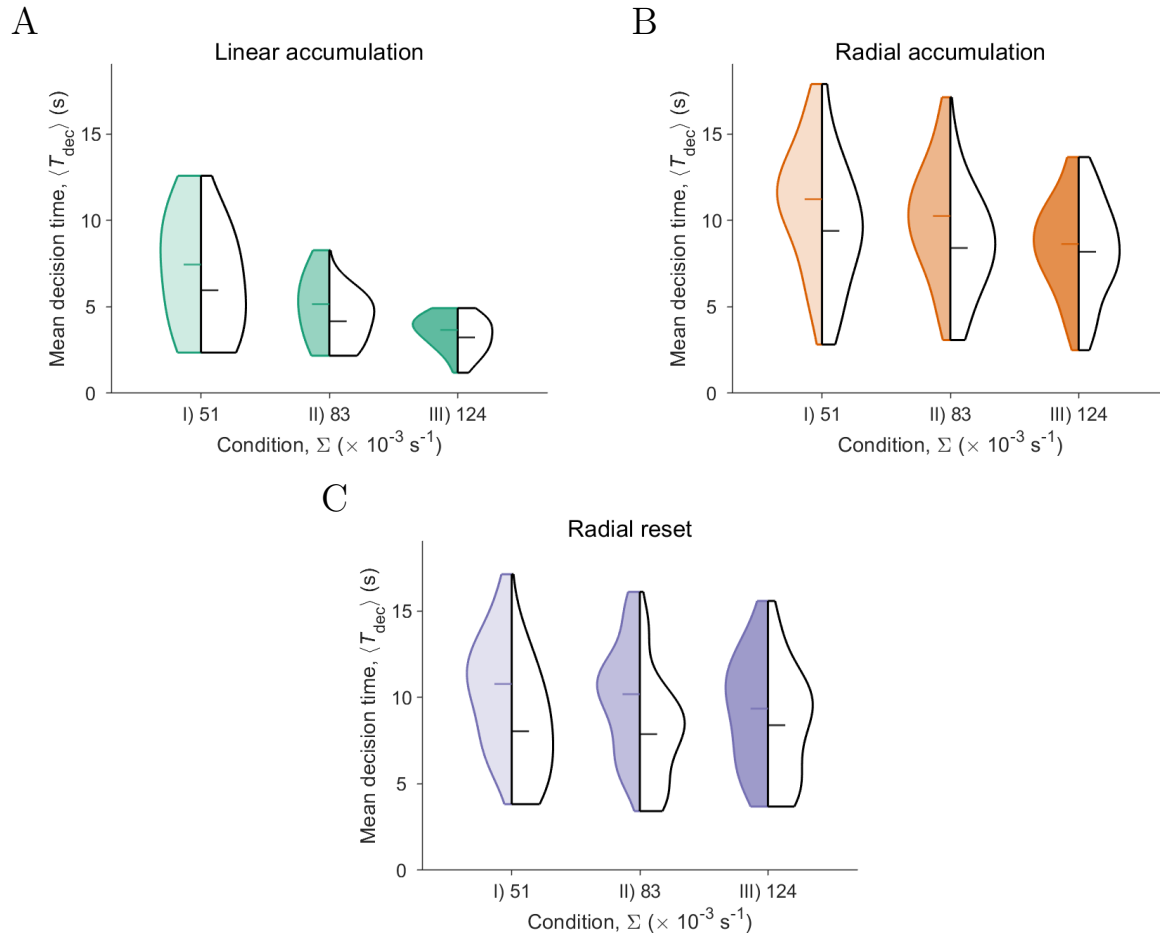


Figure 4.4: **Experiment C, statistics of the participants' mean decision time $\langle T_{\text{dec}} \rangle$ as a function of entropy production rate Σ levels.** (A) Participant-by-participant mean decision time $\langle T_{\text{dec}} \rangle$ distributions in linear accumulation condition, (B) the same metric in radial accumulation, and (C) in radial reset condition. The colored violin plots refer to the experimental data; whereas, the empty violin plots illustrate the distribution of the theoretical predictions from SPRT for each participant. SPRT predictions of mean decision times were obtained by plugging Péclet number and entropy production rate in Eq. (2.9). Central horizontal lines show the respective means.

the rest was not significant. This suggests a gradual diminishing of the modulation of decision times with entropy across stimulus formats. Bonferroni-corrected multiple comparisons are provided in Supplementary Tables S58, and S59.

Together, these results demonstrate that i) decision times were overall fastest in the linear condition and slowest in the radial formats with no difference between radial accumulation and reset conditions; and ii) entropy led to faster decisions across all formats; and this effect was strongest in linear accumulation, weaker in radial accumulation, and minimal in radial reset.

4.3.4 Summary across decision metrics

In sum, presenting the stimulus in a radial format without externally imposed boundaries altered the dynamics of perceptual decision-making compared to linear accumulation. In both radial accumulation and radial reset formats, participants adopted substantially higher thresholds than in linear accumulation, with stronger entropy-related modulation of thresholds. Despite the additional memory demands introduced by resets, thresholds remained comparable between radial accumulation and radial reset, indicating a shared decision threshold setting across the two radial formats. Accuracy, however, diverged: while accuracy showed minimal modulation by entropy across all three stimulus formats, it was consistently higher in radial accumulation than in radial reset. One possible interpretation is that resets reduced effective evidence quality by fragmenting temporal continuity. Finally, decision times showed a graded sensitivity to entropy: entropy led to faster decisions most strongly in linear accumulation, less so in radial accumulation, and only minimally in radial reset. In the next section, we investigate the implications of these behavioral results in terms of decision optimality.

4.3.5 Comparison with an optimal decision-maker

The Sequential Probability Ratio Test (SPRT) provides a normative benchmark for perceptual decision-making by specifying how an ideal observer accumulates evidence to discriminate between competing hypotheses [21], [23]. Importantly, SPRT does not generate predictions for decision thresholds X_{dec} themselves, but instead links thresholds and accuracies: it predicts decision accuracy for a given threshold (Eq. (2.10)) and the expected decision time corresponding to that accuracy (Eq. (2.9); Section 4.5.4.2). Thus, when comparing human behavior to SPRT, we focused on accuracies and decision times.

To quantify deviations from optimality, we computed the absolute error (AE) between observed performance and SPRT predictions (see Section 4.5.4.3). These absolute errors served as a measure of suboptimality: smaller values indicate closer alignment with the normative model, whereas larger values reflect systematic departures from optimal decision-making.

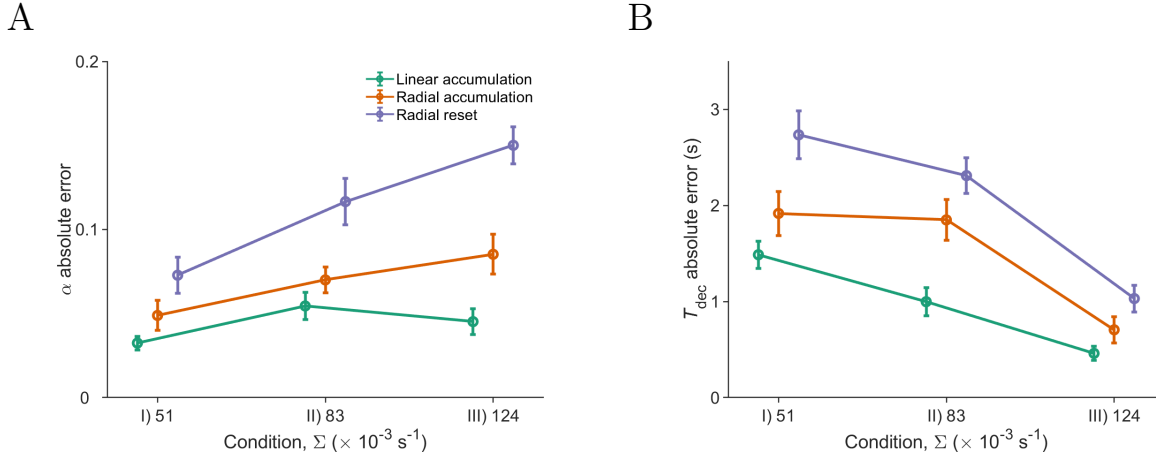


Figure 4.5: **Absolute error (AE) between observed performance and SPRT predictions.** Participant-by-participant absolute errors between (A) observed accuracy and SPRT predictions of accuracy as a function of entropy production rate levels for each stimulus format, and (B) the same metric for decision times, T_{dec} . The data points represent mean AE of all participants, whereas error bars represent SEM.

4.3.5.1 Accuracies α

For accuracies, absolute errors (AE_{α}) were tested using a LME (see Section 4.5.4.3, and Fig. 4.5A). Type III ANOVA on model estimates revealed significant main effects of stimulus format ($F_{2,176} = 46.02$, $p < .0001$) and entropy ($F_{2,176} = 17.38$, $p < .0001$), as well as an interaction ($F_{4,176} = 3.51$, $p = .0087$). ANOVA results are provided in Supplementary Table S60.

Post-hoc multiple comparisons showed that observed accuracy was closest to SPRT predictions in linear accumulation, and the deviation between observed and SPRT-predicted accuracy increased in radial accumulation and even further in radial reset (all $|t_{184}| > 3.2$, $p = .0048$; see Supplementary Tables S61 and S62).

The interaction showed that entropy effects differed by stimulus format. In linear accumulation, absolute errors remained uniformly low with no significant difference across entropy levels. In radial accumulation, absolute errors increased moderately, with a significant increase from low to high entropy ($t_{184.38} = -2.81$, $p = .017$). In radial reset, absolute errors increased sharply and monotonically across all entropy contrasts (all $|t_{184.38}| > 2.584$, $p < .03$). Bonferroni-corrected multiple comparisons are provided in Supplementary Tables S63, and S64.

Together, these results demonstrate that participants' accuracies approximated SPRT predictions most closely in the linear accumulation condition, less so in radial accumulation, and least in radial reset. Moreover, entropy selectively increased deviations from optimality in radial formats, with the strongest divergence observed in radial reset.

4.3.5.2 Decision times T_{dec}

We next tested whether participants' decision times align with SPRT predictions by analyzing absolute errors in mean decision times ($AE_{\langle T_{\text{dec}} \rangle}$), using a LME (see Section 4.5.4.3; and Fig. 4.5B). Type III ANOVA on model estimates revealed significant main effects of stimulus format ($F_{2,176} = 33.34$, $p < .0001$) and entropy level ($F_{2,176} = 69.60$, $p < .0001$), with no significant interaction. ANOVA results are provided in Supplementary Table S65.

Post-hoc comparisons showed that the difference between the observed and SPRT-predicted decision times increased from linear accumulation to radial reset gradually (all $|t_{184.38}| > 3.98$, $p < .0003$). Moreover, it decreased with increasing entropy production (all $|t_{184.38}| > 8.64$, $p < .0001$; except low vs. medium entropy was only reaching significance with $p = .069$). Bonferroni-corrected multiple comparisons are provided in Supplementary Tables S66, and S67.

Together, these results showed that deviations from SPRT-predicted decision times were smallest when the stimulus was presented linearly and increased under radial formats, particularly under radial reset. At the same time, increasing entropy production rates substantially reduced the deviations across all formats.

4.3.5.3 Linking back to behavioral results

The comparison with SPRT clarifies how behavioral deviations from optimality relate to the patterns observed in raw decision metrics. Suboptimality was smallest in the linear accumulation condition and grew in radial formats, especially in radial reset. This matches the behavioral finding that externally imposed bounds in linear accumulation format may promote lower thresholds and faster decisions. Another possible interpretation is that in linear accumulation, the mapping

from the stimulus display to the decision variable was more direct and more reliable compared to the radial conditions due to less memory demands, and this could be the underlying factor explaining why suboptimality was smallest in linear accumulation.

The role of entropy, however, requires careful interpretation. The behavioral analyses showed that higher entropy production accelerated decision times across all formats and increased thresholds in the radial conditions, consistent with the fact that trajectories providing stronger directional cues under high entropy. These changes in the underlying distributions may partly explain why absolute errors relative to SPRT across entropy conditions diverged for accuracy but converged for decision times. For accuracies, higher entropy produced more variability in decision thresholds, which propagated into greater across-trial variance and thus larger deviations from predicted performance. For decision times, the faster and less variable response distributions observed at high entropy may have yielded smaller errors relative to SPRT predictions. In this way, the entropy dependence of suboptimality may mirror the statistical structure of the raw behavioral data.

Taken together, these results showed that stimulus format was the dominant determinant of whether human decision-making approximates optimality: behavior approached closest to the SPRT benchmark in the linear accumulation condition, while the largest deviations was produced under radial reset condition. Entropy production, in turn, modulated suboptimality in distinct ways for accuracy and decision time, which may be due to the impact of decision threshold and decision time distributions.

4.3.6 A closer look into stimulus dynamics influencing decision

The analyses of decision thresholds, accuracies, decision times, and their deviations from SPRT predictions showed that participants' decision strategies vary substantially across stimulus formats, even though the underlying physical motion is identical. However, they did not yet reveal how participants process the ongoing stimulus information within a trial to arrive at their choices.

To address this gap, we turned to two complementary approaches that allow us to probe the fine-grained temporal dynamics of evidence integration. First, we estimated temporal weighting

profiles using reverse correlation, which quantify how different moments within a stimulus trajectory contribute to the final decision. This analysis has been widely used to distinguish between primacy- and recency-weighted strategies, and it provides a direct test of whether participants integrate evidence uniformly or whether they exhibit time-dependent weighting [78].

Second, we examined participants' sensitivity to salient motion segments embedded in the trajectories, focusing on brief but consistent motion towards a direction, or "jumps". Prior work on Experiments A and B showed that participants can exploit such cues to guide their decisions (see Chapter 3). Here, we tested whether the reliance on these jumps varies when trajectories are presented linearly, radially, or with resets.

Together, these complementary analyses provide a more mechanistic picture of how evidence was temporally weighted and how specific trajectory features were used during decision-making. By combining them with the previous findings, we aimed to reveal not only how well participants approximated optimality, but also how they constructed decisions from stochastic, nonequilibrium dynamics.

4.3.6.1 Temporal weighting profile

To characterize how participants weighted stimulus information over time within each trial, we estimated temporal weighting profiles using a reverse-correlation analysis (see Section 4.5.4.4). For each participant, we computed a reverse-correlation kernel that captured the contribution of each moment in the stimulus trajectory to the final decision. To do so, for each trial, we extracted the decision-relevant displacement of the stimulus, defined as the horizontal displacement Δx_t in the linear accumulation condition and the arc-length displacement Δs_t in the radial conditions. To enable comparison across entropy levels, we z-scored these displacements within each participant \times stimulus format \times entropy level. We then aligned the resulting standardized displacements to the decision, multiplied them by the signed choice (i.e., ± 1 for right/left or clockwise/anticlockwise), and averaged across trials. This procedure yielded decision-locked reverse correlation kernels for each participant, reflecting their temporal weighting profile. Fig. 4.6A shows these kernels for

each stimulus format, averaged across participants for visualization. From each kernel, we then derived a recency index (RI), which summarizes whether choices were driven more strongly by early (primacy) or late (recency) portions of the trajectory. To compute this metric, we split each kernel into two equal temporal halves and calculated the difference between their mean weights. Positive RI values indicate recency, reflecting stronger weighting of late evidence, whereas negative values indicate primacy, reflecting stronger weighting of early evidence. Fig. 4.6B shows the RI values derived from each participant’s kernel across stimulus formats.

Finally, we tested participants’ temporal weighting profiles, quantified by decision-locked recency indices (RI_{dec}), using a LME model (see Section 4.5.4.4). Type III ANOVA on model estimates revealed a significant main effect of stimulus format ($F_{2,176} = 43.46, p < 0.001$). In contrast, entropy level did not yield a significant effect, and we likewise did not find a significant interaction between stimulus format and entropy. ANOVA results are provided in Supplementary Table S68.

Post-hoc multiple comparisons revealed that, recency indices were significantly higher in the linear accumulation condition compared to both radial accumulation and radial reset (all $|t_{184.38}| > 7.18, p < .0001$), while radial accumulation and radial reset did not differ significantly. Bonferroni-corrected multiple comparisons are provided in Supplementary Table S69.

These results suggested that participants’ temporal weighting strategy depended strongly on stimulus format. In the linear accumulation condition, participants exhibited a more pronounced recency effect, meaning that they placed greater weight on late evidence. In contrast, temporal weighting was reduced in both radial accumulation and radial reset.

Taken together, one possible interpretation of these results is that the stronger recency effect observed in the linear format may reflect the external boundary constraints imposed by the task: as the disk was approaching the screen edge, each new displacement provided increasingly more directional information and participants faced a clear deadline for responding. By contrast, in both radial formats the absence of a terminating boundary may yield flatter temporal weighting.

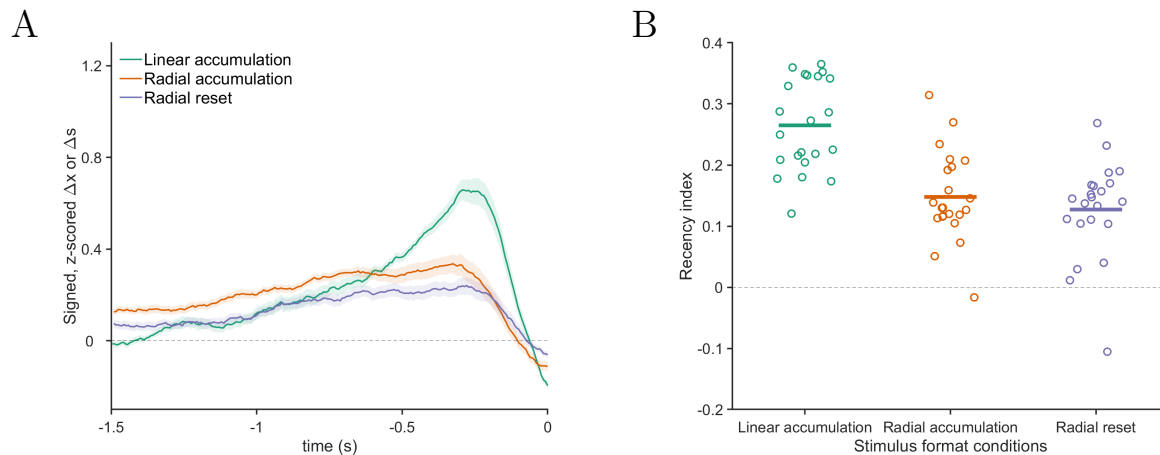


Figure 4.6: **Temporal weighting profiles across stimulus formats.** (A) Reverse correlation kernels at the offset of the trial, showing the last 1.5 seconds before decision time (0 in the x -axis refers to the response). In this figure, for each stimulus format, kernels were computed by averaging decision-locked trial trajectories across entropy levels and participants. Shaded area refers to the SEM across participant kernels. (B) Recency indices obtained from each participant’s reverse correlation kernels defined as the difference between the mean kernel weight in the second and first halves. Positive values indicate a recency effect. Colored horizontal bars show the mean RI for each stimulus format.

4.3.6.2 Salient features in stimulus trajectories (jump analysis)

We then asked whether brief, salient motion segments towards a direction (i.e., “jumps”) served as decision-relevant cues. If observers noticed and used these jumps then features of the jumps that are congruent with the eventual choice should differ systematically from incongruent ones. This analysis complements the temporal weighting results by targeting discrete within-trial segments of high directional evidence and testing whether participants leveraged them, and whether this depends on how the same physical motion was presented (linear vs. radial formats).

Fig. 4.7 illustrates the procedure. For each participant and each trajectory, we first smoothed the trajectory along the horizontal axis, and then scanned it with a moving segment until a decision was made. We marked a segment as a jump when it showed a sustained movement in one direction that lasted at least ~ 200 ms, and covered more than 150 pixels. For each detected jump segment, we fitted a line to obtain the following jump’s features: duration, distance, slope, and latest position. Finally, depending on its direction we labeled each jump as decision-congruent or incongruent (see Section 4.5.4.5).




Figure 4.7: **Jump detection illustration of an example trajectory.** Trajectory until decision (gray line) was smoothed (black line), and it was scanned with a moving segment until decision time. Each segment that had a consistent derivative with a minimum length and duration was detected as jump. In this trajectory two segments were defined as decision-congruent jumps shown within the gray shaded area. As shown here, the jumps were brief motion segments in a direction (either congruent or incongruent with the decision direction). Then, we fitted lines to the detected jumps (red lines) to obtain features of the jumps, such as slope.

Linear accumulation. Fig. 4.8 shows box plots, for each of the identified feature of the jumps (i.e., duration, slope, length, position), the difference between congruent and incongruent segments. As expected, congruent jumps occurred more frequently overall than incongruent jumps (see Supplementary Fig. S13).

The quantification of the differences between decision-congruent and decision-incongruent jumps across the four features revealed that congruent jumps exceeded incongruent ones for all features (two-sided Wilcoxon rank-sum tests: all $U_C > 363/484$, all $p_{\text{Holm}} < .005$; see Supplementary Table S70). Among these, the latest jump position showed the largest effect size (Cohen’s $d = 0.304$), whereas jump slope showed the smallest effect size (Cohen’s $d = 0.058$; see Supplementary Table S79). We confirmed these findings using a paired sign-flip permutation test on the mean difference between congruent and incongruent jumps (all $p_{\text{Holm}} < .0001$; see Supplementary Table S71).

We next wondered whether the influence of the jumps accumulates over a trial. Consistent with the recency effect, we expected later jumps to be more predictive of participants’ choices. To test this,

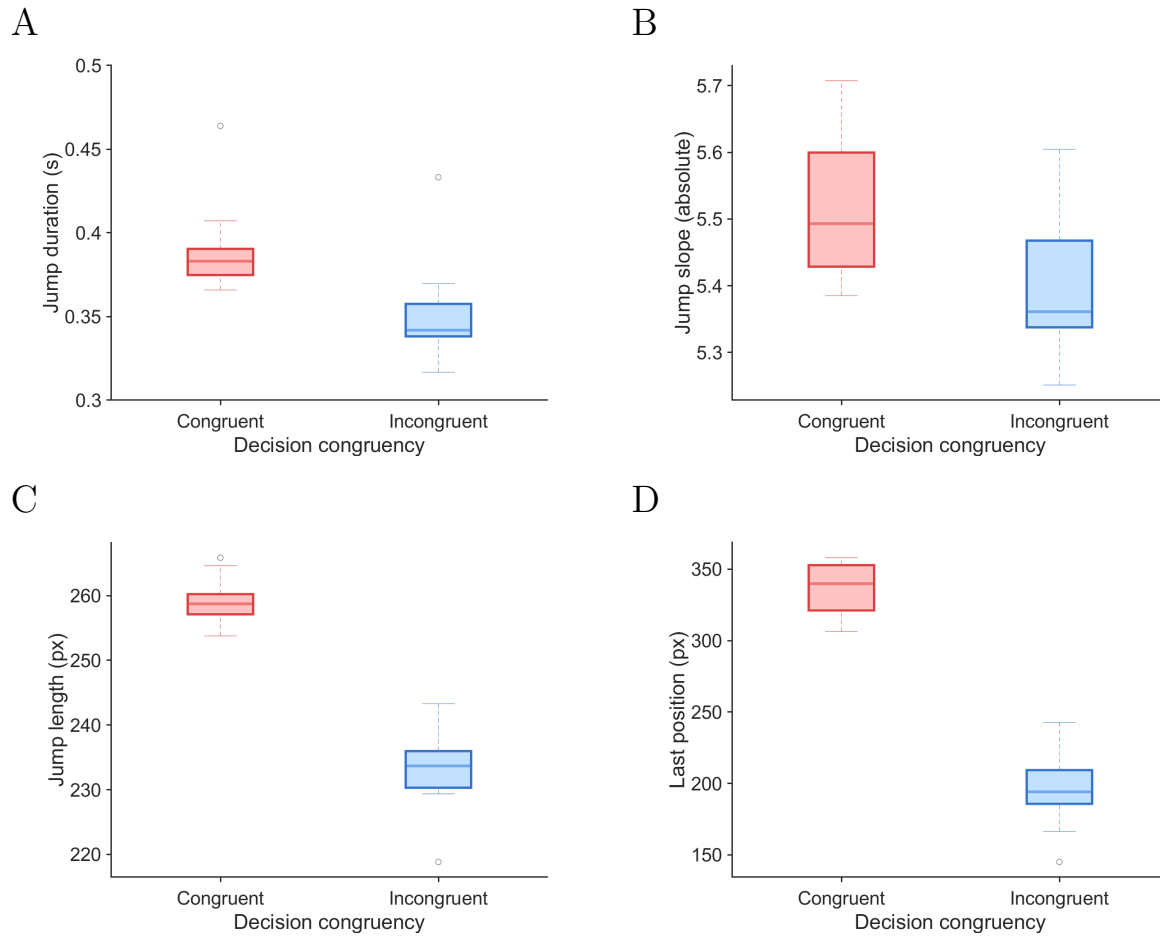


Figure 4.8: **Features of the jumps grouped by decision-congruency in linear accumulation condition.** (A) The duration of the jumps, (B) the slope of the jumps, (C) the spatial length of the jumps, and (D) the last position that the jump arrives (calculated as distance to the starting point of the disk), were divided into whether the jumps' direction was consistent with participants' response (decision-congruent jumps), or not (decision-incongruent jumps). Data points of the box plots were obtained by taking average of the relevant feature within each participant.

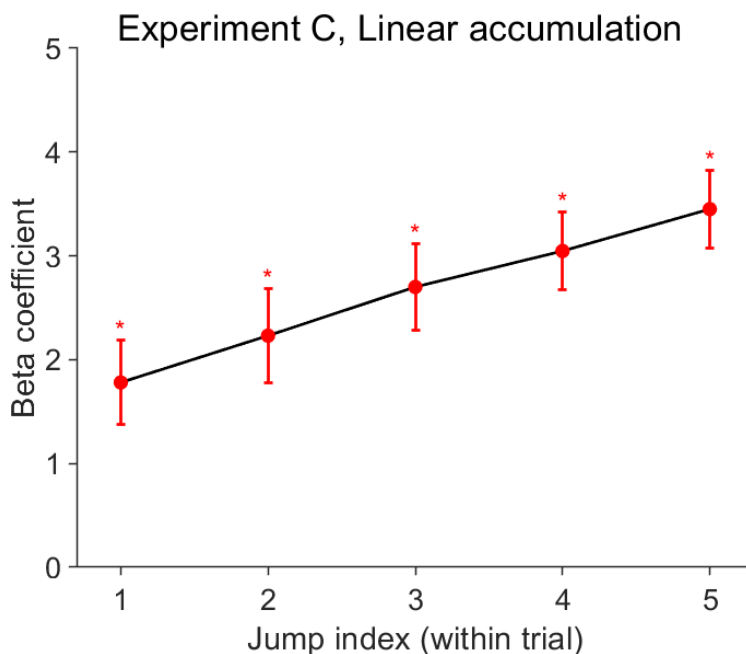


Figure 4.9: **Jump order effect in the linear accumulation condition.** In order to test whether the influence of jumps on decision accumulates over a trial we fitted a logistic regression that predicts the decision from a binary indicator for each jump’s direction. In the linear accumulation format, all beta coefficients were positive and significant (shown with red data points and stars). No such pattern was observed in the radial conditions.

we focused on trials with the same number of detected jumps ($J = 5$), which allowed us to compare jump order directly (first through fifth). We then fitted a logistic regression to predict participant’s decision from the direction of each of the five jumps (see Section 4.5.4.5). As can be seen in Fig. 4.9, all five coefficients were positive and significant (all $|t_{401}| > 4.36$, all $p < 0.0001$), suggesting that all five jumps reliably biased participants’ choices in their direction (See Supplementary Table S72); the model was also significant overall ($F_{5,401} = 155.87$, $p < 0.0001$; Supplementary Table S73). However, multiple comparisons revealed that, after correction, only the first jump was significantly lower from the last jump ($t_{401} = -3.25$, $p = 0.012$). Bonferroni-corrected multiple comparisons are provided in Supplementary Table S74). These results suggest that while jumps seem to have biased participants’ choices, the gradual increase in their effect within a trial was less evident than in the previous experiments (see Chapter 3).

Radial accumulation and radial reset. In the radial accumulation and the radial reset formats, there was no significant difference between congruent and incongruent jumps for any of the jump features. Full results of Wilcoxon rank-sum tests and the paired sign-flip permutation tests are provided in Supplementary Tables S75, S76 for radial accumulation, and in S77, S78 for radial reset.

These results are in line with the reduced recency observed in the radial formats. One possible interpretation is that, in the radial formats, mapping the trajectory onto angular motion may reduce positional cues, and that, in the radial reset format, the interruptions in the trajectory may further disrupt temporal continuity. Thus, in these two formats the same physical motion may be less diagnostic, and observers may not be able to exploit these cues.

4.4 Discussion

Experiment C asked whether sensitivity to the thermodynamic arrow of time—the nonequilibrium signatures embedded in noisy motion—generalizes across stimulus formats that preserve identical trajectory statistics while changing the spatiotemporal layout of stimulus and therefore, plausibly, altering the mapping from motion to decision evidence.

We found that i) in the radial formats participants adopted higher decision thresholds than in the linear accumulation format. Notably, thresholds were closely matched between radial accumulation and radial reset despite the additional disruption introduced by resets. ii) Accuracy showed little systematic modulation by entropy, particularly in both radial formats, and was consistently lower in radial reset than in radial accumulation. iii) Decision times decreased with entropy in all formats, but with a clear gradient: the decrease was strongest in linear accumulation, weaker in radial accumulation, and minimal under radial reset. iv) Reverse correlation analysis revealed that the strongest recency effect was found in the linear accumulation format, with attenuated recency in both radial formats. v) Jump analysis indicated that participants were sensitive to the brief, salient motion dynamics within the trajectory only in the linear accumulation format.

Taken together, these results indicate that nonequilibrium sensitivity is shaped not only by stimulus physics but also by how the same dynamics are rendered into usable evidence. In linear accumulation, proximity to the screen edges provided an additional positional cue correlated with the latent drift direction and imposed an externally terminating boundary. One possible interpretation is that this combination may increase the influence of late evidence, yielding stronger recency, narrower thresholds, and faster decisions—consistent with the closer alignment to SPRT in the absolute-error analysis. In radial accumulation, remapping the trajectory onto an angular display removed the terminating boundary. This may flatten the late-trial advantage and thus reduce recency. In this condition, observers nevertheless maintained high thresholds to keep accuracy relatively stable across entropy levels, while deviating more from SPRT. In radial reset, interruptions in temporal continuity may further introduce memory demands, degrading accuracy and weakening entropy-related modulation of decision time even though thresholds remained similar to radial accumulation.

Several factors should be considered when interpreting these results. First, although the underlying displacement time series in the trajectories were identical across formats, the information available to the observer may not have been equally straightforward to use [90], [91]. Remapping linear motion onto an angular display can change how clearly small displacements and brief salient motion segments stand out, which could contribute to the reduced recency and the absence of jump signatures in the radial formats. Second, the linear format provides an additional positional cue (i.e., distance to the screen edges) and a terminal boundary (i.e., screen edges). These features may encourage participants to adjust their decision strategy over time within a trial (for example, becoming more likely to commit late in the trial), which could strengthen recency. Third, the reset manipulation likely combines several costs: it breaks temporal continuity, may force participants to re-engage attention after each reset, and may encourage them to treat a trial as a set of short segments rather than one continuous stream. Finally, Sequential Probability Ratio Test (SPRT) is a useful benchmark, but it does not include these format-specific constraints, so deviations from SPRT may reflect a mismatch between the model and the task demands rather than a simple failure to use the evidence.

These limitations motivate targeted follow-ups. A direct test of the role of boundaries and positional cues would be to remove edge-based information in the linear format (e.g., by making position non-informative) or, conversely, to introduce a matched terminal constraint in the radial format (e.g., an angular boundary or a fixed stopping rule). This would help separate effects of spatiotemporal layout (stimulus format) from effects of time pressure or termination. To isolate the effect of resetting, future work could vary reset frequency and predictability (rare vs. frequent; regular vs. irregular timing). On the modeling side, a natural next step is to fit format-aware decision models that can capture both accuracy and decision times, such as leaky integration with an explicit penalty or degradation after resets. Finally, the jump results can be strengthened by robustness checks: varying the jump detection thresholds (in time and space), defining jumps in the native coordinate of each format (linear displacement vs. angular change).

In sum, sensitivity to nonequilibrium structure is not a fixed property of the observer but an interaction between stimulus physics and its spatiotemporal layout. When the format provides spatial cues and a terminal boundary, human decisions approached the normative trade-off; when those cues are removed—or temporal continuity is disrupted—behavior remained sensitive to entropy but departed systematically from the SPRT benchmark, revealing the conditions under which the brain can effectively harness the arrow of time.

4.5 Materials and Methods

4.5.1 Participants

All the participants were tested in the Cognitive Neuroscience laboratory of SISSA. Prior to the experiment, the written consent form, approved by the Ethics Committee of the SISSA was obtained from the participants. At the end of the experiments, all participants were provided with a debriefing session.

In Experiment C, twenty-two participants (14F, 8M, mean age=28.77, SD= 4.91, 2 left-handed) with normal or corrected-to-normal vision were recruited. None of them have participated in

Experiment A or B. The participants were tested for approximately 4 hours, depending on their own pace, separated into 2 sessions conducted on different days. They received 10 euros per hour for taking part.

4.5.2 Stimuli and apparatus

At each trial a dark gray disk (diameter = 1° of visual angle) was displayed on a light gray background. From its starting position, at time 0, the disk began to move in a trajectory defined by the nonequilibrium dynamics of an overdamped Brownian particle in two dimensions. The details of the disk motion is explained separately for each of the three stimulus format conditions (linear accumulation, radial accumulation, radial reset) below. All the stimuli were shown on a BenQ ZOWIE XL2731 monitor (resolution = 1920×1080 pixels, refresh rate = 144 Hz). The discrete time step in all numerical simulations was chosen as the inverse of the screen refresh rate. Throughout the experiment, a chin rest was used to restrict head movements and maintain a fixed viewing distance of 57 cm from the computer screen.

Linear accumulation. In this condition, the motion of the disk was generated by Euler numerical integration of the Langevin equation:

$$\dot{X}_t = Hv + \sqrt{2D} \xi_t, \quad (4.1)$$

$$\dot{Y}_t = \sqrt{2D} \eta_t, \quad (4.2)$$

where $v > 0$ is the drift velocity along the x -axis, $D > 0$ is the diffusion coefficient (equal along the x and y axes), and $H \in \{1, -1\}$ specifies the leftward/rightward drift direction, set at random with equal probability in each trial. The quantities ξ_t and η_t are independent Gaussian white noises with zero mean. This simulation mimics the fluctuating motion of an overdamped Brownian particle dragged at constant net speed v in a Newtonian fluid with viscosity $\gamma = k_B T / D$, where k_B is Boltzmann's constant and T the temperature of the fluid.

A trial finishes when the disk reaches the left or right edges of the screen along the x -axis. Reflecting

boundary conditions were implemented along the y -axis, so that the disk bounces back when contacting the vertical edges.

Radial accumulation. In this condition, we used the same Brownian particle motion, but the horizontal displacement from the linear accumulation was mapped onto an angular displacement along the circumference (shown as a visible black circle in the screen). At the beginning of each trial, the disk was placed at a random position on the circle. The circle's center coordinates $(X_{\text{center}}, Y_{\text{center}})$ correspond to the screen center, and its nominal radius was $r_0 = 400$ pixels. To introduce decision-unrelated variability analogous to the vertical motion noise present in the linear accumulation condition, the circle's radius was jittered on every frame. Specifically, the instantaneous radius was defined as $r_t = r_0 + \varepsilon_t$, where ε_t is drawn from a Gaussian distribution with zero mean and a standard deviation of 5 pixels. The disk's position at each time step was then computed as

$$X_{t+\Delta t} = r_t \cos(\theta_{t+\Delta t}) + X_{\text{center}}, \quad (4.3)$$

$$Y_{t+\Delta t} = r_t \sin(\theta_{t+\Delta t}) + Y_{\text{center}}. \quad (4.4)$$

The angular position was updated from the previous step according to

$$\theta_{t+\Delta t} = \theta_t + \frac{\Delta s_t}{r_t}, \quad (4.5)$$

where Δs_t is the arc-length displacement during the time interval Δt . The term Δs_t is derived directly from the horizontal displacement in the linear accumulation condition,

$$\Delta s_t = H v \Delta t + \sqrt{2D} \Delta W_t, \quad (4.6)$$

where v is the drift velocity, $H \in \{-1, +1\}$ is the randomly assigned drift direction, D is the diffusion coefficient, and ΔW_t is Gaussian white noise with zero mean. This formulation ensures that the motion in the radial accumulation condition was generated from the same underlying stochastic

process as in the linear accumulation condition, with the horizontal displacement reinterpreted as an arc-length displacement along the circle’s circumference. A red cross (± 10 pixels in horizontal and vertical extent, line width = 3 px) is displayed at the center of the circle throughout the trial.

Radial reset. This condition shares the same circular geometry, jittering, and motion transformation as the radial accumulation condition. However, the disk resets to its starting position on the circle after a randomly chosen reset interval of approximately 0.4–0.5 seconds. The starting position for each trial is marked with a red cross identical in size and line width to the one described above, but here placed directly at the disk’s starting point rather than the circle center.

4.5.3 Task and experimental design

Experiment C was designed as a two-alternative forced choice task. In the linear accumulation condition, participants decided whether the disk’s motion is rightward or leftward. In the radial accumulation and radial reset conditions, participants decided whether the disk’s motion is clockwise or anticlockwise along the circle’s circumference. In all conditions, responses were collected via a two-button pad, with the right button corresponding to right or clockwise decisions depending on the condition.

For each trial and each participant, we recorded: i) the true hypothesis (ground truth), i.e., the actual direction of the net drift $H \in \{1, -1\}$; ii) the decision outcome $D_{\text{dec}} \in \{1, -1\}$; iii) the decision time T_{dec} ; iv) the position of the disk along the x - and y -axes of the screen at the decision time, $(X_{\text{dec}}, Y_{\text{dec}}) = (X(t = T_{\text{dec}}), Y(t = T_{\text{dec}}))$, from which decision thresholds (the disk position at decision time along the decision-relevant axis) were later determined; and v) the full stochastic trajectory $(X_{[0, T_{\text{dec}}]}, Y_{[0, T_{\text{dec}}]})$ of the disk up to T_{dec} .

After participants completed training session to become familiar with the task, experimental data were collected. In both training and experimental sessions, participants received immediate feedback after each trial, on any of the three possible outcomes they can get from their response together with an accumulating score: i) correct responses (+40 points), ii) incorrect responses (-20 points),

and iii) missed trials (-50 points), accompanied by a sound and color of the feedback text according to the outcome. The next trial started automatically after the feedback screen.

Each of the three conditions (linear accumulation, radial accumulation, radial reset) was tested at three levels of entropy production rate in the stimulus dynamics, a stimulus parameter defined by the ratio of squared drift velocity to diffusion coefficient (see the next Modeling Section 4.5.4.2 for details). The particular velocity and diffusivity values are presented in Supplementary Table S50. For each condition, participants performed 6 blocks of 36 trials (total of 648 trials), with trials from all three entropy production rate levels intermixed within each block (intermixed design). This results in 18 blocks in total per participant, with the order of the three main conditions counterbalanced across participants.

To explore the decision consistency across participants given a particular stimulus, we presented the same set of trials to each participant, but randomized the presentation order of the trials across participants.

All displays and data collection are coded using the MATLAB software and Psychophysics Toolbox Version 3 (PTB-3) [66].

4.5.4 Data Analysis

4.5.4.1 Behavioral performance analysis

The behavioral analysis aimed to assess the effects of experimental manipulation (stimulus format, entropy production rate) on the decision metrics: decision threshold, accuracy, and decision time. In this set of analyses, we tested these decision metrics in Experiment C. Prior to the analysis, we removed the "missed trials" ($\sim 2\%$ for Experiment C) from the dataset, the cases where the participant did not take a decision before the trial ends. All decision metrics (decision thresholds, accuracy, and decision times) were averaged across each participant, stimuli format, and entropy level.

Decision thresholds X_{dec} . The mean decision thresholds $\langle X_{\text{dec}} \rangle$ were tested using a linear mixed-effects (LME) model with fixed effects of i) stimulus format (linear accumulation, radial accumulation, radial reset), and ii) entropy production rate (referred to as entropy); as well as their interaction, and with a random effect of participant (id) (marginal R^2 : 0.55, conditional R^2 : 0.83):

$$\langle X_{\text{dec}} \rangle \sim \text{stimulus format} * \text{entropy} + (1 \mid \text{id}). \quad (4.7)$$

Accuracies α . Accuracy was analyzed using a LME model with the same fixed and random effects used as before (marginal R^2 : 0.18, conditional R^2 : 0.49):

$$\alpha \sim \text{stimulus format} * \text{entropy} + (1 \mid \text{id}). \quad (4.8)$$

Decision times T_{dec} . Log-transformed mean decision times $\langle T_{\text{dec}} \rangle$ were also analyzed using a LME model with the same fixed and random effects used as before (marginal R^2 : 0.50, conditional R^2 : 0.85):

$$\langle T_{\text{dec}} \rangle \sim \text{stimulus format} * \text{entropy} + (1 \mid \text{id}). \quad (4.9)$$

All LME models were fitted with the `lme4` package [67]. LME model variance explained was computed using the `MuMIn` package [68]. The Satterwaite’s method [69] implemented in the `lmerTest` package [70] was used to estimate the degrees of freedom for the LME model ANOVA, and estimated marginal means were compared using the `emmeans` package [71] using Bonferroni adjustment.

4.5.4.2 Sequential probability ratio test (SPRT)

The details of the SPRT model presented in this chapter are provided in Section 2.5.4.2.

4.5.4.3 Comparison between observed data and SPRT predictions

To assess the extent to which human decisions approximate optimality, we quantified suboptimality as the deviation between participants' behavioral performance and the predictions of SPRT.

Accuracies α . First, we calculated the absolute error between predicted and observed accuracy for each participant in each stimulus format and each entropy level. Then, we analyzed the absolute errors in accuracies (shown as AE_α) using a LME with fixed effects of stimulus format, and entropy, as well as their interaction; and a random effect of participant (id), with the following formula (marginal $R^2 = 0.39$, conditional $R^2 = 0.45$):

$$AE_\alpha \sim \text{stimulus format} * \text{entropy} + (1 \mid \text{id}). \quad (4.10)$$

Decision times T_{dec} . Similarly, we calculated the absolute error between the predicted and observed $\langle T_{\text{dec}} \rangle$ for each participant in each stimulus format and each entropy level. Due to the skewed error distributions, we applied a square-root transformation prior to modeling. Then, we analyzed the absolute errors in $\langle T_{\text{dec}} \rangle$ (shown as $AE_{\langle T_{\text{dec}} \rangle}$) using a LME with the same fixed and random effects used as before (marginal $R^2 = 0.47$, conditional $R^2 = 0.56$):

$$AE_{\langle T_{\text{dec}} \rangle} \sim \text{stimulus format} * \text{entropy} + (1 \mid \text{id}). \quad (4.11)$$

We applied the same methodological procedures described in Section 4.5.4.1.

4.5.4.4 Temporal weighting profile

To assess how participants used stimulus information across time within each trial, we estimated temporal weighting profiles using a decision-locked reverse-correlation analysis. This method allows us to infer the contribution of moment-to-moment fluctuations in the horizontal motion signal to the eventual choice, providing an estimate of the reverse correlation kernel for each experiment.

Reverse correlation kernels. For each trial, we extracted the displacement of the stimulus motion in decision-relevant (horizontal) axis, defined as Δx_t . We computed these displacements as running differences over a smoothing window of 0.2 s, converted into the corresponding number of frames given the screen refresh rate. To enable comparisons across entropy levels and participants, we standardized the displacement values by z -scoring within each participant \times stimulus format \times entropy level:

$$z_t = \frac{\Delta_t - \mu}{\sigma}, \quad (4.12)$$

where μ and σ are the mean and standard deviation of the displacement series across all trials for that participant and entropy level. This procedure removes scale differences in variance introduced by entropy while preserving the temporal structure of fluctuations.

Kernels were computed as decision-signed averages of the standardized displacements: each trial's z_t series was multiplied by the participant's decision sign ($D_{\text{dec}} \in \{-1, +1\}$, coding leftward/anticlockwise as -1 and rightward/clockwise as $+1$), and then averaged across N trials:

$$K(\tau) = \frac{1}{N} \sum_{i=1}^N D_{\text{dec}}^{(i)} z_{\tau}^{(i)}, \quad (4.13)$$

where i indexes trials and τ denotes time relative to the decision.

We analyzed kernels aligned to the decision moment, spanning a fixed time window of 1.5 s prior to response. For each trial, we took the last 1.5 s of the standardized displacement series, reversed its time axis so that $\tau = 0$ corresponds to the decision, and averaged across trials. Trials that have a decision time below 1.5 second were skipped, which corresponds to 12.3 % of the trials in linear accumulation; 2.3 % of the trials in radial accumulation; and 1.9 % of the trials in radial reset condition.

Recency indices. To quantify the relative influence of late versus early evidence, we computed a recency index (RI) as the difference in mean kernel amplitude between the last and the first halves

of the decision-locked window until the peak:

$$RI = \langle K_{\text{late}} \rangle - \langle K_{\text{early}} \rangle. \quad (4.14)$$

Positive RI values indicate recency (greater weight on later evidence), while negative values indicate primacy (greater weight on earlier evidence).

Linear mixed-effect models. Recency indices derived from individual reverse correlation kernels (and transformed as $\log(\text{constant} + \text{recency index})$) were analyzed using a LME with the same fixed and random effects used as before using the following formula (marginal R^2 : 0.30, conditional R^2 : 0.38):

$$RI_{\text{dec}} \sim \text{stimulus format} * \text{entropy} + (1 \mid \text{id}). \quad (4.15)$$

We applied the same methodological procedures described in Section 4.5.4.1.

4.5.4.5 Salient features in stimulus trajectories (jump analysis)

To test whether participants detect and use brief statistical cues in the stimulus trajectories, we performed a complementary analysis focusing on jumps. Jumps were defined as short temporal segments in which the disk’s velocity shows a consistent displacement in one direction, exceeding the fluctuations expected from background diffusion.

Jumps. For each trial, we extracted the trajectory along the decision-relevant axis (Δx_t for the linear format and Δs_t for the radial formats). We smoothed this trajectory using MATLAB’s `smooth` function with the Savitzky–Golay option (`'sgolay'`), applying a smoothing window of 50 frames (corresponding to ~ 0.35 s at 144 Hz). We then scanned the smoothed trajectory with a moving segment until the decision time. A segment was classified as a jump if it satisfies three conditions: 1) the derivative of the segment has a consistent sign, 2) the segment length is at least 25 frames (≈ 200 ms), 3) the net displacement between start and end points of the segment is larger

than 150 pixels. This procedure ensures that detected jumps reflect coherent, directional motion rather than random noise.

Features of jumps. To fit each jump segment, we used MATLAB’s `fit` function with a first-order polynomial model (`poly1`), which returns the slope (p_1) and intercept (p_2):

$$\hat{y}(t) = p_1 t + p_2. \tag{4.16}$$

We retrieved four parameters from these fits as features of the jumps: The duration, spatial distance, slope and spatial position at the segment end (last position) of the jumps. All the features were recorded separately for decision congruent and incongruent segments.

Testing congruent vs. incongruent jump features. Once jumps were identified, we quantified their behavioral impact in the following analyses. For each participant p , feature f (duration, spatial distance, slope and spatial position), and congruency (incongruent vs. congruent), we computed a participant-level summary by averaging the feature across all jumps detected before the decision within that cell (participant \times feature \times congruency). This produced two paired vectors for each feature: C (congruent) and I (incongruent).

For each feature, we formed paired differences:

$$d_{p,f} = C_{p,f} - I_{p,f}. \tag{4.17}$$

We tested whether the mean difference d is zero with a two-tailed paired sign-flip permutation test that is magnitude-sensitive. Under the null hypothesis, the signs of $d_{p,f}$ are exchangeable. We generated a null distribution by randomly flipping signs (Monte Carlo with 100.000 iterations) and computed the mean d as statistic, as well as the p -value. We controlled the family-wise error rate across the four features within each experiment with the Holm step-down adjustment. As an effect size in original units, we reported the Hodges-Lehmann estimate with nonparametric 95% confidence intervals from a bootstrap over participants (5.000 resamples with replacement).

We also compared the congruent and incongruent distributions with a two-sided Wilcoxon rank-sum (Mann-Whitney U) test per feature. Alongside the two-sided p -value, we reported the common-language effect size

$$\text{CL}_f = P(C > I) = \frac{U_C}{n_I n_C}, \quad (4.18)$$

where U_C is the Mann-Whitney statistic for the congruent group C against the incongruent group I (mid-ranks for ties), and n_I, n_C are the group sizes. Here, $P(C > I)$ means the probability that a randomly drawn observation from the congruent group exceeds a randomly drawn observation from the incongruent group. We quantified uncertainty in CL_f with a stratified bootstrap ($R = 5,000$): we resampled participants with replacement within each group, hold n_I, n_C fixed, and recomputed CL_f ; we take the 2.5% and 97.5% percentiles as the confidence interval. We applied the Holm adjustment across the four features to control the family-wise error rate.

For completeness, we also computed Cohen's d as effect size of each feature on the same participant-level summaries as:

$$d = \frac{\bar{X}_{\text{cong}} - \bar{X}_{\text{incong}}}{s_p}, \quad (4.19)$$

$$s_p = \sqrt{\frac{(n_{\text{cong}} - 1) \text{var}(X_{\text{cong}}) + (n_{\text{incong}} - 1) \text{var}(X_{\text{incong}})}{n_{\text{cong}} + n_{\text{incong}} - 2}}. \quad (4.20)$$

Jump integrator. Next, to test whether participants integrate jumps to guide decisions, we fitted logistic regression models using MATLAB's `fitglm` function with a binomial distribution and logit link. Trials with the same number of jumps J ($J = 5$) were selected. In the linear accumulation condition, the trials that had 5 jumps corresponded to the 9.01% of the all trials; and the same value was 5.17% and 8.66% in radial accumulation and radial reset conditions respectively. For each trial, we constructed a feature vector $X = (x_1, x_2, \dots, x_J)$, where $x_j = 1$ if the j th jump was in the participant's chosen direction, and $x_j = 0$ otherwise. The binary decision outcome $y \in \{0, 1\}$ was modeled as

$$\Pr(y = 1 \mid X) = \frac{1}{1 + \exp \left[- \left(\beta_0 + \sum_{j=1}^J \beta_j x_j \right) \right]}. \quad (4.21)$$

The regression coefficients β_j quantified the influence of each jump on the decision. A positive and significant β_j indicates that the corresponding jump reliably biases choices in its direction.

The fitted coefficients, their standard errors, and associated p -values were extracted directly from the `fitglm` output. We visualized these coefficients across jump order and highlighted statistically significant effects.

To assess whether the set of jump-index predictors jointly improved fit, we computed an ANOVA-style omnibus statistic on the model estimates and reported the F approximation. Pairwise differences between jump positions were tested as model-based coefficient contrasts using t -tests (estimate/SE with the GLM covariance), reporting $\Delta\beta$, SE, $t(\text{df})$, and Bonferroni-adjusted two-sided p -values across the 10 comparisons.

Chapter 5

General Discussion and Conclusion

This thesis examined whether perceptual decision-making is sensitive to the temporal organization of a noisy evidence stream and, in particular, to its degree of nonequilibrium structure. To address this question, we introduced a stimulus framework in which visual motion trajectories are generated by drifted Brownian dynamics whose entropy production rate (Σ) can be manipulated parametrically. Across three behavioral experiments with human participants, we examined how graded changes in nonequilibrium fluctuations shape behavioral metrics (i.e., accuracy, decision times, and decision thresholds), and we linked these behavioral signatures to computational accounts spanning normative Sequential Probability Ratio Test (SPRT) and an Evidence Integration Model (EIM) with memory constraints. We then asked which aspects of the stimulus dynamics influence participants' choices, focusing on i) how observers weight different moments of the stimulus trajectory and ii) observers' sensitivity to the salient motion segments in the trajectory. Experiment A led us to investigate observers' perceptual decision-making mechanisms under nonequilibrium fluctuations using the above-mentioned model-based and data-driven analyses. Experiment B further shed light on such mechanisms using i) a larger stimulus parameter space that allowed us to assess individual contributions of velocity (v) and diffusivity (D), and ii) an intermixed design where the entropy production rate changed on a trial-by-trial basis. Finally, in Experiment C, we investigated how changing the spatiotemporal layout of the same physical evidence might affect the decision performance.

5.1 Summary of behavioral signatures and comparison to an optimal observer

The behavioral results across the three experiments demonstrated that decision times, decision thresholds, and accuracy depended jointly on the nonequilibrium dynamics of the stimulus and on the constraints imposed by the task format (Chapters 2 , 3, and 4). First, reaction times decreased with increasing entropy (Σ) in all three experiments. Second, in the linear accumulation conditions decision thresholds were largely stable across entropy levels, whereas, in the radial conditions they increased with higher entropy production rate. Third, accuracy was related to both decision thresholds and the velocity-to-diffusivity ratio. Accuracy increased with higher decision thresholds, and also with higher velocity and lower diffusivity. Then, we compared these behavioral metrics with the optimal decision-maker using the Sequential Probability Ratio Test (SPRT) (Chapters 2 , and 4). While the SPRT provided a strong normative benchmark for perceptual decision-making, the observed behavioral metrics (e.g., decision times) showed clear deviations from its predictions. Specifically, participants' decision times were consistently longer than those predicted by SPRT given their accuracy levels, pointing to suboptimal decision-making strategies in human observers. Importantly, this suboptimality increased further in Experiment B. In Experiment C, we again observed suboptimality in decision times, with deviations increasing gradually from linear accumulation to radial accumulation (i.e., where the stimulus was presented along a radial trajectory), and further to radial reset (i.e., where the trajectory was not constantly available to participants).

One way to interpret these patterns is that higher entropy production makes the evidence stream contain more usable structure over time (i.e., more moments where the motion is consistently informative) so observers are more likely to encounter a “decisive” stretch of evidence earlier and can commit sooner. At the same time, the fact that thresholds changed with entropy mainly in the radial formats suggests that the effect of nonequilibrium structure depends on how the evidence is represented: when the mapping from the stimulus display to the decision variable is less direct (e.g., in radial accumulation and radial reset conditions), observers may compensate by adopting a more

cautious criterion as entropy increases. Finally, the growing mismatch between SPRT predictions and behavior in the intermixed design (Experiment B) and in the radial formats (Experiment C) is consistent with limits that are not part of the ideal observer—such as imperfect calibration to changing stimulus statistics (i.e., intermixed-design), and extra processing demands introduced by the stimulus format (i.e., radial accumulation and radial reset)—which can preserve accuracy while lengthening decision times relative to a normative benchmark.

5.2 Memory-limited integration: insights from the EIM

We used an Evidence Integration Model (EIM) with memory constraints to capture aspects of evidence accumulation that are absent from the SPRT (Chapter 2). In particular, the SPRT assumes lossless integration of momentary likelihood increments, whereas human observers may rely on a finite temporal window that effectively “forgets” older evidence. In Experiment A, the EIM provided a better account of trial-by-trial variability than the SPRT, reproducing both the broad distributions of decision times and the observed spread in decision thresholds. Moreover, the temporal integration window τ_m^* estimated by the EIM increased with entropy production rate. While this is only an interpretation, it is tempting to assume that participants may have benefited from longer integration when the stimulus dynamics were more strongly driven, potentially because informative structure persists over longer timescales.

Notably, this entropy-dependent modulation of the integration window was not evident in Experiment B. One plausible interpretation is that unpredictable trial-to-trial changes in Σ (as in the intermixed design) may interfere with systematic tuning of the integration window with Σ , such that observers rely on a relatively short τ_m^* that is not finely tuned to the stimulus statistics of the current trial. In this view, the difference between Experiments A and B reflects how task structure (blocked vs. intermixed) shapes the ability to tune evidence integration to the temporal structure of the evidence stream.

We could not apply the EIM to Experiment C (Chapter 4). While the physical evidence of the stimulus trajectory was matched across stimulus formats, the radial accumulation and radial reset

displays introduced additional perceptual constraints because the display changes the spatial layout and continuity of the stimulus trajectory. Under these conditions we may require extending the EIM model—for example, to accommodate more flexible mappings from the stimulus trajectory to the decision variable. This leaves an important limitation: the current EIM formalization was validated primarily in the linear format and may not generalize without modification to radial representations. One possibility is that radial formats change how observers encode and retain evidence—for instance, by making evidence updates depend on where the stimulus is in the radial display (e.g., due to nonuniform encoding/attention around the circle), rather than assuming that the motion fluctuations translate into the same incremental evidence update everywhere along the circular trajectory. Another possibility is that resets disrupt the perceived continuity of the trajectory, weakening the sense that each momentary evidence belongs to the same ongoing process. Future work should therefore develop and test a more generalized EIM that explicitly links the stimulus format to the dynamics of accumulation.

5.3 How stimulus dynamics shape choices: perspectives from temporal weighting profile and jump analysis

We used decision-locked reverse correlation method to quantify participants' temporal weighting profiles, that is, the observers tendency to weight differently different moments of the stimulus trajectory (Chapters 3, and 4). This analysis revealed a robust recency effect across all experiments, indicating that late portions of the stimulus trajectory exerted a stronger influence on choice than early portions. However, the entropy-dependent modulation of recency (clearly expressed in Experiment A) was attenuated in Experiment B and was not evident in Experiment C. Thus, while observers reliably weighted recent evidence, their tendency to tune this temporal weighting to nonequilibrium structure depended strongly on whether the Σ was stable across trials (i.e., blocked design) and on how the evidence was presented (i.e., stimulus format).

A plausible interpretation is that entropy-dependent changes in temporal weighting reflect an adaptive strategy that requires consistent expectations about the statistics of the evidence stream. In

the blocked design of Experiment A, observers can gradually learn and maintain an integration process that exploits the temporal organization of fluctuations, resulting in a graded increase in recency with entropy. In contrast, in the intermixed design (Experiment B), frequent trial-to-trial changes in the driving statistics may discourage such tuning, promoting a more fixed weighting policy and thereby flattening the recency–entropy relationship. These results point to additional adaptive mechanisms—such as learning of environmental statistics, strategy switching, or context-dependent settings—that have been emphasized in related work on adaptive evidence accumulation [4], [7], [92]. In Experiment C, changes in the stimulus format (linear vs. radial) may introduce additional representational demands (e.g., transforming the stimulus into a decision-relevant coordinate system) because radial patterns require integrating motion with position [90], [91], and if this process is noisy or costly, it could compress temporal-weighting profiles and attenuate any entropy modulation.

We also asked whether observers were sensitive to brief, salient motion segments in the stimulus trajectory (“jumps”; Chapters 3, and 4). Across experiments, jump-related signatures were qualitatively consistent in the linear accumulation format: i) the features of the jumps (i.e., their duration, slope, length, and position) contributed to the decision, ii) jumps biased choices toward their direction, and iii) their influence increased when they occurred closer to the decision, indicating that participants did not treat the evidence stream as homogeneous in time. However, jump effects were reduced when task demands increased. In Experiment B, jump-based influences appeared weaker, suggesting reduced reliance on brief, salient motion segments under the intermixed design. In Experiment C, jump effects were further attenuated in the linear condition and did not generalize clearly to the radial formats. In the radial accumulation and radial reset conditions, we did not observe the same systematic jump–choice relationship seen in linear accumulation. Thus, sensitivity to salient local motion events is present and reproducible under the standard linear mapping, but it becomes less evident when the environment varies unpredictably across trials or when the stimulus format alters how evidence is expressed.

One interpretation is that sensitivity to jumps reflects a tendency to pick up brief, event-like

motion segments in addition to integrating evidence gradually over time. Whether this shows up in behavior may depend on how directly the display maps motion onto the decision-relevant direction. In the linear format, jumps are aligned with the left–right decision axis, so they are easy to identify and can be used as strong evidence. By contrast, in the intermixed design (Experiment B), trial-to-trial variability in Σ may limit observers’ ability to tune their expectations to the current stimulus statistics, leading to less well-matched attentional allocation and reduced sensitivity to brief features such as jumps [86], [87]. In the radial formats, the same physical motion may be transformed by the display geometry, so that a coherent “jump” is harder to perceive as a single directional event. On this account, weak or absent jump effects in Experiment B and in the radial conditions of Experiment C do not mean that observers ignore the trajectory; rather, it suggests that the detectability of these salient, directional events depends on how the evidence is displayed and on whether observers can keep a consistent strategy for using such events.

5.4 Generality and limitations of the stimulus family

A key limitation of this work is that the experimental manipulation of nonequilibrium structure was implemented within a specific stimulus family—drifted Brownian motion with parametrically controlled entropy production—and it remains open how broadly the observed signatures generalize beyond this generative model. Although this choice offered tight experimental control and a direct link to stochastic thermodynamics [13], real-world sensory streams often involve richer structure [93] that could interact with decision strategies in qualitatively different ways. For instance, i) natural signals may show non-Gaussian fluctuations, with occasional large events occurring more frequently than would be expected under Gaussian noise [93], [94]; ii) natural inputs can also exhibit long-range temporal correlations, such that fluctuations separated by substantial time intervals remain statistically related [94]; iii) sensory streams can be nonstationary, meaning that the stimulus statistics may undergo sudden transitions [4] (e.g., changing the v or D multiple times within a trial); iv) moreover, such changes may unfold on multiple timescales simultaneously [94] (e.g., relatively fast fluctuations in D superimposed on slower variations in v). Extending the approach

to other stimulus classes would test whether the evidence integration processes reflect a general sensitivity to nonequilibrium structure or a property of the particular evidence statistics used in this work.

5.5 Neural predictions and future directions

A natural next step is to connect these behavioral and computational signatures of nonequilibrium sensitivity to their underlying neural implementations. Future work can combine the present stimulus framework with time-resolved recordings of brain activity (i.e., EEG/MEG) to test whether entropy production modulates canonical decision signals—such as the build-up rate and amplitude of a centroparietal decision variable (CPP-like dynamics) [95], [96], or motor-preparatory lateralization (LRP) [97]. Similarly, one can test the predictions of the temporal weighting profiles using neural reverse correlation by regressing neural decision signals onto stimulus history [95], [98]. Beyond observing the neural correlates of decision making under nonequilibrium dynamics, it may be equally interesting to use pupillometry and pharmacology-linked arousal manipulations, to probe whether entropy-related effects are mediated by neuromodulatory mechanisms of gain (i.e., the effective strength of evidence impact) and urgency (i.e., a rising drive to commit a decision as time elapses) [99]–[101]. Together, these approaches would help to identify which processing stage (e.g., sensory encoding, evidence accumulation, or motor preparation/execution) is mostly affected by nonequilibrium structure.

5.6 Concluding remarks

Overall, the results of this thesis support the view that perceptual decisions are shaped by the nonequilibrium structure of fluctuations and by the task constraints. By bringing tools from stochastic thermodynamics into an experimentally controllable perceptual decision paradigm, this thesis provides a quantitative route for studying how humans adapt evidence integration to temporally structured noise. More broadly, it motivates a mechanistic approach linking the physics of sensory fluctuations to adaptive computation in the brain, and suggests that understanding

decision-making in natural settings will require models that jointly account for stimulus dynamics, context-dependent tuning of integration, and spatiotemporal format of the stimulus.

References

- [1] R. Ratcliff, P. L. Smith, S. D. Brown, and G. McKoon, “Diffusion decision model: Current issues and history,” *Trends Cogn. Sci.*, vol. 20, no. 4, pp. 260–281, 2016.
- [2] R. Ratcliff and G. McKoon, “The diffusion decision model: Theory and data for two-choice decision tasks,” *Neural Comput.*, vol. 20, no. 4, pp. 873–922, 2008.
- [3] B. U. Forstmann, R. Ratcliff, and E.-J. Wagenmakers, “Sequential sampling models in cognitive neuroscience: Advantages, applications, and extensions,” *Annu. Rev. Psychol.*, vol. 67, p. 641, 2016.
- [4] C. M. Glaze, J. W. Kable, and J. I. Gold, “Normative evidence accumulation in unpredictable environments,” *Elife*, vol. 4, e08825, 2015.
- [5] A. T. Piet, A. El Hady, and C. D. Brody, “Rats adopt the optimal timescale for evidence integration in a dynamic environment,” *Nature communications*, vol. 9, no. 1, p. 4265, 2018.
- [6] Z. P. Kilpatrick, W. R. Holmes, T. L. Eissa, and K. Josić, “Optimal models of decision-making in dynamic environments,” *Current opinion in neurobiology*, vol. 58, pp. 54–60, 2019.
- [7] P. R. Murphy, N. Wilming, D. C. Hernandez-Bocanegra, G. Prat-Ortega, and T. H. Donner, “Adaptive circuit dynamics across human cortex during evidence accumulation in changing environments,” *Nature neuroscience*, vol. 24, no. 7, pp. 987–997, 2021.
- [8] M. Ruesseler, L. A. Weber, T. R. Marshall, J. O’Reilly, and L. T. Hunt, “Quantifying decision-making in dynamic, continuously evolving environments,” *Elife*, vol. 12, e82823, 2023.
- [9] P. R. Murphy, K. Krkovic, G. Monov, N. Kudlek, T. Lincoln, and T. H. Donner, “Individual differences in belief updating and phasic arousal are related to psychosis proneness,” *Communications Psychology*, vol. 2, no. 1, p. 88, 2024.

- [10] H. Park, J.-M. Lueckmann, K. von Kriegstein, S. Bitzer, and S. J. Kiebel, “Spatiotemporal dynamics of random stimuli account for trial-to-trial variability in perceptual decision making,” *Sci. Rep.*, vol. 6, no. 1, pp. 1–17, 2016.
- [11] P. R. Fard, S. Bitzer, S. Pannasch, and S. J. Kiebel, “Stochastic motion stimuli influence perceptual choices in human participants,” *Front. Neurosci.*, vol. 15, p. 1893, 2022.
- [12] O. Ossmy, R. Moran, T. Pfeffer, K. Tsetsos, M. Usher, and T. H. Donner, “The timescale of perceptual evidence integration can be adapted to the environment,” *Current biology*, vol. 23, no. 11, pp. 981–986, 2013.
- [13] É. Roldán, I. Neri, M. Dörpinghaus, H. Meyr, and F. Jülicher, “Decision making in the arrow of time,” *Phys. Rev. Lett.*, vol. 115, no. 25, p. 250 602, 2015.
- [14] A. Seif, M. Hafezi, and C. Jarzynski, “Machine learning the thermodynamic arrow of time,” *Nat. Phys.*, vol. 17, no. 1, pp. 105–113, 2021.
- [15] G. Lan, P. Sartori, S. Neumann, V. Sourjik, and Y. Tu, “The energy–speed–accuracy trade-off in sensory adaptation,” *Nature Physics*, vol. 8, pp. 422–428, 2012. DOI: 10 . 1038 / nphys2276.
- [16] S. Bo, M. Del Giudice, and A. Celani, “Thermodynamic limits to information harvesting by sensory systems,” *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2015, no. 01, P01014, 2015. DOI: 10.1088/1742-5468/2015/01/P01014.
- [17] X. Fang, K. Kruse, T. Lu, and J. Wang, “Nonequilibrium physics in biology,” *Rev. Mod. Phys.*, vol. 91, no. 4, p. 045 004, 2019.
- [18] M. Reinl and A. Bartels, “Perception of temporal asymmetries in dynamic facial expressions,” *Frontiers in psychology*, vol. 6, p. 1107, 2015.
- [19] K. Meding, I. Shehet, B. Schölkopf, F. A. Wichmann, and D. Janzing, “Perceiving the arrow of time in autoregressive motion,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [20] N. Hanyu, K. Watanabe, and S. Kitazawa, “Ready to detect a reversal of time’s arrow: A psychophysical study using short video clips in daily scenes,” *Royal Society open science*, vol. 10, no. 4, p. 230 036, 2023.

- [21] A. Wald and J. Wolfowitz, “Optimum character of the sequential probability ratio test,” *Ann. Math. Stat.*, pp. 326–339, 1948.
- [22] J. I. Gold, M. N. Shadlen, *et al.*, “The neural basis of decision making,” *Annu. Rev. Neurosci.*, vol. 30, no. 1, pp. 535–574, 2007.
- [23] A. Wald, “Sequential tests of statistical hypotheses,” *Ann. Math. Stat.*, vol. 16, no. 2, pp. 117–186, 1945.
- [24] A. Tartakovsky, I. Nikiforov, and M. Basseville, *Sequential Analysis: Hypothesis Testing and Changepoint Detection*. CRC press, London, 2015.
- [25] T. J. Kobayashi, “Implementation of dynamic bayesian decision making by intracellular kinetics,” *Physical review letters*, vol. 104, no. 22, p. 228 104, 2010.
- [26] E. D. Siggia and M. Vergassola, “Decisions on the fly in cellular sensory systems,” *Proc. Natl. Acad. Sci. U.S.A.*, vol. 110, no. 39, E3704–E3712, 2013.
- [27] J. Desponds, M. Vergassola, and A. M. Walczak, “A mechanism for hunchback promoters to readout morphogenetic positional information in less than a minute,” *eLife*, vol. 9, e49758, 2020.
- [28] R. Ratcliff, P. L. Smith, and G. McKoon, “Modeling regularities in response time and accuracy data with the diffusion model,” *Curr. Dir. Psychol. Sci.*, vol. 24, no. 6, pp. 458–470, 2015.
- [29] C. D. Salzman, K. H. Britten, and W. T. Newsome, “Cortical microstimulation influences perceptual judgements of motion direction,” *Nature*, vol. 346, no. 6280, pp. 174–177, 1990.
- [30] M. N. Shadlen and W. T. Newsome, “Neural basis of a perceptual decision in the parietal cortex (area lip) of the rhesus monkey,” *J. Neurophysiol.*, vol. 86, no. 4, pp. 1916–1936, 2001.
- [31] J. D. Roitman and M. N. Shadlen, “Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task,” *J. Neurosci.*, vol. 22, no. 21, pp. 9475–9489, 2002.
- [32] J. Rigosa, A. Lucantonio, G. Noselli, A. Fassihi, E. Zorzini, F. Manzano, F. Pulecchi, and M. E. Diamond, “Dye-enhanced visualization of rat whiskers for behavioral studies,” *eLife*, vol. 6, e25290, 2017.

- [33] W. Gerstner, W. M. Kistler, R. Naud, and L. Paninski, *Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition*. Cambridge University Press, Cambridge, 2014.
- [34] A. Insabato, L. Dempere-Marco, M. Pannunzi, G. Deco, and R. Romo, “The influence of spatiotemporal structure of noisy stimuli in decision making,” *PLoS Comput. Biol.*, vol. 10, no. 4, e1003492, 2014.
- [35] K. Wimmer, A. Compte, A. Roxin, D. Peixoto, A. Renart, and J. De La Rocha, “Sensory integration dynamics in a hierarchical network explains choice probabilities in cortical area mt,” *Nat. Commun.*, vol. 6, no. 1, p. 6177, 2015.
- [36] A. E. Urai, A. Braun, and T. H. Donner, “Pupil-linked arousal is driven by decision uncertainty and alters serial choice bias,” *Nat. Commun.*, vol. 8, no. 1, p. 14637, 2017.
- [37] A. Zylberberg, P. Barttfeld, and M. Sigman, “The construction of confidence in a perceptual decision,” *Front. Integr. Neurosci.*, vol. 6, p. 79, 2012.
- [38] R. Kiani, T. D. Hanks, and M. N. Shadlen, “Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment,” *J. Neurosci.*, vol. 28, no. 12, pp. 3017–3029, 2008.
- [39] D. Peixoto, J. R. Verhein, R. Kiani, J. C. Kao, P. Nuyujukian, C. Chandrasekaran, J. Brown, S. Fong, S. I. Ryu, K. V. Shenoy, *et al.*, “Decoding and perturbing decision states in real time,” *Nature*, vol. 591, no. 7851, pp. 604–609, 2021.
- [40] B. W. Brunton, M. M. Botvinick, and C. D. Brody, “Rats and humans can optimally accumulate evidence for decision-making,” *Science*, vol. 340, no. 6128, pp. 95–98, 2013.
- [41] S. A. Los, “Identifying stimuli of different perceptual categories in mixed blocks of trials: Evidence for cost in switching between computational processes,” *Journal of Experimental Psychology: Human Perception and Performance*, vol. 25, no. 1, p. 3, 1999.
- [42] A. C. Barato and U. Seifert, “Thermodynamic uncertainty relation for biomolecular processes,” *Phys. Rev. Lett.*, vol. 114, p. 158101, 15 Apr. 2015.
- [43] G. Falasco and M. Esposito, “Dissipation-time uncertainty relation,” *Phys. Rev. Lett.*, vol. 125, p. 120604, 12 Sep. 2020.

- [44] J. Horowitz and T. Gingrich, “Thermodynamic uncertainty relations constrain non-equilibrium fluctuations,” *Nat. Phys.*, vol. 16, pp. 15–20, 2020.
- [45] M. Dorpinghaus, I. Neri, E. Roldán, and F. Julicher, “Optimal information usage in binary sequential hypothesis testing,” *Theory of Probability & Its Applications*, vol. 68, no. 1, pp. 77–87, 2023.
- [46] M. E. Diamond and A. Toso, “Tactile cognition in rodents,” *Neuroscience & Biobehavioral Reviews*, p. 105161, 2023.
- [47] A. Toso, A. Fassihi, L. Paz, F. Pulecchi, and M. E. Diamond, “A sensory integration account for time perception,” *PLoS Computational Biology*, vol. 17, no. 1, e1008668, 2021.
- [48] I. Hachen, S. Reinartz, R. Brasselet, A. Stroligo, and M. Diamond, “Dynamics of history-dependent perceptual judgment,” *Nature communications*, vol. 12, no. 1, p. 6036, 2021.
- [49] A. Fassihi, Y. Zuo, and M. E. Diamond, “Making sense of sensory evidence in the rat whisker system,” *Current Opinion in Neurobiology*, vol. 60, pp. 76–83, 2020.
- [50] Y. Zuo and M. E. Diamond, “Texture identification by bounded integration of sensory cortical signals,” *Current Biology*, vol. 29, no. 9, pp. 1425–1435, 2019.
- [51] V. Esmaili and M. E. Diamond, “Neuronal correlates of tactile working memory in prefrontal and vibrissal somatosensory cortex,” *Cell reports*, vol. 27, no. 11, pp. 3167–3181, 2019.
- [52] A. Fassihi, A. Akrami, F. Pulecchi, V. Schönfelder, and M. E. Diamond, “Transformation of perception from sensory to motor cortex,” *Current Biology*, vol. 27, no. 11, pp. 1585–1596, 2017.
- [53] F. Balci, P. Simen, R. Niyogi, A. Saxe, J. A. Hughes, P. Holmes, and J. D. Cohen, “Acquisition of decision making criteria: Reward rate ultimately beats accuracy,” *Attention, Perception, & Psychophysics*, vol. 73, pp. 640–657, 2011.
- [54] R. Bogacz, E. Brown, J. Moehlis, P. Holmes, and J. D. Cohen, “The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks,” *Psychological review*, vol. 113, no. 4, p. 700, 2006.

- [55] R. Bogacz, P. T. Hu, P. J. Holmes, and J. D. Cohen, “Do humans produce the speed–accuracy trade-off that maximizes reward rate?” *Quarterly journal of experimental psychology*, vol. 63, no. 5, pp. 863–891, 2010.
- [56] J. J. Starns and R. Ratcliff, “The effects of aging on the speed–accuracy compromise: Boundary optimality in the diffusion model,” *Psychology and aging*, vol. 25, no. 2, p. 377, 2010.
- [57] M. Zacksenhouse, R. Bogacz, and P. Holmes, “Robust versus optimal strategies for two-alternative forced choice tasks,” *Journal of mathematical psychology*, vol. 54, no. 2, pp. 230–246, 2010.
- [58] K.-F. Wong and A. C. Huk, “Temporal dynamics underlying perceptual decision making: Insights from the interplay between an attractor model and parietal neurophysiology,” *Front. Neurosci.*, vol. 2, p. 383, 2008.
- [59] X.-J. Wang, “Probabilistic decision making by slow reverberation in cortical circuits,” *Neuron*, vol. 36, no. 5, pp. 955–968, 2002.
- [60] K.-F. Wong and X.-J. Wang, “A recurrent network mechanism of time integration in perceptual decisions,” *J. Neurosci.*, vol. 26, no. 4, pp. 1314–1328, 2006.
- [61] C. W. Lynn, E. J. Cornblath, L. Papadopoulos, M. A. Bertolero, and D. S. Bassett, “Broken detailed balance and entropy production in the human brain,” *Proc. Natl. Acad. Sci. U.S.A.*, vol. 118, no. 47, e2109889118, 2021.
- [62] E. G-Guzmán, Y. S. Perl, J. Vohryzek, A. Escrichs, D. Manasova, B. Türker, E. Tagliazucchi, M. Kringelbach, J. D. Sitt, and G. Deco, “The lack of temporal brain dynamics asymmetry as a signature of impaired consciousness states,” *Interface Focus*, vol. 13, no. 3, p. 20220086, 2023.
- [63] L. A. de la Fuente, F. Zamberlan, H. Bocaccio, M. Kringelbach, G. Deco, Y. S. Perl, C. Pallavicini, and E. Tagliazucchi, “Temporal irreversibility of neural dynamics as a signature of consciousness,” *Cereb. Cortex*, vol. 33, no. 5, pp. 1856–1865, 2023.
- [64] J. Cruzat, R. Herzog, P. Prado, Y. Sanz-Perl, R. Gonzalez-Gomez, S. Moguilner, M. L. Kringelbach, G. Deco, E. Tagliazucchi, and A. Ibañez, “Temporal irreversibility of large-

- scale brain dynamics in Alzheimer’s disease,” *J. Neurosci.*, vol. 43, no. 9, pp. 1643–1656, 2023.
- [65] M. Zanin, B. Güntekin, T. Aktürk, L. Hanoğlu, and D. Papo, “Time irreversibility of resting-state activity in the healthy brain and pathology,” *Frontiers in physiology*, vol. 10, p. 1619, 2020.
- [66] M. Kleiner, D. Brainard, and D. Pelli, “What’s new in psychtoolbox-3?” *Perception*, 2007.
- [67] D. Bates, M. Mächler, B. Bolker, and S. Walker, “Fitting linear mixed-effects models using lme4,” *arXiv preprint arXiv:1406.5823*, 2014.
- [68] K. Bartoń, “Mumin: Multi-model inference,” *R package version 1.46. 0.*, 2022.
- [69] F. E. Satterthwaite, “An approximate distribution of estimates of variance components,” *Biometrics bulletin*, vol. 2, no. 6, pp. 110–114, 1946.
- [70] A. Kuznetsova, P. B. Brockhoff, and R. H. Christensen, “Lmertest package: Tests in linear mixed effects models,” *Journal of statistical software*, vol. 82, pp. 1–26, 2017.
- [71] R. Lenth, “Emmeans: Estimated marginal means, aka least-squares means...,” *R package version 1.8. 5*, 2023.
- [72] U. Seifert, “Entropy production along a stochastic trajectory and an integral fluctuation theorem,” *Phys. Rev. Lett.*, vol. 95, no. 4, p. 040 602, 2005.
- [73] T. R. Gingrich, J. M. Horowitz, N. Perunov, and J. L. England, “Dissipation bounds all steady-state current fluctuations,” *Phys. Rev. Lett.*, vol. 116, p. 120 601, 12 Mar. 2016.
- [74] I. Neri, “Universal tradeoff relation between speed, uncertainty, and dissipation in nonequilibrium stationary states,” *SciPost Phys.*, 2021.
- [75] P. Neri, A. J. Parker, and C. Blakemore, “Probing the human stereoscopic system with reverse correlation,” *Nature*, vol. 401, no. 6754, pp. 695–698, 1999.
- [76] H. Nienborg and B. G. Cumming, “Decision-related activity in sensory neurons reflects more than a neuron’s causal effect,” *Nature*, vol. 459, no. 7243, pp. 89–92, 2009.
- [77] B. W. Brunton, M. M. Botvinick, and C. D. Brody, “Rats and humans can optimally accumulate evidence for decision-making,” *Science*, vol. 340, no. 6128, pp. 95–98, 2013. DOI: 10.1126/science.1233912.

- [78] G. Okazawa, L. Sha, B. A. Purcell, and R. Kiani, “Psychophysical reverse correlation reflects both sensory and decision-making processes,” *Nature communications*, vol. 9, no. 1, p. 3479, 2018.
- [79] A. E. Urai, J. W. de Gee, K. Tsetsos, and T. H. Donner, “Choice history biases subsequent evidence accumulation,” *eLife*, vol. 8, e46331, 2019. DOI: 10.7554/eLife.46331.
- [80] M. Glickman, R. Moran, and M. Usher, “Evidence integration and decision confidence are modulated by stimulus consistency,” *Nature Human Behaviour*, vol. 6, no. 7, pp. 988–999, 2022.
- [81] T. E. J. Behrens, M. W. Woolrich, M. E. Walton, and M. F. S. Rushworth, “Learning the value of information in an uncertain world,” *Nature Neuroscience*, vol. 10, pp. 1214–1221, 2007. DOI: 10.1038/nn1954.
- [82] M. R. Nassar, R. C. Wilson, B. Heasly, and J. I. Gold, “An approximately bayesian delta-rule model explains the dynamics of belief updating in a changing environment,” *The Journal of Neuroscience*, vol. 30, pp. 12 366–12 378, 2010. DOI: 10.1523/JNEUROSCI.0822-10.2010.
- [83] O. Ossmy, R. Moran, T. Pfeffer, K. Tsetsos, M. Usher, and T. H. Donner, “The timescale of perceptual evidence integration can be adapted to the environment,” *Current Biology*, vol. 23, pp. 981–986, 2013. DOI: 10.1016/j.cub.2013.04.039.
- [84] S. Monsell, “Task switching,” *Trends in Cognitive Sciences*, vol. 7, no. 3, pp. 134–140, 2003. DOI: 10.1016/S1364-6613(03)00028-7.
- [85] A. Kiesel, M. Steinhauser, M. Wendt, M. Falkenstein, K. Jost, A. M. Philipp, and I. Koch, “Control and interference in task switching—a review,” *Psychological Bulletin*, vol. 136, no. 5, pp. 849–874, 2010. DOI: 10.1037/a0019842.
- [86] C. Summerfield and T. Egnér, “Expectation (and attention) in visual cognition,” *Trends in Cognitive Sciences*, vol. 13, no. 9, pp. 403–409, 2009. DOI: 10.1016/j.tics.2009.06.003.
- [87] G. Rohenkohl, A. M. Cravo, V. Wyart, and A. C. Nobre, “Temporal expectation improves the quality of sensory information,” *The Journal of Neuroscience*, vol. 32, no. 24, pp. 8424–8428, 2012. DOI: 10.1523/JNEUROSCI.0804-12.2012.

- [88] U. Seifert, *Stochastic thermodynamics: principles and perspectives*. Springer, 2008, vol. 64, pp. 423–431. DOI: 10.1140/epjb/e2008-00001-9.
- [89] A. Durmaz, Y. Sarmiento, G. Fortunato, D. Das, M. E. Diamond, D. Buetti, and É. Roldán, “Human perceptual decision making of nonequilibrium fluctuations,” *arXiv preprint arXiv:2311.12692*, 2023.
- [90] A. L. Lee and H. Lu, “A comparison of global motion perception using a multiple-aperture stimulus,” *Journal of Vision*, vol. 10, no. 4, pp. 9–9, 2010.
- [91] D. Burr and P. Thompson, “Motion psychophysics: 1985–2010,” *Vision research*, vol. 51, no. 13, pp. 1431–1456, 2011.
- [92] J. Drugowitsch, R. Moreno-Bote, A. K. Churchland, M. N. Shadlen, and A. Pouget, “The cost of accumulating evidence in perceptual decision making,” *The Journal of Neuroscience*, vol. 32, no. 11, pp. 3612–3628, 2012. DOI: 10.1523/JNEUROSCI.4010-11.2012. [Online]. Available: <https://pubmed.ncbi.nlm.nih.gov/22423085/>.
- [93] G. Felsen and Y. Dan, “A natural approach to studying vision,” *Nature Neuroscience*, vol. 8, no. 12, pp. 1643–1646, 2005. DOI: 10.1038/nn1608. [Online]. Available: <https://pubmed.ncbi.nlm.nih.gov/16306891/>.
- [94] R. Metzler and J. Klafter, “The random walk’s guide to anomalous diffusion: A fractional dynamics approach,” *Physics Reports*, vol. 339, no. 1, pp. 1–77, 2000. DOI: 10.1016/S0370-1573(00)00070-3. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0370157300000703>.
- [95] R. G. O’Connell, P. M. Dockree, and S. P. Kelly, “A supramodal accumulation-to-bound signal that determines perceptual decisions in humans,” *Nature Neuroscience*, vol. 15, no. 12, pp. 1729–1735, 2012. DOI: 10.1038/nn.3248. [Online]. Available: <https://doi.org/10.1038/nn.3248>.
- [96] D. M. Twomey, P. R. Murphy, S. P. Kelly, and R. G. O’Connell, “The classic P300 encodes a build-to-threshold decision variable,” *European Journal of Neuroscience*, vol. 42, no. 1, pp. 1636–1643, 2015. DOI: 10.1111/ejn.12936. [Online]. Available: <https://doi.org/10.1111/ejn.12936>.

- [97] M. K. Van Vugt, P. Simen, L. Nystrom, P. Holmes, and J. D. Cohen, “Lateralized readiness potentials reveal properties of a neural mechanism for implementing a decision threshold,” *PloS one*, vol. 9, no. 3, e90943, 2014.
- [98] V. Wyart, V. de Gardelle, J. Scholl, and C. Summerfield, “Rhythmic fluctuations in evidence accumulation during decision making in the human brain,” *Neuron*, vol. 76, no. 4, pp. 847–858, 2012. DOI: 10.1016/j.neuron.2012.09.015. [Online]. Available: <https://doi.org/10.1016/j.neuron.2012.09.015>.
- [99] G. Aston-Jones and J. D. Cohen, “An integrative theory of locus coeruleus-norepinephrine function: Adaptive gain and optimal performance,” *Annual Review of Neuroscience*, vol. 28, pp. 403–450, 2005. DOI: 10.1146/annurev.neuro.28.061604.135709. [Online]. Available: <https://doi.org/10.1146/annurev.neuro.28.061604.135709>.
- [100] D. Thura, J. Beauregard-Racine, C.-W. Fradet, and P. Cisek, “Decision making by urgency gating: Theory and experimental support,” *Journal of neurophysiology*, vol. 108, no. 11, pp. 2912–2930, 2012.
- [101] M. A. Carland, D. Thura, and P. Cisek, “The urge to decide and act: Implications for brain function and dysfunction,” *The Neuroscientist*, vol. 25, no. 5, pp. 491–511, 2019.
- [102] S. Redner, *A Guide to First-Passage Processes*. Cambridge university press, Cambridge, 2001.

Appendix A

Supplementary Materials

Sequential Probability Ratio Test on the Direction of Time's Arrow

Consider a time series of snapshots characterized by a nonequilibrium stationary stochastic process X_t . By observing this evidence, one must make a decision D_{dec} by choosing between two competing hypotheses, each corresponding to a physical model of the process, e.g., $H_1 = 1$ and $H_2 = -1$.

When judging such a stochastic process repeatedly across a sequence of independent trials, how should one process the evidence to optimally achieve the minimal decision time (averaged across trials) for a given prescribed level of accuracy?

The first rigorous answer to this question emerges from Wald's SPRT:

the optimal decision time T_{dec} to select $D_{\text{dec}} = 1$ ($D_{\text{dec}} = -1$) is when the cumulative log-likelihood ratio $\mathcal{L}_t = \ln[\mathcal{P}(X_{[0,t]}|H_1)/\mathcal{P}(X_{[0,t]}|H_2)]$ first exceeds (first falls below) a fixed decision threshold L_+ (L_-) determined by the prescribed accuracy [21], [23]. Here, the decision accuracies are the conditional probabilities $\alpha_+ = P(D_{\text{dec}} = 1|H_1)$ and $\alpha_- = P(D_{\text{dec}} = -1|H_2)$, and denote, respectively, the following. The probability that the decision $D_{\text{dec}} = 1$ ($D_{\text{dec}} = -1$) when the hypothesis H_1 (H_2) is indeed true. On the other hand, the quantity $\mathcal{P}(X_{[0,t]}|H_1)$ ($\mathcal{P}(X_{[0,t]}|H_2)$) denotes the conditional probability of observing the trajectory $X_{[0,t]}$ given hypothesis H_1 (H_2) is true. For a broad class of continuous stochastic processes, the decision thresholds ensuring symmetric accuracies $\alpha_{\pm} = \alpha$ are given by $L_{\pm} = \pm\lambda$ with $\lambda = \ln[\alpha/(1 - \alpha)]$ [24]. Wald's SPRT is thus a first-passage-time problem for \mathcal{L}_t with the decision time $T_{\text{dec}} = \inf\{t \geq 0 | \mathcal{L}_t \notin (-\lambda, \lambda)\}$ given by the first-exit time from the interval $(-\lambda, \lambda)$.

In thermodynamics, an enduring question is how the time irreversibility of nonequilibrium processes gives rise to entropy production in the form of, e.g., heat dissipation. Consider the paradigmatic 1D drift-diffusion model (1DDM) for an overdamped Brownian particle in one dimension $\dot{X}_t = v + \sqrt{2D}\xi_t$, where X_t denotes the particle's position at time t ; v and D are the drift velocity and the diffusion constant, respectively. The quantity ξ_t is a zero-mean Gaussian white noise with autocorrelation $\langle \xi_t \xi_s \rangle = \delta(t - s)$. Suppose the task of an SPRT is to decide as soon as possible, achieving some prescribed average accuracy across trials, whether an observed stochastic path $X_{[0,t]}$ (with initial position $X_0 = 0$ and periodic boundary conditions) is produced by a 1DDM with positive speed $v > 0$ (hypothesis H_1) or instead by a time-reversed 1DDM with negative speed $v < 0$ (hypothesis H_2). The log-likelihood ratio associated with such SPRT reads

$$\mathcal{L}_t = \ln \left[\frac{\mathcal{P}(X_{[0,t]}|H_1)}{\mathcal{P}(X_{[0,t]}|H_2)} \right] = \left(\frac{v}{D} \right) X_t, \quad (\text{S1})$$

and coincides with the stochastic entropy production [72] associated with the stochastic trajectory $X_{[0,t]}$ [13], which for the 1DDM is proportional to the position. Note that we have $\mathcal{P}(X_{[0,t]}|H_2) = \mathcal{P}(\tilde{X}_{[0,t]}|H_1)$, where $\mathcal{P}(\tilde{X}_{[0,t]}|H_1)$ is the probability of observing the time-reversed trajectory $\tilde{X}_{[0,t]}$ under hypothesis H_1 ($v > 0$) true. Therefore \mathcal{L}_t for this problem encodes the degree of time-reversal asymmetry associated with the statistics of the model H_1 . These relations imply that the SPRT evidence accumulation rate is given by

$$\langle \dot{\mathcal{L}}_t \rangle = \frac{v \langle \dot{X}_t \rangle}{D} = \frac{v^2}{D}. \quad (\text{S2})$$

Interestingly, we have $\langle \dot{\mathcal{L}}_t \rangle = \Sigma$, where Σ is the steady-state rate of entropy production associated with the 1DDM (in units of the Boltzmann constant k_B). The latter is given by the steady-state average power divided by $k_B T$ with T being the temperature of the environment, i.e.,

$$\Sigma = \frac{\langle \dot{W} \rangle}{k_B T} = \frac{v^2}{D}. \quad (\text{S3})$$

In Eq. (S3), we use $\langle \dot{W} \rangle = \gamma v^2$ as the power dissipated by a constant friction force γv and Einstein's

relation $D = k_B T / \gamma$.

Furthermore, the SPRT on the direction of time's arrow maps into the problem of first exit time for X_t to reach any of the two decision thresholds located in L and $-L$ with $L = \lambda D / v = D \ln[\alpha_{\text{SPRT}} / (1 - \alpha_{\text{SPRT}})] / v$. Consequently, the decision time is given by $T_{\text{dec}} = \inf \{t \geq 0 | X_t \notin (-L, L)\}$, whose mean value retrieved from the first-passage theory [102] reads $\langle T_{\text{dec}} \rangle_{\text{SPRT}} = (L/v) \tanh(vL/2D)$. Using Eq. (S3) and introducing the dimensionless Péclet number

$$\text{Pe} \equiv \frac{vL}{D}, \quad (\text{S4})$$

we obtain the SPRT mean decision time

$$\langle T_{\text{dec}} \rangle_{\text{SPRT}} = \frac{\text{Pe}}{\Sigma} \tanh\left(\frac{\text{Pe}}{2}\right). \quad (\text{S5})$$

Moreover, we obtain that the predicted accuracy set by the absorption probability $\alpha_{\text{SPRT}} = P(X_{T_{\text{dec}}} = L)$ is equal to

$$\alpha_{\text{SPRT}} = \frac{1}{1 + \exp(-\text{Pe})}. \quad (\text{S6})$$

Note that Eqs. (S5) and (S6) reveal that the mean decision time and the accuracy in the SPRT depend only on two parameters, Σ and Pe .

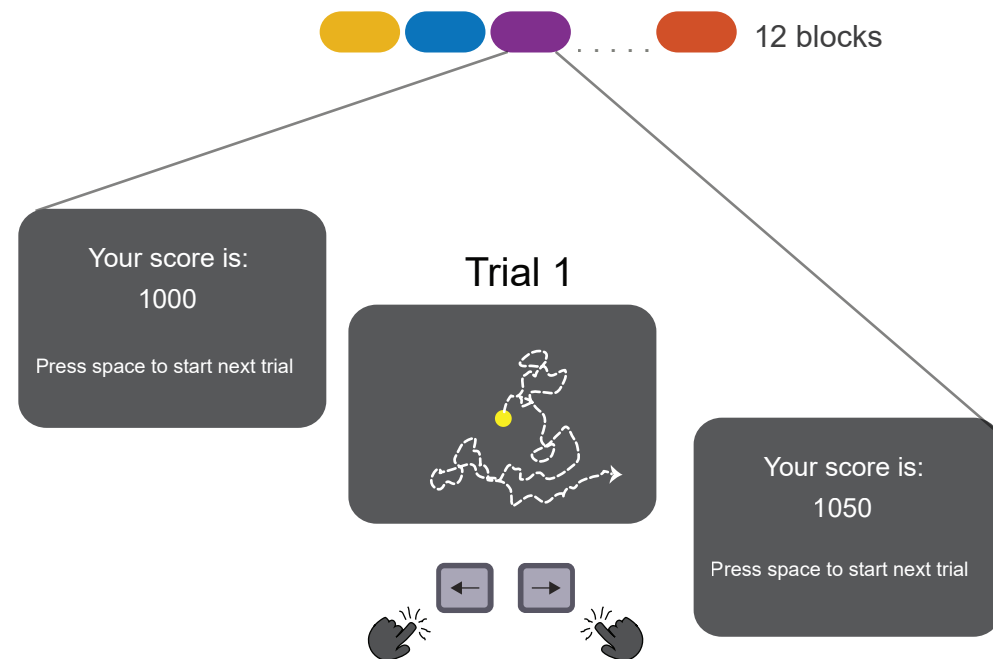
Let us now consider a suboptimal decision maker that collects stochastic evidence from the DDM and takes decisions based on a first-exit problem of \mathcal{J}_t , a functional of $X_{[0,t]}$ that may be in general different from the log-likelihood ratio (e.g. $\mathcal{J}_t = \int_0^t ds f(s) X_s$, with f a prescribed real function). Let $(-L_J, L_J)$ be the interval chosen such that the accuracy of the decision maker equals that of the SPRT [Eq. (S6)], leading to the mean decision time $\langle T_{\text{dec}} \rangle$. Under such assumptions, the SPRT's optimality implies the relation $\langle T_{\text{dec}} \rangle \geq \langle T_{\text{dec}} \rangle_{\text{SPRT}}$, which can be generalized to a broad class of

stochastic processes [24]. For the 1DDM, this relation together with Eq. (S5) implies that

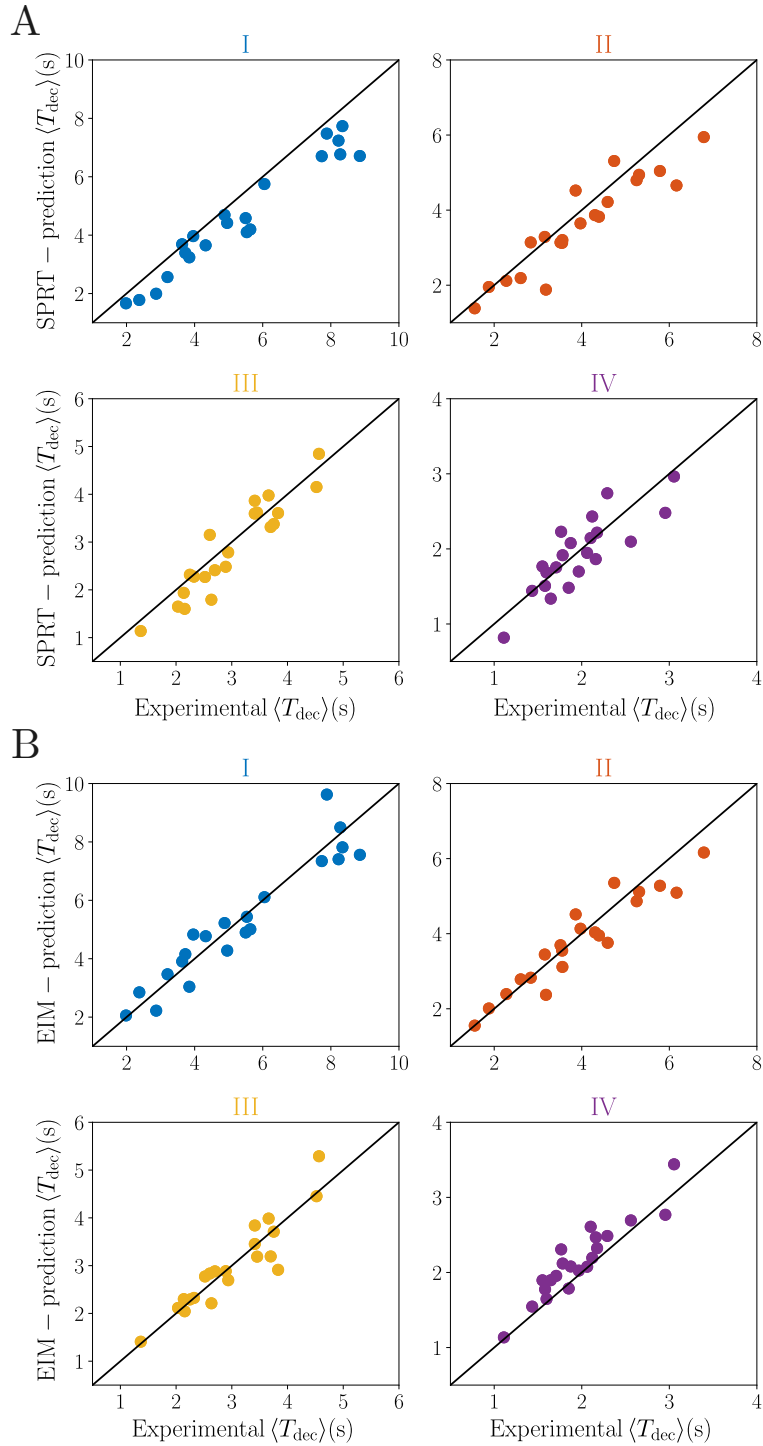
$$\Sigma \langle T_{\text{dec}} \rangle \geq \text{Pe} \tanh\left(\frac{\text{Pe}}{2}\right), \quad (\text{S7})$$

i.e. an inverse relation between dissipation (entropy production) and (mean decision) time Eq. (S7) belongs to the emerging collection of quantitative predictions of thermodynamic constraints on the accuracy of nonequilibrium currents in the burgeoning field of thermodynamic uncertainty relations (TURs) [42]–[44], [73]. A fundamental lower limit to the mean decision time of the form of Eq. (S7) was first conjectured for Markovian processes [13] and later proved and generalized to generic nonequilibrium stationary processes [74], but here tested in experimental scenarios for the first time.

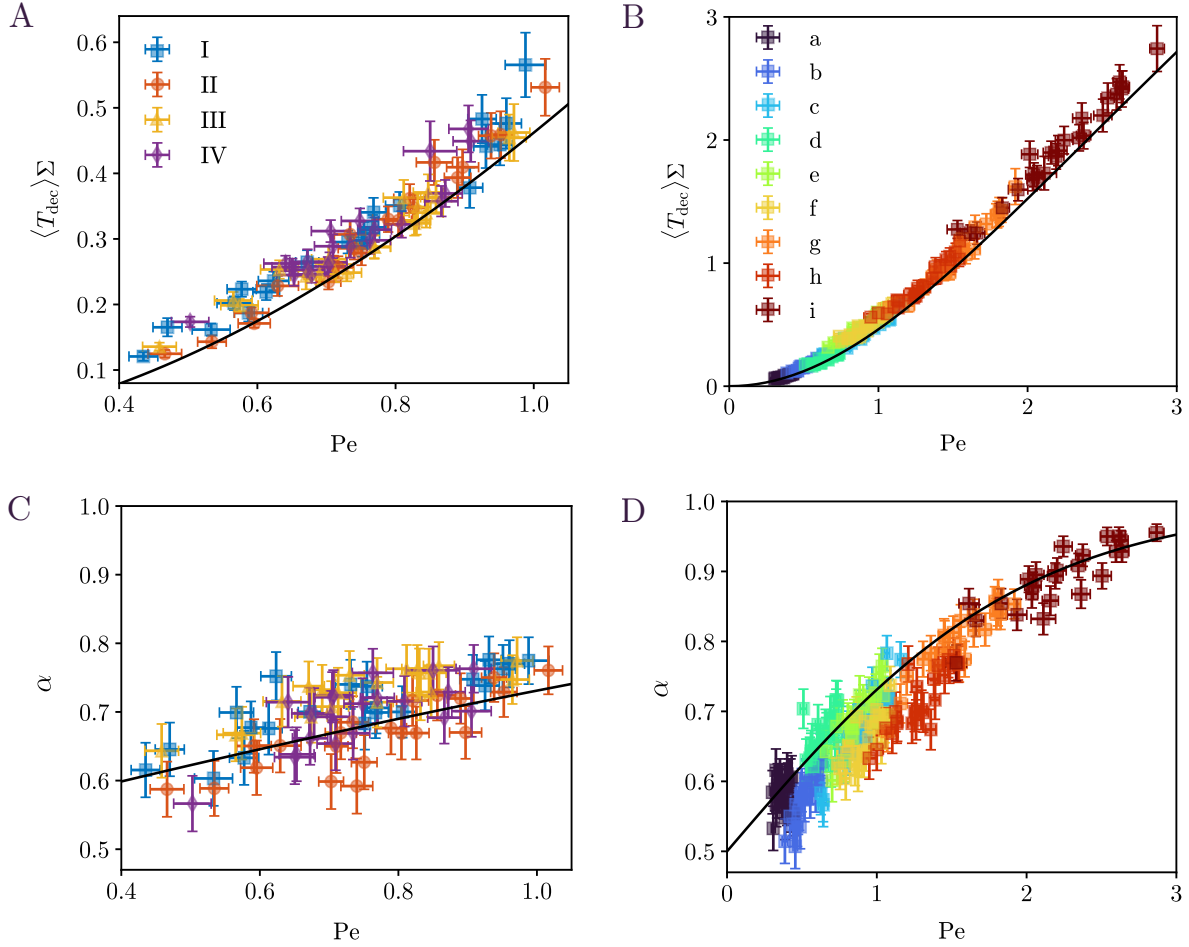
Supplementary Figures



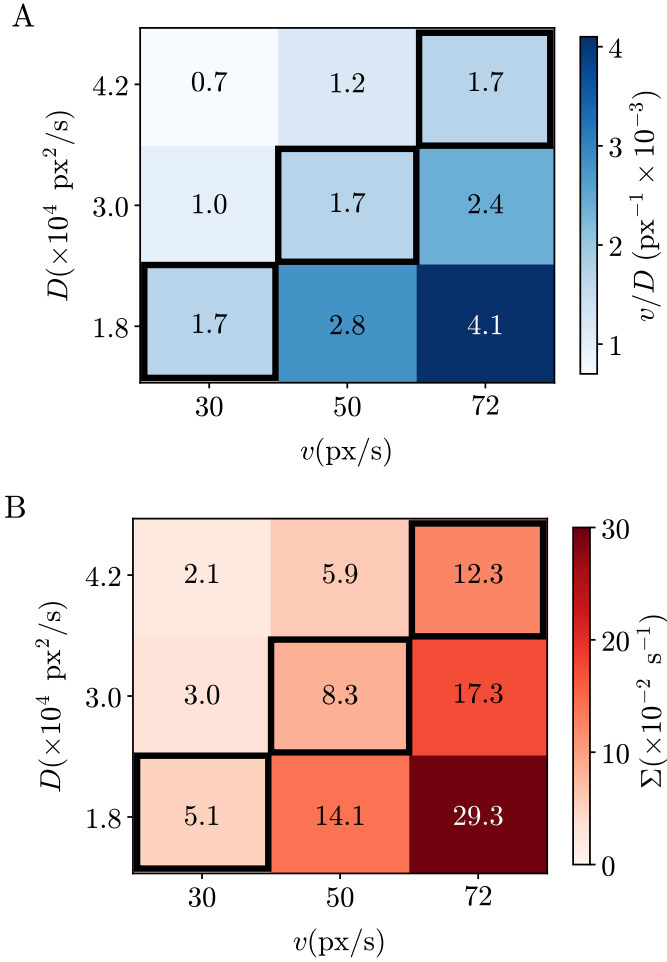
Supplementary Figure S1: **Schematics of visual displays in Experiment A.** Each of the twelve blocks consists of trials from the same condition. For each condition, the trials are shown in three blocks. In the beginning of each block, participants start with 1000 points. The score is accumulated within each block as participants give correct, incorrect, or missed responses, and in the next block participants start with 1000 points afresh.



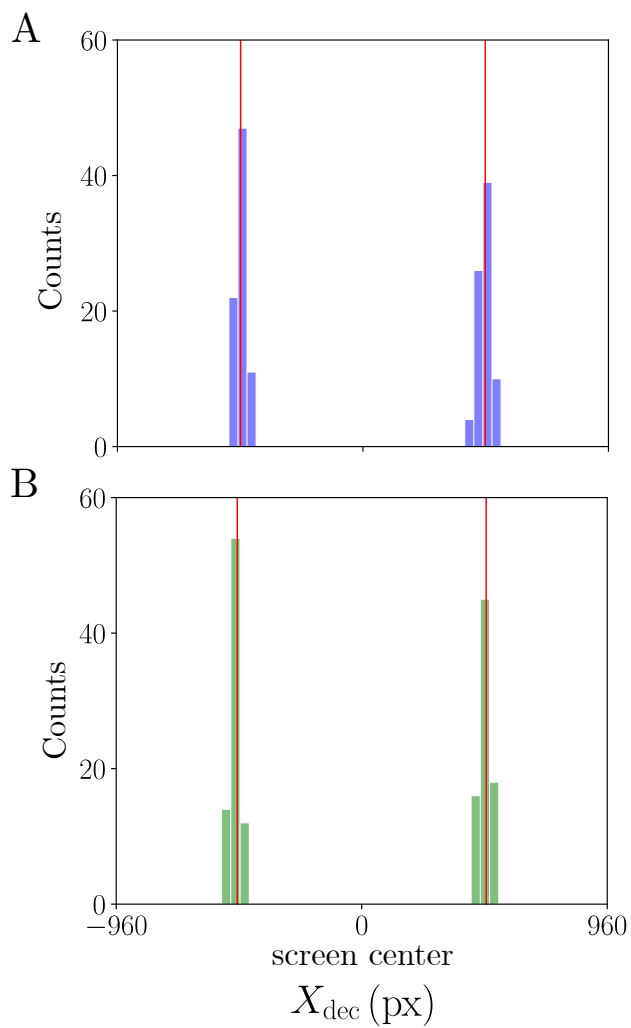
Supplementary Figure S2: **Cross validation of the model predictions of T_{dec} for (A) SPRT, and (B) EIM in Experiment A.** We find that the average experimental decision time of the training data accurately matches the predicted decision time of the test data for both (A) SPRT and (B) EIM. Each data point represents participants and the lines the best prediction.



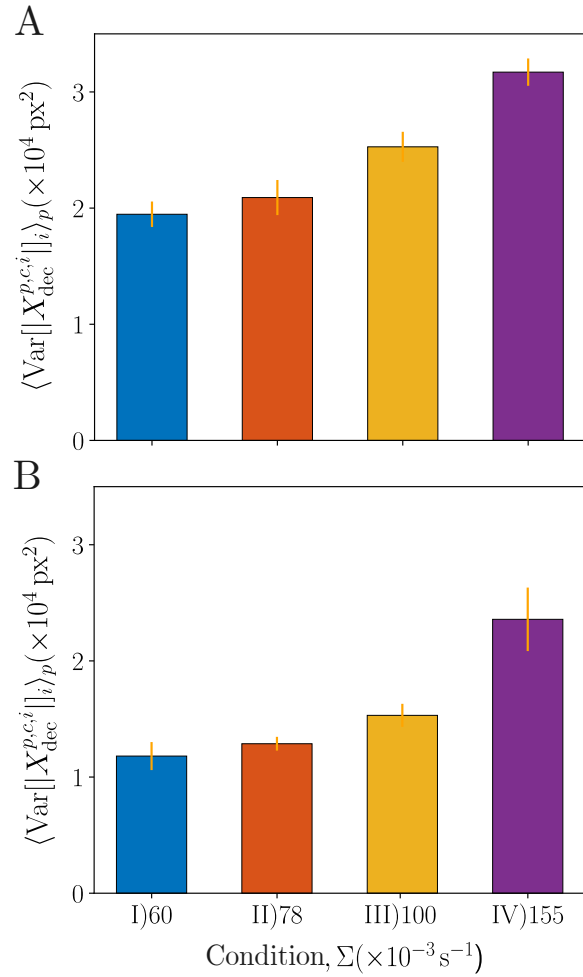
Supplementary Figure S3: **Experimental test of the stochastic-thermodynamics relations Eq. (2.11) and (2.10) in human perceptual decision making.** (A, B) Mean decision time (scaled by the rate of dissipation $\Sigma = v^2/D$) as a function of the participants' effective Péclet number for Experiment A (A), and Experiment B (B). The black line is given by $Pe \tanh(Pe/2)$ [Eq. (2.11)]. (C, D) Mean accuracy as a function of the participants' effective Péclet number for Experiment A (C), and Experiment B (D). The black line is given by the right-hand side in Eq. (2.10). (A-D) Each symbol is obtained from the statistics of a single participant, and different shapes and colors are used to denote different conditions. The effective Péclet number for each participant is estimated from the mean absolute value of the participant's decision threshold in each condition as $Pe = v\langle |X_{\text{dec}}| \rangle / D$. The error bars correspond to SEM.




Supplementary Figure S4: **Parameter values for conditions in Experiment B.**(A) Parameter values for drift velocity v , diffusivity D , and their ratio used in the nine experimental conditions in Experiment B. Each condition is a combination of the v and D of the corresponding column and row of the table, shown with its v/D value color-coded with its numerical value shown in text. The diagonal conditions with the black frame refer to the conditions with a fixed value v/D ($\simeq 1.69 \times 10^{-3} \text{ px}^{-1}$), and these three conditions have parameters similar to that of Experiment A. (B) Same experimental conditions but numerical values and color-code show the corresponding entropy production rate Σ for each experimental condition.




Supplementary Figure S5: **Control experiment results of an example participant.** (A) Control experiment, no diffusion version. The variance in the x -axis position of the disk at the moment of response is shown. Participants are instructed to press the correspondent key (right/left) as soon as the disk reaches either of the visible red lines that are 500 pixels away from the center. The x -axis represents the horizontal axis of the screen with 1920 pixels. The red symmetrical lines are the visual boundaries that the participants see throughout the whole trial. The number of observations of different positions of the disk at the response moment is shown in the y -axis. (B) Control experiment, Y -axis diffusion version. The variance in the x -axis position of the disk at the moment of response is shown. All participants show similar behavior.



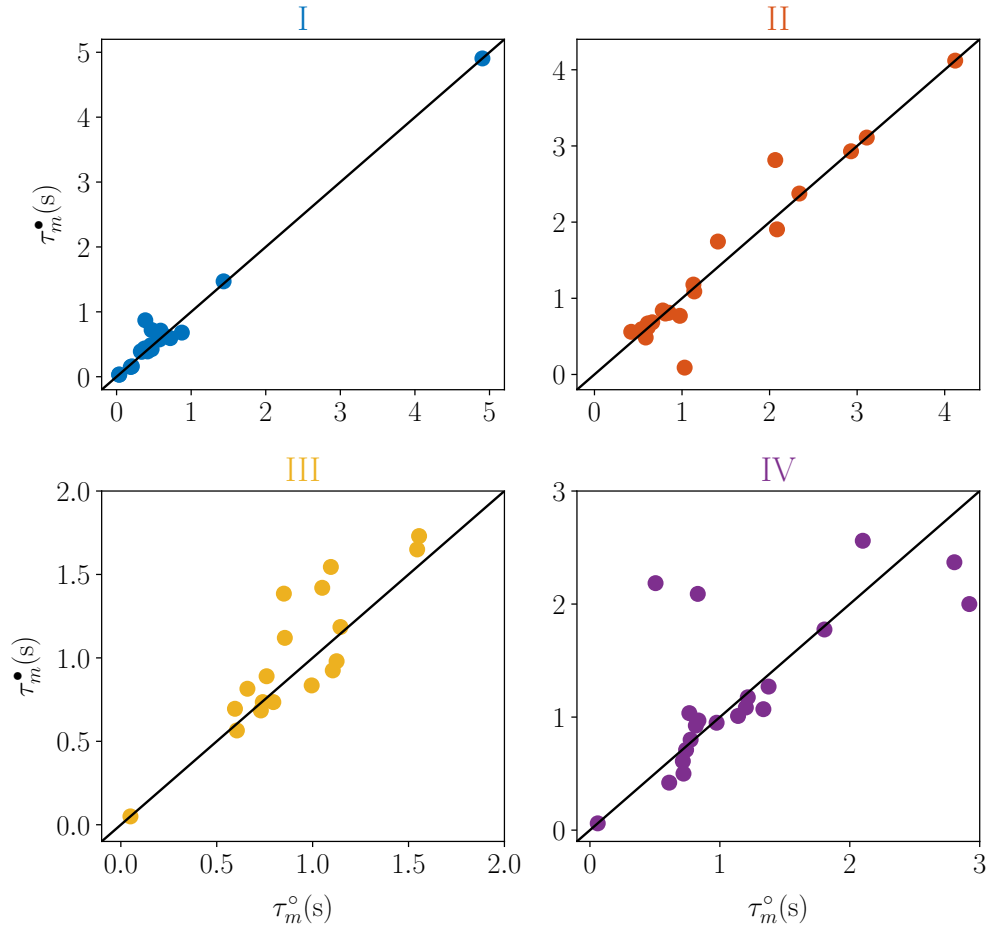
Supplementary Figure S6: **Variance in the X_{dec} .** (A) The variance of X_{dec} averaged over participants across different experimental conditions in Experiment A. The variance in the decision threshold increases with entropy production rate. (B) Experimental results from the control experiment original parameters version. Participants are instructed to press the decision key as soon as the disk reaches visible red lines on the screen, representing Wald's SPRT-like fixed thresholds. Even though overall variance is decreased as opposed to Experiment A in panel A, there is still some variance, and it varies across conditions in a similar pattern as in panel A.



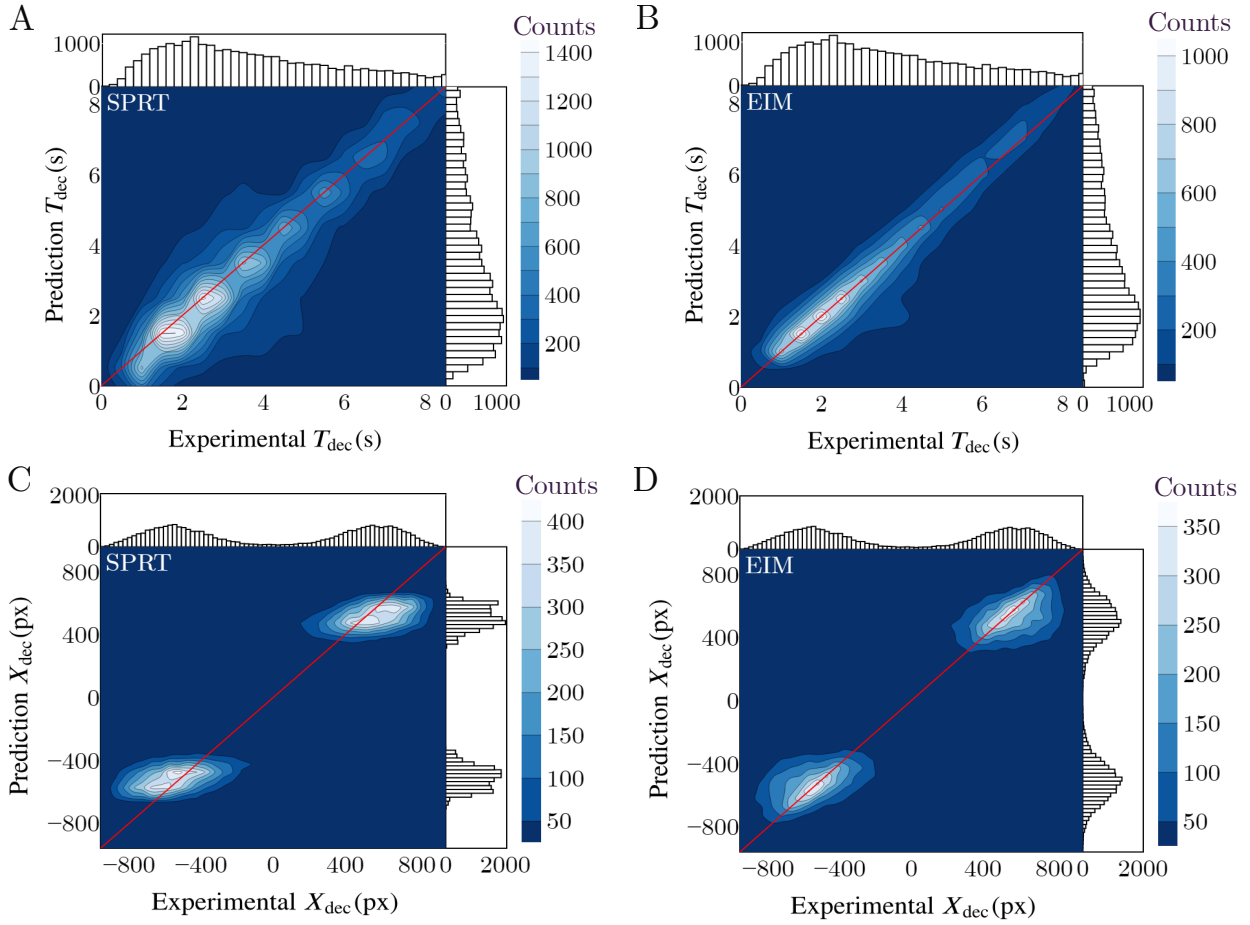
Supplementary Figure S7: **Participants respond consistently for the same trials.** Given the same trial (i), we calculate the normalized change in each participant's decision threshold with respect to their mean. Each data point represents the mean of the normalized change over all participants in Experiment A. Error bars show SEM. The fluctuating horizontal lines are the 95% confidence intervals.



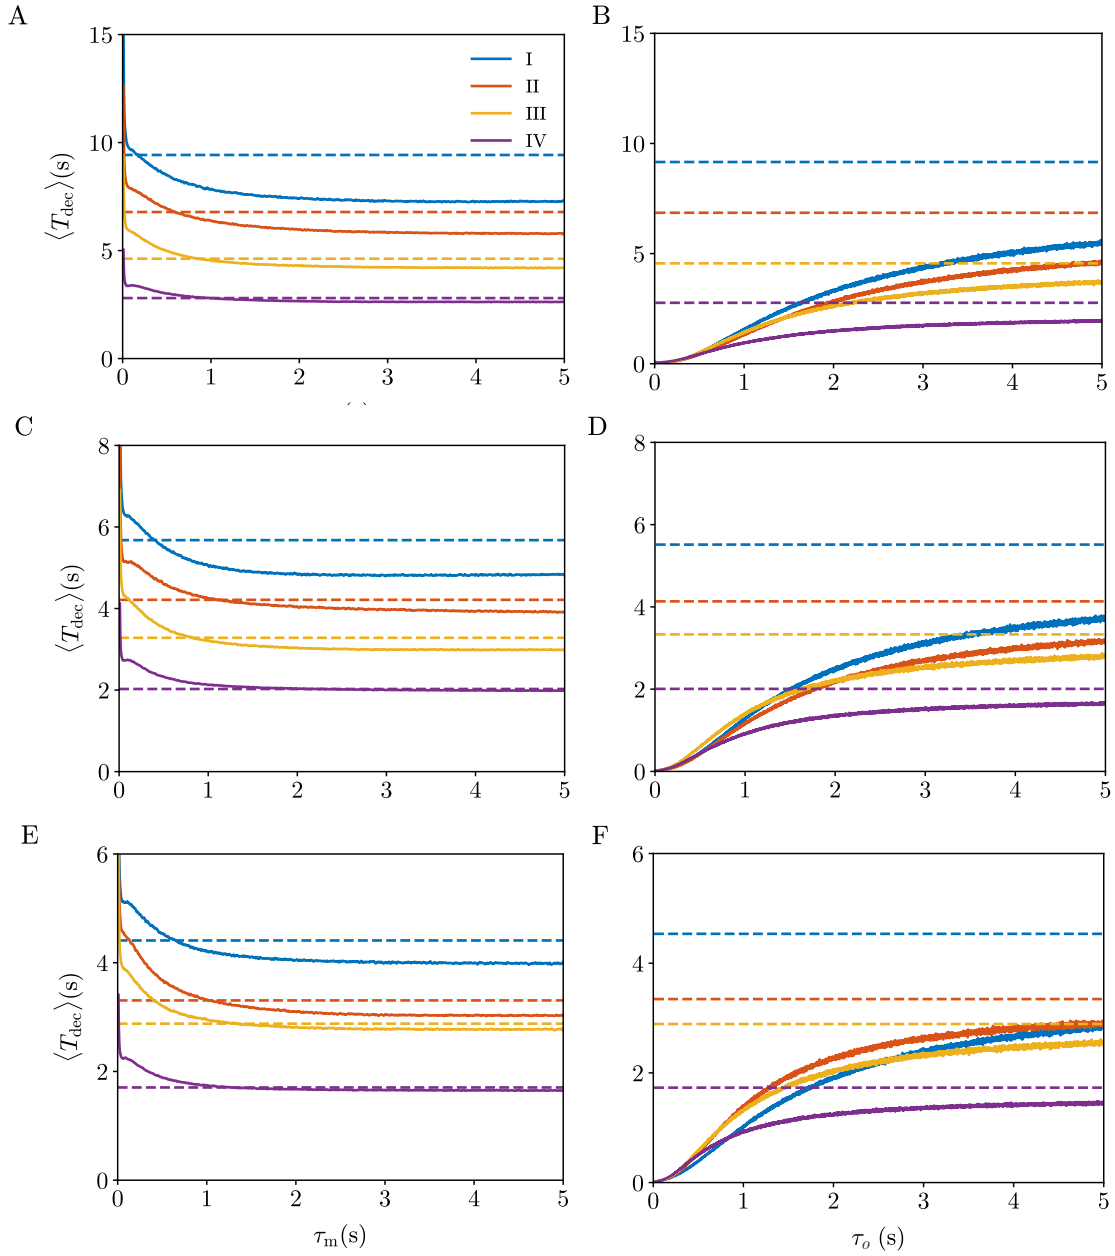
Supplementary Figure S8: **Probability density of $W_{\text{dec}} = W_{T_{\text{dec}}}$ in Experiment A.** The histograms show the probability density of W_{dec} obtained from the trial-by-trial analysis of a single participant using Eq. (2.14) across experimental conditions. The dashed lines in each condition display the corresponding mean values $-L_W^*$ and L_W^* . Note that the mean decision thresholds of EIM (L_W^*) change substantially with experimental conditions.



Supplementary Figure S9: **Cross validation of the integration time window parameter.** Conditions I (top left), II (top right), III (bottom left), and IV (Experiment A). We divide the trial-by-trial data into two equal-sized parts at random and find two memory relaxation times τ_m^* and τ_m° that accurately matches the average decision time for each respective portion of the experimental data. Points correspond to a scatter plot of these two parameters for each participant and lines with slope one to the ideal best fit.




Supplementary Figure S10: **Comparison of decision-making models (SPRT and EIM) with experimental decision times and decision thresholds for Experiment B.** (A-B) Joint histogram of the trial-by-trial experimental decision time T_{dec} for all participants and conditions, and the numerical predictions of T_{dec} by Wald's SPRT (A) and the EIM (B). (C-D) Joint histogram of the trial-by-trial experimental decision threshold X_{dec} for all participants and conditions, and the numerical predictions of X_{dec} by Wald's SPRT (C) and the EIM (D). The red lines have slope one, representing the region with perfect agreement between the experimental output and model predictions.

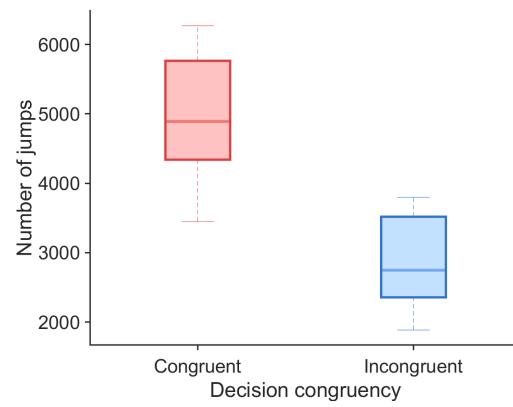


Supplementary Figure S11: **Inference of the optimal memory relaxation time τ_m^* for EIM (A, C, E) and τ_o^* for OU model (B,D,F) in Experiment A.** Three example participants are shown in the four conditions. Each dashed line corresponds to the experimental mean decision time. The solid lines denote the mean escape time of the evidence accumulator W_t from the interval $(-L_W, L_W)$ extracted from the numerical simulations as a function of the corresponding memory relaxation time.


A



B



Supplementary Figure S12: **Total number of decision-congruent and decision-incongruent jumps.** Box plots show the distribution of jump counts across participants, separated by whether the jump direction was consistent with the participant's response (decision-congruent) or not (decision-incongruent), for (A) Experiment A and (B) Experiment B.



Supplementary Figure S13: **Total number of decision-congruent and decision-incongruent jumps in Experiment C linear accumulation format.** Box plots show the distribution of jump counts across participants, separated by whether the jump direction was consistent with the participant's response (decision-congruent) or not (decision-incongruent)

Supplementary Tables

Condition	v (px/s)	D (px ² /s)	Σ (s ⁻¹)
I	36	21600	0.060
II	47	28200	0.078
III	60	36000	0.100
IV	93	55800	0.155

Supplementary Table S1: Experiment A, parameter values for the drift velocity v , diffusivity D , and entropy production rate $\Sigma = v^2/D$ used in the experimental conditions (I, II, III, and IV). The parameter values in Experiment A are chosen such that across conditions, the entropy production rate v^2/D varies, yet the ratio $v/D \simeq 1.67 \times 10^{-3} \text{ px}^{-1}$ remains constant.

	Sum Sq	Mean Sq	NumDF	DenDF	F value	p value
condition	1.8812	0.62706	3	60	1.2122	0.3131

Supplementary Table S2: Experiment A: X_{dec} - type III ANOVA on LME model estimates with Satterthwaite's method for degrees of freedom

	Sum Sq	Mean Sq	NumDF	DenDF	F value	p value
condition	0.049351	0.01645	3	60	25.954	6.901×10^{-11}

Supplementary Table S3: Experiment A: Accuracy – type III ANOVA on LME model estimates with Satterthwaite's method for degrees of freedom.

	estimate	SE	df	t stat	p value
I - II	0.045	0.008	60	5.804	<.0001
I - III	-0.022	0.008	60	-2.832	0.0377
I - IV	0.011	0.008	60	1.432	0.9447
II - III	-0.067	0.008	60	-8.636	<.0001
II - IV	-0.034	0.008	60	-4.372	0.0003
III - IV	0.033	0.008	60	4.264	0.0004

Supplementary Table S4: Multiple comparisons of condition pairs for $\text{Model}_A(\alpha)$. Degrees-of-freedom method: kenward-roger. P-value adjustment: bonferroni method for 6 tests.

	Sum Sq	Mean Sq	NumDF	DenDF	F value	p value
condition	1.9139	0.63797	3	60	170.26	$< 2.2 \times 10^{-16}$

Supplementary Table S5: Experiment A: T_{dec} – type III ANOVA on LME model estimates with Satterthwaite’s method for degrees of freedom.

	estimate	SE	df	t stat	p value
I - II	0.073	0.019	60	3.889	0.0015
I - III	0.178	0.019	60	9.431	$<.0001$
I - IV	0.401	0.019	60	21.198	$<.0001$
II - III	0.105	0.019	60	5.541	$<.0001$
II - IV	0.327	0.019	60	17.308	$<.0001$
III - IV	0.222	0.019	60	11.767	$<.0001$

Supplementary Table S6: Multiple comparisons of condition pairs for $\text{Model}_A(T_{\text{dec}})$. Degrees-of-freedom method: kenward-roger. P-value adjustment: bonferroni method for 6 tests.

	Sum Sq	Mean Sq	NumDF	DenDF	F value	p value
condition	0.19292	0.064305	3	63	10.297	1.319×10^{-5}

Supplementary Table S7: Experiment A: logit absolute error of accuracy (discrepancy between observed accuracy and SPRT predictions) – type III ANOVA on LME model estimates with Satterthwaite’s method for degrees of freedom.

	estimate	SE	df	t stat	p value
I - II	0.058	0.025	66.15	2.316	0.1420
I - III	-0.068	0.025	66.15	-2.712	0.0511
I - IV	0.038	0.025	66.15	1.518	0.8022
II - III	-0.126	0.025	66.15	-5.028	$<.0001$
II - IV	-0.020	0.025	66.15	-0.798	1.0000
III - IV	0.106	0.025	66.15	4.230	0.0004

Supplementary Table S8: Multiple comparisons of condition pairs for model AE_α in Experiment A. Degrees-of-freedom method: kenward-roger. P-value adjustment: bonferroni method for 6 tests.

	Sum Sq	Mean Sq	NumDF	DenDF	F value	p value
condition	0.11382	0.037939	3	63	18.274	1.212×10^{-8}

Supplementary Table S9: Experiment A: relative absolute error (AE) of decision times (discrepancy between observed T_{dec} and SPRT predictions) – type III ANOVA on LME model estimates with Satterthwaite’s method for degrees of freedom.

	estimate	SE	df	t stat	p value
I - II	0.086	0.014	66.15	5.984	<.0001
I - III	0.080	0.014	66.15	5.523	<.0001
I - IV	0.022	0.014	66.15	1.545	0.7620
II - III	-0.007	0.014	66.15	-0.461	1.0000
II - IV	-0.064	0.014	66.15	-4.438	0.0002
III - IV	-0.057	0.014	66.15	-3.977	0.0011

Supplementary Table S10: Multiple comparisons of condition pairs for model $AE_{T_{dec}}$ in Experiment A. Degrees-of-freedom method: kenward-roger. P-value adjustment: bonferroni method for 6 tests.

	Sum Sq	Mean Sq	NumDF	DenDF	F value	p value
experiment	119186	119186	1	43.073	4.3544	0.04287
entropy	26428	26428	1	44.700	0.9656	0.33108
experiment:entropy	31464	31464	1	44.700	1.1495	0.28940

Supplementary Table S11: Comparing Experiments A and B: X_{dec} – type III ANOVA on LME model estimates with Satterthwaite’s method for degrees of freedom.

	estimate	SE	df	t stat	p value
Experiment A – Experiment B	-88.1	22.2	43	-3.972	0.0003

Supplementary Table S12: Comparing Experiments A and B: pairwise contrast for X_{dec} .

	Sum Sq	Mean Sq	NumDF	DenDF	F value	p value
experiment	0.0002161	0.0002161	1	52.39	0.2046	0.65291
entropy	0.0014956	0.0014956	1	87.80	1.4161	0.23726
experiment:entropy	0.0038077	0.0038077	1	87.80	3.6052	0.06088

Supplementary Table S13: Comparing Experiments A and B: Accuracy – type III ANOVA on LME model estimates with Satterthwaite’s method for degrees of freedom.

	Sum Sq	Mean Sq	NumDF	DenDF	F value	p value
experiment	6.18	6.18	1	43.286	12.6425	0.0009274
entropy	333.48	333.48	1	45.102	681.7821	$< 2.2 \times 10^{-16}$
experiment:entropy	0.17	0.17	1	45.102	0.3544	0.5545809

Supplementary Table S14: Comparing Experiments A and B: T_{dec} – type III ANOVA on LME model estimates with Satterthwaite’s method for degrees of freedom.

	estimate	SE	df	t stat	p value
Experiment A – Experiment B	-0.426	0.0756	43	-5.628	<.0001

Supplementary Table S15: Comparing Experiments A and B: pairwise contrast for T_{dec} (values in log unit).

	Sum Sq	Mean Sq	NumDF	DenDF	F value	p value
v	5.8946	2.94728	2	184	14.7886	1.107×10^{-6}
D	2.0493	1.02465	2	184	5.1414	0.006718
$v:D$	0.6140	0.15349	4	184	0.7702	0.545901

Supplementary Table S16: Experiment B: X_{dec} – type III ANOVA on LME model estimates with Satterthwaite’s method for degrees of freedom.

	estimate	SE	df	t stat	p value
low v - medium v	0.326	0.074	184	4.381	0.0001
low v - high v	0.371	0.074	184	4.981	<.0001
medium v - high v	0.045	0.074	184	0.600	1.0000

Supplementary Table S17: Multiple comparisons of velocity pairs for $\text{Model}_B(X_{\text{dec}})$. Degrees-of-freedom method: kenward-roger. P-value adjustment: bonferroni method for 3 tests.

	estimate	SE	df	t stat	p value
low D - medium D	-0.239	0.074	184	-3.206	0.0048
low D - high D	-0.115	0.074	184	-1.548	0.3698
medium D - high D	0.123	0.074	184	1.658	0.2972

Supplementary Table S18: Multiple comparisons of diffusivity pairs for $\text{Model}_B(X_{\text{dec}})$. Degrees-of-freedom method: kenward-roger. P-value adjustment: bonferroni method for 3 tests.

	Sum Sq	Mean Sq	NumDF	DenDF	F value	p value
v	0.77715	0.38858	2	184	564.85	$< 2.2 \times 10^{-16}$
D	0.97347	0.48673	2	184	707.54	$< 2.2 \times 10^{-16}$
$v:D$	0.11277	0.02819	4	184	40.98	$< 2.2 \times 10^{-16}$

Supplementary Table S19: Experiment B: Accuracy – type III ANOVA on LME model estimates with Satterthwaite’s method for degrees of freedom.

	estimate	SE	df	t stat	p value
<i>v</i> = low					
low <i>D</i> - medium <i>D</i>	0.101	0.008	184	13.391	<.0001
low <i>D</i> - high <i>D</i>	0.096	0.008	184	12.692	<.0001
medium <i>D</i> - high <i>D</i>	-0.005	0.008	184	-0.699	1.0000
<i>v</i> = medium					
low <i>D</i> - medium <i>D</i>	0.109	0.008	184	14.341	<.0001
low <i>D</i> - high <i>D</i>	0.136	0.008	184	17.985	<.0001
medium <i>D</i> - high <i>D</i>	0.028	0.008	184	3.644	.0010
<i>v</i> = high					
low <i>D</i> - medium <i>D</i>	0.174	0.008	184	22.987	<.0001
low <i>D</i> - high <i>D</i>	0.228	0.008	184	30.102	<.0001
medium <i>D</i> - high <i>D</i>	0.054	0.008	184	7.115	<.0001

Supplementary Table S20: Multiple comparisons of diffusivity pairs grouped by velocity for $\text{Model}_B(\alpha)$. Degrees-of-freedom method: kenward-roger. P-value adjustment: bonferroni method for 3 tests.

	estimate	SE	df	t stat	p value
<i>D</i> = low					
low <i>v</i> - medium <i>v</i>	-0.120	0.008	184	-15.881	<.0001
low <i>v</i> - high <i>v</i>	-0.210	0.008	184	-27.717	<.0001
medium <i>v</i> - high <i>v</i>	-0.090	0.008	184	-11.836	<.0001
<i>D</i> = medium					
low <i>v</i> - medium <i>v</i>	-0.113	0.008	184	-14.931	<.0001
low <i>v</i> - high <i>v</i>	-0.137	0.008	184	-18.121	<.0001
medium <i>v</i> - high <i>v</i>	-0.024	0.008	184	-3.190	0.0050
<i>D</i> = high					
low <i>v</i> - medium <i>v</i>	-0.080	0.008	184	-10.588	<.0001
low <i>v</i> - high <i>v</i>	-0.078	0.008	184	-10.307	<.0001
medium <i>v</i> - high <i>v</i>	0.002	0.008	184	0.281	1.0000

Supplementary Table S21: Multiple comparisons of velocity pairs grouped by diffusivity for $\text{Model}_B(\alpha)$. Degrees-of-freedom method: kenward-roger. P-value adjustment: bonferroni method for 3 tests.

	Sum Sq	Mean Sq	NumDF	DenDF	F value	p value
v	0.2060	0.1030	2	184	22.104	2.493×10^{-9}
D	14.4644	7.2322	2	184	1551.716	$< 2.2 \times 10^{-16}$
$v:D$	0.6319	0.1580	4	184	33.894	$< 2.2 \times 10^{-16}$

Supplementary Table S22: Experiment B: T_{dec} – type III ANOVA on LME model estimates with Satterthwaite’s method for degrees of freedom.

	estimate	SE	df	t stat	p value
$v = \text{low}$					
low D - medium D	0.43	0.02	184	21.758	<.0001
low D - high D	0.73	0.02	184	37.258	<.0001
medium D - high D	0.31	0.02	184	15.500	<.0001
$v = \text{medium}$					
low D - medium D	0.25	0.02	184	12.648	<.0001
low D - high D	0.66	0.02	184	33.723	<.0001
medium D - high D	0.42	0.02	184	21.075	<.0001
$v = \text{high}$					
low D - medium D	0.35	0.02	184	17.538	<.0001
low D - high D	0.50	0.02	184	25.412	<.0001
medium D - high D	0.16	0.02	184	7.874	<.0001

Supplementary Table S23: Multiple comparisons of diffusivity pairs grouped by velocity for $\text{Model}_B(T_{\text{dec}})$. Degrees-of-freedom method: kenward-roger. P-value adjustment: bonferroni method for 3 tests.

	estimate	SE	df	t stat	p value
<i>D</i> = low					
low <i>v</i> - medium <i>v</i>	0.11	0.02	184	5.676	<.0001
low <i>v</i> - high <i>v</i>	0.18	0.02	184	9.160	<.0001
medium <i>v</i> - high <i>v</i>	0.07	0.02	184	3.484	0.0019
<i>D</i> = medium					
low <i>v</i> - medium <i>v</i>	-0.07	0.02	184	-3.434	0.0022
low <i>v</i> - high <i>v</i>	0.10	0.02	184	4.940	<.0001
medium <i>v</i> - high <i>v</i>	0.17	0.02	184	8.374	<.0001
<i>D</i> = high					
low <i>v</i> - medium <i>v</i>	0.04	0.02	184	2.141	0.1008
low <i>v</i> - high <i>v</i>	-0.05	0.02	184	-2.686	0.0237
medium <i>v</i> - high <i>v</i>	-0.10	0.02	184	-4.827	<.0001

Supplementary Table S24: Multiple comparisons of velocity pairs grouped by diffusivity for Model_B(T_{dec}). Degrees-of-freedom method: kenward-roger. P-value adjustment: bonferroni method for 3 tests.

	Sum Sq	Mean Sq	NumDF	DenDF	F value	p value
condition	6.8147	2.2716	3	60	12.45	1.967×10^{-6}

Supplementary Table S25: Experiment A: τ_m^* – type III ANOVA on LME model estimates with Satterthwaite’s method for degrees of freedom.

	estimate	SE	df	t stat	p value
I - II	-0.431	0.132	60	-3.270	0.0107
I - III	-0.383	0.132	60	-2.904	0.0309
I - IV	-0.804	0.132	60	-6.100	<.0001
II - III	0.048	0.132	60	0.366	1.0000
II - IV	-0.373	0.132	60	-2.830	0.0379
III - IV	-0.421	0.132	60	-3.196	0.0133

Supplementary Table S26: Experiment A: Multiple comparisons of condition pairs for model τ_m^* . Degrees-of-freedom method: kenward-roger. P-value adjustment: Bonferroni method for 6 tests.

	Sum Sq	Mean Sq	NumDF	DenDF	F value	p value
condition	0.98588	0.49294	2	46	6.7428	0.002704

Supplementary Table S27: Experiment B: τ_m^* – type III ANOVA on LME model estimates with Satterthwaite’s method for degrees of freedom.

	estimate	SE	df	t stat	p value
low - medium	0.240	0.078	46	3.081	0.0104
low - high	-0.015	0.078	46	-0.190	1.0000
medium - high	-0.255	0.078	46	-3.271	0.0061

Supplementary Table S28: Experiment B, comparable conditions: Multiple comparisons of condition pairs for model τ_m^* . Degrees-of-freedom method: kenward-roger. P-value adjustment: bonferroni method for 3 tests.

	Sum Sq	Mean Sq	NumDF	DenDF	F value	p value
Σ	0.16736	0.055788	3	63	20.491	2.21×10^{-9}

Supplementary Table S29: Experiment A: RI_{dec} – type III ANOVA on LME model estimates with Satterthwaite’s method for degrees of freedom, testing differences across entropy production rate conditions.

	estimate	SE	df	t stat	p value
I - II	-0.102	0.016	66.15	-6.188	<.0001
I - III	-0.044	0.016	66.15	-2.653	0.0599
I - IV	-0.109	0.016	66.15	-6.598	<.0001
I - III	0.058	0.016	66.15	3.535	0.0045
II - IV	-0.007	0.016	66.15	-0.411	1.0000
III - IV	-0.065	0.016	66.15	-3.945	0.0012

Supplementary Table S30: Experiment A: Multiple comparisons of condition pairs for RI_{dec} model formula: $RI_{\text{dec}} \sim \Sigma + (1 \mid \text{id})$. Degrees-of-freedom method: kenward-roger. P-value adjustment: bonferroni method for 6 tests.

	Sum Sq	Mean Sq	NumDF	DenDF	F value	p value
v	0.078392	0.039196	2	192	14.721	1.127×10^{-6}
D	0.087628	0.043814	2	192	16.456	2.535×10^{-7}
$v : D$	0.137329	0.034332	4	192	12.894	2.563×10^{-9}

Supplementary Table S31: Experiment B: RI_{dec} – type III ANOVA on LME model estimates with Satterthwaite’s method for degrees of freedom, with drift velocity (v) and diffusivity (D) as fixed effects.

	estimate	SE	df	t stat	p value
low v - medium v	0.0106	0.00879	200	1.206	0.6873
low v - high v	-0.0341	0.00879	200	-3.877	0.0004
medium v - high v	-0.0447	0.00879	200	-5.083	<.0001

Supplementary Table S32: Experiment B: multiple comparisons of velocity groups for RI_{dec} (averaged over diffusivity levels). Degrees-of-freedom method: kenward-roger. P-value adjustment: bonferroni method for 3 tests.

	estimate	SE	df	t stat	p value
low D - medium D	-0.0167	0.00879	200	-1.905	0.1745
low D - high D	-0.0486	0.00879	200	-5.528	<.0001
medium D - high D	-0.0318	0.00879	200	-3.622	0.0011

Supplementary Table S33: Experiment B: multiple comparisons of diffusivity groups for RI_{dec} (averaged over velocity levels). Degrees-of-freedom method: kenward-roger. P-value adjustment: bonferroni method for 3 tests.

	estimate	SE	df	t stat	p value
$D = \text{low}$					
low v - medium v	-0.069	0.015	200.35	-4.562	<.0001
low v - high v	-0.090	0.015	200.35	-5.928	<.0001
medium v - high v	-0.021	0.015	200.35	-1.365	0.5213
$D = \text{medium}$					
low v - medium v	0.043	0.015	200.35	2.827	0.0155
low v - high v	0.011	0.015	200.35	0.753	1.0000
medium v - high v	-0.032	0.015	200.35	-2.074	0.1181
$D = \text{high}$					
low v - medium v	0.058	0.015	200.35	3.825	0.0005
low v - high v	-0.023	0.015	200.35	-1.540	0.3753
medium v - high v	-0.082	0.015	200.35	-5.365	<.0001

Supplementary Table S34: Experiment B: Multiple comparisons of velocity pairs grouped by diffusivity for RI_{dec} model formula: $RI_{\text{dec}} \sim v * D + (1 | \text{id})$. Degrees-of-freedom method: kenward-roger. P-value adjustment: bonferroni method for 3 tests.

	estimate	SE	df	t stat	p value
<i>v</i> = low					
low <i>D</i> - medium <i>D</i>	-0.088	0.015	200.35	-5.790	<.0001
low <i>D</i> - high <i>D</i>	-0.113	0.015	200.35	-7.450	<.0001
medium <i>D</i> - high <i>D</i>	-0.025	0.015	200.35	-1.660	0.2955
<i>v</i> = medium					
low <i>D</i> - medium <i>D</i>	0.024	0.015	200.35	1.599	0.3341
low <i>D</i> - high <i>D</i>	0.014	0.015	200.35	0.938	1.0000
medium <i>D</i> - high <i>D</i>	-0.010	0.015	200.35	-0.661	1.0000
<i>v</i> = high					
low <i>D</i> - medium <i>D</i>	0.014	0.015	200.35	0.890	1.0000
low <i>D</i> - high <i>D</i>	-0.047	0.015	200.35	-3.063	0.0075
medium <i>D</i> - high <i>D</i>	-0.060	0.015	200.35	-3.953	0.0003

Supplementary Table S35: Experiment B: Multiple comparisons of diffusivity pairs grouped by velocity for RI_{dec} model formula: $RI_{\text{dec}} \sim v * D + (1 | \text{id})$. Degrees-of-freedom method: kenward-roger. P-value adjustment: bonferroni method for 3 tests.

	Sum Sq	Mean Sq	NumDF	DenDF	F value	p value
Σ	0.23324	0.11662	2	48	56.045	2.786×10^{-13}

Supplementary Table S36: Experiment B (comparison conditions): RI_{dec} - type III ANOVA on LME model estimates with Satterthwaite's method for degrees of freedom, testing differences across entropy production rate conditions.

	estimate	SE	df	t stat	p value
low Σ - medium Σ	-0.045	0.013	50.09	-3.352	0.0046
low Σ - high Σ	-0.137	0.013	50.09	-10.169	<.0001
medium Σ - high Σ	-0.092	0.013	50.09	-6.817	<.0001

Supplementary Table S37: Experiment B (comparison conditions): Multiple comparisons of velocity pairs for RI_{dec} model formula: $RI_{\text{dec}} \sim \Sigma + (1 | \text{id})$. Degrees-of-freedom method: kenward-roger. P-value adjustment: bonferroni method for 3 tests.

Feature	U_C	max	CL	p value
Jump duration	441	441	1.00 [1.00, 1.00]	<.0001
Jump slope	327	441	0.74 [0.58, 0.88]	0.0077
Jump length	441	441	1.00 [1.00, 1.00]	<.0001
Jump last position	441	441	1.00 [1.00, 1.00]	<.0001

Supplementary Table S38: Experiment A: Wilcoxon rank-sum tests (two-sided; congruent vs. incongruent) across jump features ($n_C = n_I = 21$). U_C is the Mann-Whitney U for the congruent group (maximum $n_C n_I = 441$). CL is the common-language effect size with 95% CI. Holm correction is applied across features.

Feature	Cohen's d
Jump duration	0.163
Jump slope	0.040
Jump length	0.166
Jump last position	0.380

Supplementary Table S39: Experiment A: Standardized effect sizes (Cohen's d) for congruent-incongruent differences across jump features.

Feature	mean diff \bar{d}	HL	CI low	CI high	p value
Jump duration	0.026	0.025	0.024	0.027	<.0001
Jump slope	0.080	0.085	0.058	0.106	<.0001
Jump length	18.55	18.54	17.73	20.10	<.0001
Jump last position	64.75	65.59	60.21	68.94	<.0001

Supplementary Table S40: Experiment A: Paired sign-flip permutation tests on the mean within-subject difference between congruent and incongruent jumps ($C - I$), Holm-corrected across features. \bar{d} denotes the mean difference; HL is the Hodges-Lehmann median shift with nonparametric 95% CI (original units).

	estimate	SE	df	t stat	p value
(Intercept)	-8.366	0.752	785	-11.133	<.0001
x_1	1.686	0.317	785	5.309	<.0001
x_2	1.357	0.407	785	3.333	0.0009
x_3	2.655	0.391	785	6.799	<.0001
x_4	5.158	0.547	785	9.420	<.0001
x_5	5.888	0.543	785	10.850	<.0001

Supplementary Table S41: Experiment A: Binomial logistic GLM coefficients predicting choice from jump-index indicators ($\text{logit}(y) \sim 1 + x_1 + x_2 + x_3 + x_4 + x_5$). Estimates and tests are Wald (log-odds scale).

Source	Sum Sq	df	Mean Sq	F value	p value
Model (Regression)	151.73	5	30.347	518.73	6.3519×10^{-246}
Error (Residual)	45.924	785	0.058503	–	–
Total	197.66	790	–	–	–

Supplementary Table S42: Experiment A: Overall regression ANOVA-style decomposition (omnibus test of the full model (jump integrator) vs. intercept-only).

	estimate	SE	df	t stat	p value
$x_1 - x_2$	0.329	0.410	785	0.803	1.0000
$x_1 - x_3$	-0.970	0.443	785	-2.189	0.2887
$x_1 - x_4$	-3.472	0.597	785	-5.815	<.0001
$x_1 - x_5$	-4.203	0.604	785	-6.960	<.0001
$x_2 - x_3$	-1.299	0.340	785	-3.824	0.0014
$x_2 - x_4$	-3.801	0.630	785	-6.037	<.0001
$x_2 - x_5$	-4.532	0.657	785	-6.896	<.0001
$x_3 - x_4$	-2.502	0.584	785	-4.285	0.0002
$x_3 - x_5$	-3.233	0.599	785	-5.398	<.0001
$x_4 - x_5$	-0.731	0.314	785	-2.326	0.2029

Supplementary Table S43: Experiment A: Multiple comparisons of jump-index coefficients (super-subject binomial GLM). Pairwise t-tests on GLM coefficient contrasts ($\Delta\beta$ on log-odds scale). Degrees-of-freedom from model fit. P-value adjustment: Bonferroni for 10 tests.

Feature	U_C	max	CL	p value
Jump duration	437	576	0.759 [0.610, 0.891]	0.0044
Jump slope	238	576	0.413 [0.255, 0.578]	0.307
Jump length	576	576	1.00 [1.00, 1.00]	<.0001
Jump last position	576	576	1.00 [1.00, 1.00]	<.0001

Supplementary Table S44: Experiment B: Wilcoxon rank-sum tests (two-sided; congruent vs. incongruent) across jump features ($n_C = n_I = 24$). U_C is the Mann–Whitney U for the congruent group (maximum $n_C n_I = 576$). CL is the common-language effect size with 95% CI. Holm correction is applied across features.

Feature	mean diff \bar{d}	HL	CI low	CI high	p value
Jump duration	0.025	0.0230	0.0220	0.0240	<.0001
Jump slope	-0.009	-0.0079	-0.0136	0.0007	0.02
Jump length	13.07	13.35	12.69	13.53	<.0001
Jump last position	54.90	53.20	50.70	59.01	<.0001

Supplementary Table S45: Experiment B: Paired sign-flip permutation tests on the mean within-subject difference between congruent and incongruent jumps ($C - I$), Holm-corrected across features. \bar{d} denotes the mean difference; HL is the Hodges–Lehmann median shift with nonparametric 95% CI (original units).

Feature	Cohen's d
Jump duration	0.1380
Jump slope	-0.0049
Jump length	0.1300
Jump last position	0.2890

Supplementary Table S46: Experiment B: Standardized effect sizes (Cohen's d) for congruent–incongruent differences across jump features.

	estimate	SE	df	t stat	p value
(Intercept)	-9.289	0.420	2677	-22.104	<.0001
x_1	2.646	0.189	2677	14.018	<.0001
x_2	2.498	0.200	2677	12.497	<.0001
x_3	2.669	0.189	2677	14.109	<.0001
x_4	4.964	0.277	2677	17.911	<.0001
x_5	5.674	0.260	2677	21.858	<.0001

Supplementary Table S47: Experiment B: Binomial logistic GLM coefficients for $\text{logit}(y) \sim 1 + x_1 + x_2 + x_3 + x_4 + x_5$. Estimates are on the log-odds scale; tests are Wald t with model df.

Source	Sum Sq	df	Mean Sq	F value	p value
Model (Regression)	516.18	5	103.24	1812.3	$< 5 \times 10^{-324}$
Error (Residual)	152.50	2677	0.056966	–	–
Total	668.68	2682	–	–	–

Supplementary Table S48: Experiment B: Overall regression ANOVA-style decomposition (omnibus test of the full model (jump integrator) vs. intercept-only).

	estimate	SE	df	t stat	p value
$x_1 - x_2$	0.148	0.192	2677	0.769	1.0000
$x_1 - x_3$	-0.023	0.220	2677	-0.106	1.0000
$x_1 - x_4$	-2.318	0.274	2677	-8.454	<.0001
$x_1 - x_5$	-3.029	0.262	2677	-11.557	<.0001
$x_2 - x_3$	-0.171	0.203	2677	-0.843	1.0000
$x_2 - x_4$	-2.466	0.268	2677	-9.194	<.0001
$x_2 - x_5$	-3.176	0.272	2677	-11.661	<.0001
$x_3 - x_4$	-2.295	0.254	2677	-9.044	<.0001
$x_3 - x_5$	-3.005	0.269	2677	-11.183	<.0001
$x_4 - x_5$	-0.711	0.175	2677	-4.052	0.0005

Supplementary Table S49: Experiment B: Multiple comparisons of jump-index coefficients (super-subject binomial GLM). Pairwise t-tests on GLM coefficient contrasts ($\Delta\beta$ on log-odds scale). Degrees-of-freedom from model fit. P-value adjustment: bonferroni for 10 tests.

Condition	v (px/s)	D (px ² /s)	Σ (s ⁻¹)
I) low	30	17700	0.051
II) medium	50	30000	0.083
III) high	72	42000	0.124

Supplementary Table S50: Experiment C, parameter values for the drift velocity v , diffusivity D , and entropy production rate $\Sigma = v^2/D$ used in the experimental conditions (I, II, and III).

	Sum Sq	Mean Sq	NumDF	DenDF	F value	p value
stimulus format	2423707	1211854	2	176	153.039	$< 2.2 \times 10^{-16}$
Σ	2037355	1018677	2	176	128.644	$< 2.2 \times 10^{-16}$
stimulus format: Σ	524732	131183	4	176	16.566	1.537×10^{-11}

Supplementary Table S51: Experiment C: decision thresholds $\langle X_{\text{dec}} \rangle$ – type III ANOVA on LME model estimates with Satterthwaite’s method for degrees of freedom.

	estimate	SE	df	t stat	p value
$\Sigma = \text{low}$					
linear accumulation - radial accumulation	-132.074	27.462	184.38	-4.809	<.0001
linear accumulation - radial reset	-85.615	27.462	184.38	-3.118	0.0063
radial accumulation - radial reset	46.459	27.462	184.38	1.692	0.2771
$\Sigma = \text{medium}$					
linear accumulation - radial accumulation	-236.787	27.462	184.38	-8.622	<.0001
linear accumulation - radial reset	-209.051	27.462	184.38	-7.612	<.0001
radial accumulation - radial reset	27.736	27.462	184.38	1.010	0.9415
$\Sigma = \text{high}$					
linear accumulation - radial accumulation	-363.740	27.462	184.38	-13.245	<.0001
linear - radial reset	-376.970	27.462	184.38	-13.727	<.0001
radial accumulation - radial reset	-13.229	27.462	184.38	-0.482	1.0000

Supplementary Table S52: Experiment C: Multiple comparisons of stimulus format condition pairs grouped by entropy Σ for LME model $\langle X_{\text{dec}} \rangle \sim \text{stimulus_format} * \text{entropy} + (1 | \text{id})$. Degrees-of-freedom method: kenward-roger. P-value adjustment: bonferroni method for 3 tests.

	estimate	SE	df	t stat	p value
stimulus format = linear accumulation					
$\Sigma \text{ low} - \Sigma \text{ medium}$	-49.872	27.462	184.38	-1.816	0.2130
$\Sigma \text{ low} - \Sigma \text{ high}$	-74.123	27.462	184.38	-2.699	0.0228
$\Sigma \text{ medium} - \Sigma \text{ high}$	-24.251	27.462	184.38	-0.883	1.0000
stimulus format = radial accumulation					
$\Sigma \text{ low} - \Sigma \text{ medium}$	-154.585	27.462	184.38	-5.629	<.0001
$\Sigma \text{ low} - \Sigma \text{ high}$	-305.790	27.462	184.38	-11.135	<.0001
$\Sigma \text{ medium} - \Sigma \text{ high}$	-151.205	27.462	184.38	-5.506	<.0001
stimulus format = radial reset					
$\Sigma \text{ low} - \Sigma \text{ medium}$	-173.308	27.462	184.38	-6.311	<.0001
$\Sigma \text{ low} - \Sigma \text{ high}$	-365.478	27.462	184.38	-13.309	<.0001
$\Sigma \text{ medium} - \Sigma \text{ high}$	-192.170	27.462	184.38	-6.998	<.0001

Supplementary Table S53: Experiment C: Multiple comparisons of entropy Σ pairs grouped by stimulus format condition for LME model $\langle X_{\text{dec}} \rangle \sim \text{stimulus_format} * \text{entropy} + (1 | \text{id})$. Degrees-of-freedom method: kenward-roger. P-value adjustment: bonferroni method for 3 tests.

	Sum Sq	Mean Sq	NumDF	DenDF	F value	p value
stimulus format	0.114075	0.057037	2	176	25.7022	1.61×10^{-10}
Σ	0.016286	0.008143	2	176	3.6693	0.02746
stimulus format: Σ	0.028671	0.007168	4	176	3.2299	0.01377

Supplementary Table S54: Experiment C: accuracy α – type III ANOVA on LME model estimates with Satterthwaite’s method for degrees of freedom.

	estimate	SE	df	t stat	p value
$\Sigma = \text{low}$					
linear accumulation - radial accumulation	-0.005	0.015	184.38	-0.360	1.0000
linear accumulation - radial reset	0.052	0.015	184.38	3.576	0.0013
radial accumulation - radial reset	0.057	0.015	184.38	3.936	0.0004
$\Sigma = \text{medium}$					
linear accumulation - radial accumulation	-0.052	0.015	184.38	-3.560	0.0014
linear accumulation - radial reset	-0.000	0.015	184.38	-0.029	1.0000
radial accumulation - radial reset	0.051	0.015	184.38	3.531	0.0016
$\Sigma = \text{high}$					
linear accumulation - radial accumulation	-0.065	0.015	184.38	-4.454	<.0001
linear accumulation - radial reset	-0.002	0.015	184.38	-0.131	1.0000
radial accumulation - radial reset	0.063	0.015	184.38	4.323	0.0001

Supplementary Table S55: Experiment C: Multiple comparisons of stimulus format condition pairs grouped by entropy Σ for LME model $\alpha \sim \text{stimulus.format} * \text{entropy} + (1 | \text{id})$. Degrees-of-freedom method: kenward-roger. P-value adjustment: bonferroni method for 3 tests.

	estimate	SE	df	t stat	p value
stimulus format = linear accumulation					
Σ low - Σ medium	0.045	0.015	184.38	3.123	0.0062
Σ low - Σ high	0.028	0.015	184.38	1.931	0.1650
Σ medium - Σ high	-0.017	0.015	184.38	-1.192	0.7044
stimulus format = radial accumulation					
Σ low - Σ medium	-0.001	0.015	184.38	-0.077	1.0000
Σ low - Σ high	-0.031	0.015	184.38	-2.163	0.0955
Σ medium - Σ high	-0.030	0.015	184.38	-2.087	0.1149
stimulus format = radial reset					
Σ low - Σ medium	-0.007	0.015	184.38	-0.482	1.0000
Σ low - Σ high	-0.026	0.015	184.38	-1.777	0.2319
Σ medium - Σ high	-0.019	0.015	184.38	-1.294	0.5914

Supplementary Table S56: Experiment C: Multiple comparisons of entropy Σ pairs grouped by stimulus format condition for LME model $\alpha \sim \text{stimulus_format} * \text{entropy} + (1 | \text{id})$. Degrees-of-freedom method: kenward-roger. P-value adjustment: bonferroni method for 3 tests.

	Sum Sq	Mean Sq	NumDF	DenDF	F value	p value
stimulus format	19.4414	9.7207	2	175	251.7801	$< 2.2 \times 10^{-16}$
Σ	3.9262	1.9631	2	175	50.8466	$< 2.2 \times 10^{-16}$
stimulus format: Σ	1.3053	0.3263	4	175	8.4521	2.943×10^{-6}

Supplementary Table S57: Experiment C: decision times $\langle T_{\text{dec}} \rangle$ – type III ANOVA on LME model estimates with Satterthwaite’s method for degrees of freedom.

	estimate	SE	df	t stat	p value
$\Sigma = \text{low}$					
linear accumulation - radial accumulation	-0.460	0.061	183.38	-7.587	<.0001
linear accumulation - radial reset	-0.432	0.061	183.38	-7.120	<.0001
radial accumulation - radial reset	0.028	0.061	183.38	0.467	1.0000
$\Sigma = \text{medium}$					
linear accumulation - radial accumulation	-0.702	0.061	183.38	-11.575	<.0001
linear accumulation - radial reset	-0.706	0.061	183.38	-11.635	<.0001
radial accumulation - radial reset	-0.004	0.061	183.38	-0.060	1.0000
$\Sigma = \text{high}$					
linear accumulation - radial accumulation	-0.817	0.061	183.44	-13.293	<.0001
linear accumulation - radial reset	-0.894	0.061	183.44	-14.556	<.0001
radial accumulation - radial reset	-0.078	0.061	183.38	-1.279	0.6073

Supplementary Table S58: Experiment C: Multiple comparisons of stimulus format condition pairs grouped by entropy Σ for LME model $\langle T_{\text{dec}} \rangle \sim \text{stimulus_format} * \text{entropy} + (1 | \text{id})$. Degrees-of-freedom method: kenward-roger. P-value adjustment: bonferroni method for 3 tests.

	estimate	SE	df	t stat	p value
stimulus format = linear accumulation					
$\Sigma \text{ low} - \Sigma \text{ medium}$	0.337	0.061	183.38	5.553	<.0001
$\Sigma \text{ low} - \Sigma \text{ high}$	0.619	0.061	183.44	10.081	<.0001
$\Sigma \text{ medium} - \Sigma \text{ high}$	0.283	0.061	183.44	4.601	<.0001
stimulus format = radial accumulation					
$\Sigma \text{ low} - \Sigma \text{ medium}$	0.095	0.061	183.38	1.566	0.3576
$\Sigma \text{ low} - \Sigma \text{ high}$	0.263	0.061	183.38	4.333	0.0001
$\Sigma \text{ medium} - \Sigma \text{ high}$	0.168	0.061	183.38	2.767	0.0187
stimulus format = radial reset					
$\Sigma \text{ low} - \Sigma \text{ medium}$	0.063	0.061	183.38	1.038	0.9021
$\Sigma \text{ low} - \Sigma \text{ high}$	0.157	0.061	183.38	2.586	0.0314
$\Sigma \text{ medium} - \Sigma \text{ high}$	0.094	0.061	183.38	1.548	0.3697

Supplementary Table S59: Experiment C: Multiple comparisons of entropy Σ pairs grouped by stimulus format condition for LME model $\langle T_{\text{dec}} \rangle \sim \text{stimulus_format} * \text{entropy} + (1 | \text{id})$. Degrees-of-freedom method: kenward-roger. P-value adjustment: bonferroni method for 3 tests.

	Sum Sq	Mean Sq	NumDF	DenDF	F value	p value
stimulus format	0.162457	0.081228	2	176	46.019	$< 2.2 \times 10^{-16}$
Σ	0.061351	0.030676	2	176	17.379	1.295×10^{-7}
stimulus format: Σ	0.024803	0.006201	4	176	3.513	0.008705

Supplementary Table S60: Experiment C: AE_α – type III ANOVA on LME model estimates with Satterthwaite’s method for degrees of freedom.

Contrast	Estimate	SE	df	t stat	p value
Linear accumulation – Radial accumulation	−0.0240	0.00749	184	−3.204	0.0048
Linear accumulation – Radial reset	−0.0691	0.00749	184	−9.230	< 0.0001
Radial accumulation – Radial reset	−0.0451	0.00749	184	−6.026	< 0.0001

Supplementary Table S61: Experiment C: AE_α – post-hoc pairwise contrasts across stimulus format (averaged over entropy), Kenward–Roger degrees of freedom, Bonferroni-adjusted p values (3 tests).

Contrast	Estimate	SE	df	t stat	p value
Σ low – Σ medium	−0.0290	0.00749	184	−3.873	0.0004
Σ low – Σ high	−0.0421	0.00749	184	−5.629	< 0.0001
Σ medium – Σ high	−0.0131	0.00749	184	−1.756	0.2420

Supplementary Table S62: Experiment C: AE_α – post-hoc pairwise contrasts across entropy levels (averaged over stimulus format), Kenward–Roger degrees of freedom, Bonferroni-adjusted p values (3 tests).

	estimate	SE	df	t stat	p value
$\Sigma = \text{low}$					
linear accumulation - radial accumulation	−0.016	0.013	184.38	−1.265	0.6227
linear accumulation - radial reset	−0.040	0.013	184.38	−3.111	0.0065
radial accumulation - radial reset	−0.024	0.013	184.38	−1.847	0.1992
$\Sigma = \text{medium}$					
linear accumulation - radial accumulation	−0.016	0.013	184.38	−1.196	0.6994
linear accumulation - radial reset	−0.062	0.013	184.38	−4.788	$<.0001$
radial accumulation - radial reset	−0.047	0.013	184.38	−3.591	0.0013
$\Sigma = \text{low}$					
linear accumulation - radial accumulation	−0.040	0.013	184.38	−3.089	0.0070
linear accumulation - radial reset	−0.105	0.013	184.38	−8.088	$<.0001$
Radial accumulation - Radial reset	−0.065	0.013	184.38	−5.000	$<.0001$

Supplementary Table S63: Experiment C: Multiple comparisons of stimulus format condition pairs grouped by entropy Σ for LME model $AE_\alpha \sim \text{stimulus_format} * \text{entropy} + (1 | \text{id})$. Degrees-of-freedom method: kenward-roger. P-value adjustment: bonferroni method for 3 tests.

	estimate	SE	df	t stat	p value
stimulus format = linear accumulation					
Σ low - Σ medium	-0.022	0.013	184.38	-1.700	0.2725
Σ low - Σ high	-0.013	0.013	184.38	-0.983	0.9809
Σ medium - Σ high	0.009	0.013	184.38	0.717	1.0000
stimulus format = radial accumulation					
Σ low - Σ medium	-0.021	0.013	184.38	-1.631	0.3135
Σ low - Σ high	-0.036	0.013	184.38	-2.807	0.0166
Σ medium - Σ high	-0.015	0.013	184.38	-1.175	0.7240
stimulus format = radial reset					
Σ low - Σ medium	-0.044	0.013	184.38	-3.376	0.0027
Σ low - Σ high	-0.077	0.013	184.38	-5.960	<.0001
Σ medium - Σ high	-0.034	0.013	184.38	-2.584	0.0316

Supplementary Table S64: Experiment C: Multiple comparisons of entropy Σ pairs grouped by stimulus format condition for LME model $AE_{\alpha} \sim \text{stimulus_format} * \text{entropy} + (1 | \text{id})$. Degrees-of-freedom method: kenward-roger. P-value adjustment: bonferroni method for 3 tests.

	Sum Sq	Mean Sq	NumDF	DenDF	F value	p value
stimulus format	6.3349	3.1675	2	176	33.3364	5.288×10^{-13}
Σ	13.2267	6.6133	2	176	69.6031	$< 2.2 \times 10^{-16}$
stimulus format: Σ	0.6102	0.1525	4	176	1.6055	0.1749

Supplementary Table S65: Experiment C: $AE_{\langle T_{\text{dec}} \rangle}$ - type III ANOVA on LME model estimates with Satterthwaite's method for degrees of freedom.

	estimate	SE	df	t stat	p value
linear accumulation - radial accumulation	-0.218	0.055	184.38	-3.977	0.0003
linear accumulation - radial reset	-0.438	0.055	184.38	-7.978	<.0001
radial accumulation - radial reset	-0.220	0.055	184.38	-4.001	0.0003

Supplementary Table S66: Experiment C: Multiple comparisons of stimulus format condition pairs averaged over entropy levels for LME model $AE_{\langle T_{\text{dec}} \rangle} \sim \text{stimulus_format} * \text{entropy} + (1 | \text{id})$. Degrees-of-freedom method: kenward-roger. P-value adjustment: bonferroni method for 3 tests. Contrasts are squared-root scale.

	estimate	SE	df	t stat	p value
Σ low - Σ medium	0.126	0.055	184.38	2.292	0.0691
Σ low - Σ high	0.600	0.055	184.38	10.930	<0.0001
Σ medium - Σ high	0.474	0.055	184.38	8.638	<0.0001

Supplementary Table S67: Experiment C: Multiple comparisons of entropy levels pairs averaged over stimulus format conditions for LME model $AE_{(T_{dec})} \sim \text{stimulus_format} * \text{entropy} + (1 | \text{id})$. Degrees-of-freedom method: kenward-roger. P-value adjustment: bonferroni method for 3 tests. Contrasts are squared-root scale.

	Sum Sq	Mean Sq	NumDF	DenDF	F value	p value
Σ	8.4000×10^{-8}	4.2000×10^{-8}	2	176	0.0503	0.9510
stimulus format	7.2629×10^{-5}	3.6314×10^{-5}	2	176	43.4631	4.568×10^{-16}
Σ :stimulus format	5.8880×10^{-6}	1.4720×10^{-6}	4	176	1.7617	0.1386

Supplementary Table S68: Experiment C: RI_{dec} – type III ANOVA on LME model estimates with Satterthwaite’s method for degrees of freedom.

	estimate	SE	df	t stat	p value
linear accumulation - radial accumulation	0.001	0.000	184.38	7.184	<0.0001
linear accumulation - radial reset	0.001	0.000	184.38	8.442	<0.0001
radial accumulation - radial reset	0.000	0.000	184.38	1.258	0.6301

Supplementary Table S69: Experiment C: Multiple comparisons of stimulus format condition pairs averaged over entropy levels for LME model $RI_{dec} \sim \text{stimulus_format} * \text{entropy} + (1 | \text{id})$. Degrees-of-freedom method: kenward-roger. P-value adjustment: bonferroni method for 3 tests.

Feature	U_C	max	CL	p value
Jump duration	378	484	0.78 [0.63, 0.92]	0.003
Jump length	468	484	0.97 [0.9, 1]	<.0001
Jump slope	363	484	0.75 [0.60, 0.88]	0.005
Jump last position	483	484	0.99 [0.98, 1]	<.0001

Supplementary Table S70: Experiment C (linear accumulation): Wilcoxon rank–sum tests (two-sided; congruent vs. incongruent) across jump features ($n_C = n_I = 22$). U_C is the Mann–Whitney U for the congruent group (maximum $n_C n_I = 484$). CL is the common-language effect size with 95% CI. Holm correction is applied across features.

Feature	mean diff \bar{d}	HL	CI low	CI high	p value
Jump duration	0.015	0.015	0.013	0.017	<.0001
Jump slope	0.010	0.010	0.070	0.120	<.0001
Jump length	9.5	10.3	7.26	12.54	<.0001
Jump last position	57.07	59.37	55.01	62.77	<.0001

Supplementary Table S71: Experiment C (linear accumulation): Paired sign-flip permutation tests on the mean within-subject difference between congruent and incongruent jumps ($C - I$), Holm-corrected across features. \bar{d} denotes the mean difference; HL is the Hodges–Lehmann median shift with nonparametric 95% CI (original units).

	estimate	SE	df	t stat	p value
(Intercept)	-6.465	0.675	401	-9.582	<.0001
x1	1.781	0.408	401	4.359	<.0001
x2	2.231	0.449	401	4.965	<.0001
x3	2.701	0.415	401	6.510	<.0001
x4	3.047	0.372	401	8.200	<.0001
x5	3.451	0.378	401	9.139	<.0001

Supplementary Table S72: Experiment C (Linear Accumulation): Binomial logistic GLM coefficients predicting choice from jump-index indicators ($\text{logit}(y) \sim 1 + x_1 + x_2 + x_3 + x_4 + x_5$). Estimates and tests are Wald (log-odds scale).

Source	Sum Sq	df	Mean Sq	F value	p value
Model (Regression)	66.967	5	13.393	155.87	1.1444×10^{-91}
Error (Residual)	34.458	401	0.085929	–	–
Total	101.43	406	–	–	–

Supplementary Table S73: Experiment C (linear accumulation): Overall regression ANOVA-style decomposition (omnibus test of the full model (jump integrator) vs. intercept-only).

	estimate	SE	df	t stat	p value
$x_1 - x_2$	-0.451	0.401	401	-1.125	0.2614
$x_1 - x_3$	-0.920	0.489	401	-1.883	0.6047
$x_1 - x_4$	-1.267	0.475	401	-2.665	0.0801
$x_1 - x_5$	-1.670	0.513	401	-3.253	0.0124
$x_2 - x_3$	-0.469	0.436	401	-1.076	0.2828
$x_2 - x_4$	-0.816	0.519	401	-1.572	0.1168
$x_2 - x_5$	-1.219	0.569	401	-2.143	0.3268
$x_3 - x_4$	-0.346	0.458	401	-0.756	0.4501
$x_3 - x_5$	-0.750	0.452	401	-1.659	0.9782
$x_4 - x_5$	-0.404	0.431	401	-0.935	0.3501

Supplementary Table S74: Experiment C (Linear Accumulation): Multiple comparisons of jump-index coefficients (super-subject binomial GLM). Pairwise t-tests on GLM coefficient contrasts ($\Delta\beta$ on log-odds scale). Degrees-of-freedom from model fit. P-value adjustment: Bonferroni for 10 tests.

	n_C	n_I	U_c	p value	CL [95% CI]
jump durations	22	22	259	1.00000	0.53512 [0.36570, 0.70868]
jump slopes	22	22	168	0.33794	0.34711 [0.19421, 0.52066]
jump lengths	22	22	196	0.85656	0.40496 [0.24174, 0.57851]
jump last position	22	22	241	1.00000	0.49793 [0.32231, 0.66736]

Supplementary Table S75: Experiment C — Radial Accumulation: Wilcoxon rank-sum tests (two-sided; congruent vs. incongruent; $n_C = n_I = 22$). Reported p-values are Holm-corrected for multiple comparisons.

	n	\bar{d} (C-I)	p value	HL [95% CI]
jump durations	22	0.00049	1.00000	-0.00099 [-0.00208, 0.00395]
jump slopes	22	-0.02632	0.12800	-0.02979 [-0.04652, 0.00205]
jump lengths	22	-0.79377	0.99932	-1.15674 [-2.59214, 0.58663]
jump last position	22	-2.15564	1.00000	-4.23165 [-11.23713, 5.01899]

Supplementary Table S76: Experiment C — Radial Accumulation: Paired sign-flip permutation tests on mean differences ($C-I$), Holm-corrected across features. HL = Hodges-Lehmann estimator with 95% confidence interval.

	n_C	n_I	U_c	p value	CL [95% CI]
jump durations	22	22	259	1.00000	0.53512 [0.35537, 0.70868]
jump slopes	22	22	216	1.00000	0.44628 [0.27479, 0.62397]
jump lengths	22	22	247	1.00000	0.51033 [0.33471, 0.67769]
jump last position	22	22	240	1.00000	0.49587 [0.32231, 0.66736]

Supplementary Table S77: Experiment C — Radial Reset: Wilcoxon rank-sum tests (two-sided; congruent vs. incongruent; $n_C = n_I = 22$). Reported p-values are Holm-corrected for multiple comparisons.

	n	\bar{d} (C-I)	p value	HL [95% CI]
jump durations	22	0.00023	1.00000	0.00079 [-0.00027, 0.00191]
jump slopes	22	-0.01464	1.00000	-0.02553 [-0.07325, -0.00277]
jump lengths	22	0.19217	1.00000	0.18068 [-1.63088, 1.36869]
jump last position	22	0.02807	1.00000	-0.07137 [-1.79553, 1.33771]

Supplementary Table S78: Experiment C — Radial Reset: Paired sign-flip permutation tests on mean differences ($C - I$), Holm-corrected across features. HL = Hodges-Lehmann estimator with 95% confidence interval.

	Linear Accumulation	Radial Accumulation	Radial Reset
jump durations	0.0930	-0.00003	-0.00181
jump lengths	0.1055	-0.01083	-0.00077
jump slopes	0.0579	-0.01147	-0.01034
jump last position	0.3039	-0.00632	-0.00059

Supplementary Table S79: Experiment C: Cohen's d effect sizes for each jump feature across the three accumulation conditions (Linear accumulation, Radial accumulation, and Radial reset).