Deliberate Stochastic Choice: A Mouse-tracking Experiment

Mattia Adamo*, Giorgio Coricelli†

Abstract

When individuals repeatedly choose from the same set of risky alternatives, they often select different options, a phenomenon known as stochastic choice (SC). We combine the experimental design of Agranov and Ortoleva (2017) with mouse-tracking to examine the cognitive mechanisms underlying this variability. Identical choice sets are presented under two repetition structures: distant repeats (spaced trials) and sequential repeats (immediate succession). We replicate core findings on SC across both structures and use process measures to directly link observed behavior to deliberate stochastic-choice mechanisms. Mouse trajectories reveal distinct dynamics: the Area Under the Curve (AUC), the area between the actual trajectory and an ideal linear path to the selected option, remains stable across repetitions in distant repeats but declines sharply after the first decision in sequential repeats. These patterns suggest across-trial stochastic balancing in distant repeats and planning-based deliberate randomization in sequential repeats. Moreover, AUC consistently outperforms response time as a predictor of SC, offering a more sensitive process-level indicator of stochastic behavior.

Keywords: Stochastic Choice, Randomization, Choice Under Risk, Mouse Tracking

*IMT School for Advanced Studies, Lucca. Email: `mattia.adamo@imtlucca.it`, `mattadam92@icloud.com`.

†University of Southern California, Los Angeles. Email: `coricell@usc.edu`

1. Introduction

When decision-makers are repeatedly asked to choose from the same set of options, they often select different alternatives across trials, a phenomenon known as stochastic choice (SC). This tendency towards randomization has received significant attention across economics, psychology, and decision science (Tversky, 1969; Luce, 1959; Camerer, 1989). Often observed in lottery choices (Mosteller and Nogee, 1951; Luce, 1959; Hey and Orme, 1994; Regenwetter et al., 2011), SC poses further challenges to traditional models of rationality (Samuelson, 1938; Von Neumann and Morgenstern, 1944). SC theory has hence developed into a robust interdisciplinary framework for investigating the cognitive and behavioral mechanisms underlying variability in decision-making (Cerreia-Vioglio et al., 2015; Fudenberg and Strzalecki, 2015; Caplin et al., 2019; Woodford, 2020; Allen and Rehbeck, 2023; Agranov et al., 2023).

Several theories have emerged to explain the reasons for such variability. First, individuals might randomize choice assuming multiple utility functions, subject to random shocks due to unobservable factors (Becker et al., 1964; Gul and Pesendorfer, 2006). Second, SC may be due to internal cognitive constraints in which randomness emerges from noisy and stochastic evidence accumulation (Ratcliff, 1978; Camerer, 1998; Krajbich et al., 2010). Third, individuals might deliberately decide to randomize choice, to hedge against risk or avoid regretful consequences, over imperfect preferences (Machina, 1985; Cerreia-Vioglio et al., 2015; Cubitt et al., 2015).

Empirical investigations of inconsistent behavior in economic contexts often focus on settings where decisions are repeated at distant intervals, with participants unaware of the repetitions (Tversky, 1969; Camerer, 1989; Hey and Orme, 1994; Ballinger and Wilcox, 1997). More recent research has analyzed SC in economic settings involving sequential repeats (Cubitt et al., 2015; Agranov and Ortoleva, 2017; Cerreia-Vioglio et al., 2019). This condition is expected to minimize sources of noise due to time distance and memory effects. However, persistent SC is observed in both scenarios and at different levels of difficulty.

Agranov and Ortoleva (2017) (A&O) conducted the first study that compares distant repeats and sequential repeats (called repetitions in a row in the paper) in choice under risk[1]. In the condition with distant repeats, participants are unaware of the repetition

[1] Sequential evaluations of multiple alternatives within a single decision process have been studied in economic experiments under time pressure, but these paradigms differ from true sequential repetitions, where the same 2-options decision problem is presented multiple times consecutively. The Risk-Ambi experiment conducted by Cettolin and Riedl (2019) examines decision-making under uncertainty by presenting repeated choices between a risky and an ambiguous lottery, allowing participants to opt for a randomized mix of the two. While this structure involves repeated exposures to similar decision problems, its primary focus is on the trade-off between risk and ambiguity.

structure, consistent with prior literature. In contrast, in the condition with sequential repeats, participants are explicitly informed that they will be asked to choose from the same choice set three times in a row. The authors interpreted the correlated pattern of SC in both conditions as evidence in favor of the deliberative stochastic choice model (Cerreia-Vioglio et al., 2015). Thus, SC may reflect a deliberative strategy for managing trade-offs between risk, regret, or complexity within the decision environment.

However, several unanswered questions about the nature of SC in these settings remain. Even if SC is observed under both repetition structures, can we conclude that its underlying nature is the same? Are there observable differences in the decision processes under the two conditions? And if so, how can SC under distant repeats, potentially arising from noise or error, be meaningfully associated with SC under sequential repeats? Our paper aims to answer these questions by identifying and characterizing the mechanisms underlying stochastic behavior across different repetition conditions and to test the extent to which the decision processes diverge.

Building on the experimental design of A&O, we introduce two key innovations: (1) the use of mouse-tracking as a dynamic measure of the complexity of the decision process, and (2) the explicit comparison of response time (RT) and the AUC, defined as the area between the observed mouse trajectory and an ideal path to the selected option, as complementary metrics for assessing the decision process. The analysis of our process data allows us to evaluate how different repetition structures shape the extent of SC, thereby providing insight into the underlying cognitive mechanisms. We expect more difficult choices to produce higher AUC and longer RT, and that both measures will capture action planning on the first trial of the sequential-repeats condition, consistent with A&O's findings. In their data, response times were substantially longer for the initial choice but sharply reduced for subsequent repetitions, indicating that subjects primarily executed a pre-formed plan after the first presentation. By contrast, RT patterns in the distant-repeats condition were more heterogeneous and therefore less diagnostic of the underlying source of SC.

Moreover, we expect mouse tracking to be more informative than RT for identifying the nature of SC under the two repetition structures and for discriminating between competing theoretical models. Stochasticity in repeated choices can arise from distinct computational sources. In noise-based models, such as drift–diffusion models (DDM), variability across identical trials is driven by within-trial fluctuations, either in the accumulated evidence or in latent parameters such as the drift rate or starting point. Each choice thus reflects an independent draw of diffusion noise, with no systematic dependence on prior selections. By contrast, deliberative stochastic-choice models, such as Cautious Stochastic Choice (CSC), generate variability through a history-dependent randomization process: choice probabilities are adjusted over time so that deviations in one direction are gradually offset in later trials, implementing an approximate quota-balancing mechanism.

We thus hypothesize that dynamic patterns in AUC, specifically its stability across repetitions and its invariance across stochastic and non-stochastic choices, provide a process-level diagnostic of the mechanisms generating stochastic choice. AUC is widely interpreted as a measure of within-trial response competition (see Section 2.2), and models in which stochasticity arises from independent diffusion noise (as in DDM-type accounts) naturally predict that such competition should fluctuate across trials and differ between SC and non-SC realizations. By contrast, models that generate variability through an across-trial quota-maintaining process (as in the CSC model) leave within-trial dynamics largely unchanged and therefore predict stable, choice-invariant AUC patterns. Accordingly, the observed temporal profile of AUC offers a theoretically grounded means of distinguishing between within-trial noise and across-trial deliberative control as competing explanations for stochastic choice.

Our findings are as follows. First, SC is substantial in both experimental structures, displayed by 84% of participants in distant repeats and by 65% in sequential repeats, and particularly pronounced in questions involving difficult trade-offs. Second, 24% of participants exhibit SC in first-order stochastic dominance questions during distant repeats. Third, the AUC and RT exhibit distinct patterns: AUC remains stable across repetitions in the distant condition but decreases in the sequential condition, mirroring the declining pattern observed in RT across both conditions. Finally, AUC better explains SC than RT.

The remainder of this paper is organized as follows. Section 2 reviews theoretical models of SC and their predictions, and surveys the empirical literature using process data in economics. Section 3 describes the experimental design. Section 4 presents our results. In Section 5, we conclude by discussing the implications of our findings for SC theory.

2. Related Literature

2.1. Stochastic Choice Models

Numerous models have been proposed to account for SC, reflecting a fundamental divide over whether randomness represents a by-product of cognitive or environmental noise or a deliberate component of the decision process.

Noise-based approaches include Random Utility models and Drift Diffusion Models (DDMs). Random utility models (Thurstone, 1927; Block and Marschak, 1959; McFadden and Richter, 1990; Gul and Pesendorfer, 2006; Gul et al., 2014) interpret SC as arising from unobservable heterogeneity in preferences or from utility shocks. In the Random Expected Utility (REU) model, each realization of preferences satisfies the axioms of expected utility theory but is drawn from a probability distribution over such utility functions.

DDMs (Ratcliff, 1978; Ratcliff and McKoon, 2008; Krajbich and Rangel, 2011) offer a second major noise-based account, in which choices arise from a stochastic process of evidence accumulation until a decision threshold is reached (Busemeyer and Townsend, 1993; Johnson and Busemeyer, 2005; Gold and Shadlen, 2007). In such models, variability across identical trials is generated by within-trial fluctuations, either in the accumulated evidence or in latent parameters such as the drift rate or starting point, so that each choice reflects an independent realization of diffusion noise, with no systematic dependence on prior selections. Value-based extensions further incorporate attentional and contextual influences, showing that fixation patterns and attribute salience can modulate the accumulation process (Tsetsos et al., 2010; Krajbich et al., 2010; Krajbich and Rangel, 2011; Krajbich et al., 2012; Dai and Busemeyer, 2014).

In contrast, deliberate randomization models treat stochasticity as the outcome of purposeful planning rather than cognitive imprecision. Early formulations (Machina, 1985) interpret randomization as arising from deterministic preferences over mixtures. More recent frameworks, including CSC, formalize stochastic choice functions as the product of a structured deliberation stage preceding implementation (Cerreia-Vioglio et al., 2015; Brady and Rehbeck, 2016; Cerreia-Vioglio et al., 2019). Empirical and theoretical contributions (Agranov and Ortoleva, 2017; Dwenger et al., 2012; Agranov et al., 2023) emphasize that agents may randomize to avoid regret, diversify risk, or accommodate imprecise preferences, and further show that randomization plans exhibit temporal stability.

Within this deliberate perspective, three main motives for randomization have been identified: (I) Hedging under convex or cautious preferences, when mixtures are strictly preferred to pure options (Machina, 1985; Cerreia-Vioglio et al., 2015, 2019); (II) Regret or responsibility avoidance, whereby randomization reduces anticipated regret or diffuses responsibility for adverse outcomes (Dwenger et al., 2018; Heydari, 2024);); (III) Nonlinear probability weighting, in which rank-dependent distortions increase the attractiveness of probabilistic mixtures relative to deterministic alternatives (Cerreia-Vioglio et al., 2019). Recent experimental evidence using convex budgets further documents systematic preferences for mixtures and their relation to repeated-choice stochasticity. Agranov et al. (2023) show that mixing behavior is correlated across strategic and non-strategic environments, while Feldman and Rehbeck (2022) demonstrate, using linear convex budgets and repeated discrete choices, that a preference for mixtures predicts the distribution of stochastic choices across repetitions.

2.2. Process-tracing literature in Economics

While stochastic choice models provide theoretical accounts of why decisions may appear random, process-tracing methods complement them by directly capturing the real-time temporal and dynamic features of decision-making through measures such as response times, gaze patterns, and mouse trajectories (see also Coricelli et al., 2020). When combined with choice data, these methods provide empirical leverage to disentangle whether stochasticity arises from genuine deliberation or from errors and bounded rationality. The following literature reviews how differences in attention, conflict (both exogenous and endogenous), and motor execution shed light on the mechanisms underlying economic decisions.

Response Times. Response times (RTs) are among the earliest process measures systematically employed in risky choice tasks. Typically, there is a direct relationship between choice complexity and RTs: longer RTs tend to arise when options are closer in subjective value, indicating that part of the within-subject randomness in risky choice is systematically related to time-on-task (Alós-Ferrer et al., 2016; Konovalov and Krajbich, 2016). Conversely, very short RTs are more common when one option is clearly dominated or trivially preferable (Agranov and Ortoleva, 2017). Importantly, however, short RTs may reflect either fluent recognition due to expertise and optimal conditions, implying fast and precise deliberation, or impulsive responding, suggesting random/boundedly rational choice (Börger, 2016). Similarly, long RTs may indicate careful integration of information due to task complexity, but can also reflect hesitation due to inexperience, or internal cognitive limits.

Methodological literature emphasizes the importance of adequate experimental designs and control aimed to interpret RTs correctly. Longer RTs often emerge when informational demands increase (e.g., multiple attributes, rare probabilities), consistent with higher conflict (Arieli et al., 2011; Rubinstein, 2013). Under time pressure or cognitive load, participants often display shorter RTs, and shifts in risk-taking can occur in either direction depending on framing and task structure (Kocher et al., 2013).

The evidence suggests that inference linking fast responses to intuition and slow responses to deliberation may be misleading, as non-linearities add an additional layer of complexity to the speed–accuracy trade-off (Krajbich et al., 2015), rendering RTs alone insufficient to disentangle the dynamics of choice in economic settings. Combining RTs with other process-tracing data is therefore essential to uncover how temporal dynamics interact with different sources of choice uncertainty.

Eye Movement Data While eye-tracking provides direct measures of attention allocation, its relevance within economics has been emphasized also in theoretical discussions of search and information acquisition. For instance, Caplin and Dean (2011) highlight that process data—such as those obtained from eye-tracking or MouseLab—offer crucial empirical leverage for identifying underlying decision processes, complementing traditional revealed-preference approaches. In risky choice experiments where outcomes and probabilities are presented explicitly, eye-tracking delivers granular measures of attention—such as dwell times on payoffs vs. probabilities, transition paths across attributes, and late-stage attention asymmetries as choices evolve (Arieli et al., 2011). A robust finding is that attention is selective and dynamically structured: participants tend to devote more gaze time to higher payoffs and higher probabilities, and shifts of gaze toward an option often precede and predict the eventual choice (Glöckner and Herbold, 2011; Fiedler and Glöckner, 2012; Stewart et al., 2016). Eye-tracking further refines the noisy vs. deliberate distinction by showing that similar choices can be produced through different attentional paths. Incorporating fixations into econometric models improves fit and parameter stability because attention covaries with the subjective weights applied to attributes during integration (Harrison and Swarthout, 2019). Moreover, eye-tracking has been used to disentangle how counterfactual emotions such as regret and disappointment shape visual inspection of obtained and forgone outcomes in lottery tasks, revealing distinct attentional patterns for these emotions (Bault et al., 2016).

Mouse Data The seminal study by Camerer et al. (1993) is considered one of the earliest process-tracing contributions in economics. The authors used MouseLab to infer participants' underlying cognitive strategies in a sequential bargaining game. Payoff information was hidden behind boxes, which participants could reveal by clicking with the mouse on each box. They analyzed the sequence of information acquisition, the time spent examining each payoff, and the transitions between informational elements in a sequential bargaining game involving shrinking gains or expanding losses. The results revealed a pronounced reliance on present information, with participants mainly focusing on the immediate round rather than employing backward induction across all rounds. Furthermore, the results corroborated the asymmetry between gains and losses, with greater cognitive effort but more dispersed behavior observed in the loss domain.

MouseLab captured the sequential exploration of discrete information elements, whereas later continuous mouse-tracking methods (e.g., Freeman and Ambady (2010)) shifted focus to the dynamic unfolding of decision conflict in real time, by continuously recording motor trajectories without requiring discrete information acquisition steps.

Continuous mouse-tracking has been increasingly employed to investigate real-time cognitive dynamics in the social sciences, all of which highlight a more pronounced trajectory as an indication of higher decision conflict. Spivey et al. (2005) first applied continuous mouse-tracking to language processing, showing that trajectories were dynamically attracted toward phonological competitors, revealing ongoing cognitive conflict during word recognition. Freeman et al. (2008) extended mouse-tracking to social categorization, demonstrating that mouse trajectories veered toward competing gender categories before final resolution, indicating dynamic competition between social representations. Scherbaum et al. (2010) introduced mouse-tracking into economic decision-making contexts, using a continuous Simon task to show that conflict between spatial stimulus and response mappings dynamically shaped the unfolding of motor trajectories. Dshemuchadse et al. (2013) applied mouse-tracking to intertemporal choice, observing greater curvature in trajectories when participants chose delayed larger rewards over immediate smaller ones, reflecting heightened cognitive conflict. Sullivan et al. (2015) analyzed dietary choices, revealing that taste attributes influenced mouse trajectories earlier than health attributes, supporting a dynamic model of self-control conflict resolution. Konovalov and Krajbich (2020) employed mouse-tracking in learning tasks, finding that participants' mouse movements revealed implicit structure knowledge even when overt choices appeared model-free, indicating dissociation between cognitive representations and final behavior. Stillman et al. (2018) showed that the AUC in mouse-tracking trajectories serves as a discrete proxy for continuous decision conflict, successfully predicting individual risk preferences—such as loss aversion and diminishing marginal utility—even at the single-trial level, outperforming traditional reaction time measures.

3. Methodology

3.1. Experiment Design

We replicate the experimental framework of A&O, introducing mouse tracking to record decision trajectories on a trial-by-trial basis and adding a set of additional tests to observe different behavioral and cognitive patterns of SC.

The experiment consists of seven distinct phases, four of which make the original design (phase I to phase IV). Phases I and III serve as the main parts[2]. These were designed to explore participants' SC through repeated lottery decisions with different repetitions structure: in phase I, the same choice sets are distant between each other, and occur in a pseudo-random order. Here participants are not explicitly told of the repetitions. In phase III each question is presented three times in sequence, and participants are

[2]Refer to figure S1 in Appendix A1 for a visualization of the experimental timeline. See Appendix A3 for the complete instructions.

explicitly informed. No feedback is provided in either phase.

Between phases I and III, phase II acts as a transitional stage. Its inclusion serves to reduce cognitive fatigue from the repeated decision tasks and make a clearer separation between the two main parts. This phase measures participants attitudes toward risk (Gneezy and Potters, 1997). Phase IV evaluates deviations from expected utility theory by presenting decision problems based on the variations of Allais paradox, focusing on the common consequences and common ratio effects (Allais, 1953).

Following the main and transitional phases, we included three post-experiment tasks. Phase V assesses participants preferences for convex combinations of existing lotteries. Phase VI asks participants to answer the questions from Phase I once again, and indicate their confidence as a mean to assess preference precision (Cubitt et al., 2015). Phase VII is a memory task asking to recall previous lotteries. This phase concludes with a question asking to estimate the accuracy of memory recall to each subject.

In the main phases of the experiment (Phases I and III), participants were asked to select among two-lottery choice sets displayed on the extreme left and right of the screen. Each lottery provided a payoff in tokens; the outcome was based on the roll of a simulated four-sided die with faces labeled A, B, C, and D. The outcomes of the lotteries were associated with probabilities of 0.25, 0.5, 0.75, or 1. The choice sets presented varied in difficulty, with a total of 10 unique questions clustered into three distinct groups: three characterized by First-Order Stochastic Dominance (FOSD)[3], three EASY[4] and four HARD questions[5]. The 10 question types are displayed in Table 1 below.

In Phase I, participants completed 40 decision rounds, where each of the 10 choice sets was repeated four times in a pseudo-random order. The pseudo-randomization ensured that repetitions were distributed across four blocks of 10 trials each, with each choice set presented once per block and placed to guarantee sufficient distance. For three of the four trials, the questions appeared without modifications. In one trial, pseudo-randomly ordered, participants could choose to let a simulated coin flip replace the burden of choice. The coin, if selected, costs participants 1 token from their final payoff. Table D1 in Appendix A1 presents the order of screening for Phase I.

Phase II provides a straightforward and incentive-compatible measure of individual risk aversion, à la Gneezy and Potters (1997). Participants were given 100 tokens and asked to decide how many tokens to invest. In the risky investment task, the investment had a 50% chance of success, yielding x2.5 times the investment, and a 50% chance

[3]Strictly speaking, the replicated FOSD questions correspond to 0th-order stochastic dominance: a pointwise comparison shows that one lottery never yields less than the other and yields more in at least one outcome.

[4]EASY questions are characterized by second order stochastic dominance.

[5]HARD questions featured pairs of lotteries with no obvious choice.

TABLE 1: List of Questions Asked

	Lottery 1				Lottery 2				EV	SD
Question	(.25)	(.25)	(.25)	(.25)	(.25)	(.25)	(.25)	(.25)	Difference	Difference
FOSD1	98	98	98	98	103	103	103	103	**5**	**0**
FOSD2	17	17	18	18	17	17	17	17	**0.5**	**0.5**
FOSD3	10	20	30	30	70	100	120	190	**97.5**	**35.87**
EASY1	23	23	30	30	5	5	5	31	**31**	**7.76**
EASY2	12	14	16	96	85	85	85	85	**50.5**	**35.54**
EASY3	100	100	100	100	20	20	20	101	**59.75**	**35.08**
HARD1	38	38	38	77	16	16	94	94	**7.25**	**22.12**
HARD2	10	10	90	90	32	45	45	56	**5.5**	**31.5**
HARD3	6	84	105	200	54	60	117	135	**7.25**	**33.97**
HARD4	13	30	51	81	19	32	38	86	**0**	**0.03**

Note: Each question is one row; subjects were asked to choose between lotteries 1 and 2. Each lottery is described by the four amounts of tokens that it can pay out with equal probability (i.e. $p = 0.25$). The tokens are converted into US dollars at the rate 20 tokens = \$1. The last column contains the absolute value of the difference in expected values. A screenshot is presented in Appendix A1. All 10 questions were presented in Phase I, while only the 7 questions indicated in bold were used in Phase III.

of failure, resulting in a loss of the invested tokens. In the compound lottery task the structure was similar to the previous one, but the success of the investment depended on a compound lottery that reduced to a 50% chance of success.

In Phase III, participants played a subset of seven questions from Phase I (FOSD1, EASY1, EASY2, HARD1, HARD2, HARD3, HARD4). Each question was repeated three times consecutively, summing to 21 decision rounds in total. Participants were aware of this order. Table D2 in Appendix A1 shows the order of screening.

Phase IV addressed the common consequence and common ratio effects, two classic violations of the independence axiom in expected utility theory where choices are influenced by irrelevant alternatives or proportional changes in probabilities.

In phase V, participants were asked to complete a subset of six questions from the main experimental phases (FOSD1, EASY1, HARD1, HARD2, HARD3, HARD4). Each question was repeated four times, but with an altered structure: the two original lotteries were mixed to create a convex combination at a 50/50 ratio, resulting in a third convex lottery. In three out of four trials, participants chose between two of these lotteries or opted to flip a coin to select the convex combination. In the remaining trial subjects could choose among the two original options or the convex combination.

The primary objective of this test was to assess participants' propensity to select convex mixtures as decision difficulty increased. Based on their responses, participants were categorized into behavioral profiles, reflecting their preferences and sensitivity to convexity [6].

In phase VI, participants were asked to complete an additional set of 10 decisions aimed to test for preference imprecision. These decisions mirrored the original questions, and participants were prompted to indicate their level of confidence on a scale from 1 to 10 after choosing.

In the final phase of the experiment, we asked participants to complete a memory test to assess their recall of the lotteries presented during the tasks. Participants were shown 25 lotteries, 12 of which had been used in the main tasks and 13 of which were novel but designed to resemble the originals. The novel lotteries included those with identical expected values (EV), identical variances, both attributes identical, or both differing. Subjects were asked to identify whether each lottery had been part of the main tasks with a YES/NO question. At the end of the memory test, a final question asked the percentage of lotteries subjects believed to have correctly recalled. This estimate was compared to their actual accuracy to construct an overconfidence variable, which was set to 1 if their estimated recall percentage exceeded their actual recall rate. The experiment concluded with a questionnaire asking subjects' to answer questions about the nature of their choices.

The experiment was conducted on-site between November and December 2023 at the University of Southern California. Fourty-nine (49) students - twenty-five (25) females and twenty-four (24) males, mean age 22.65 (sd = .38) years - participated in the experiment. They were grouped into study sessions of a maximum of 10 participants to ensure an appropriate distance between monitor stations during the experiment. Participants received a 5$ fixed participation fee, plus one lottery randomly drawn from phase I and phase III of the experiment and one lottery randomly drawn from phase IV. The exchange rate for the first draw was 25 tokens = 1$, while the exchange rate for the second draw was 100 tokens = 1$. As a result, the average payment was 11.2$. Subjects received no feedback about earnings status prior to the end of the experiment, which is the moment they were entitled to generate random numbers to play the lotteries.

3.1.1 Mouse Tracking Procedure

We recorded participants' trajectories on a trial-by-trial basis via MouseTracker software (Freeman and Ambady, 2010). Mouse trajectories were sampled at a 70 Hz frequency, recorded from the presentation of the options until the decision click. Participants performed the experiment with a standard computer mouse. Stimuli were presented on a white background on a 20-inches screen running at a resolution of 1,280 × 1,024 pixels (75-Hz refresh frequency). At the beginning of each trial, participants clicked a "Start"

[6] The response structure and behavioral categories are present in Appendix A3, Table D3.

button located at the bottom-center of the screen (coordinate [0,0]), from which mouse trajectories originated. The horizontal (X) axis ranged from -1 (left) to 1 (right), and the vertical (Y) axis from 0 (start position at the bottom) upwards to the top of the screen. Response options ("Left" and "Right") were located in the top corners of the screen, corresponding to the X average coordinates of -.8 and .8, respectively.

The decision options remained hidden until participants pressed the "Start" button, at which point the options appeared at the top left (Option A) and right (Option B) of the screen. To ensure the decision-making process mirrored natural evidence accumulation, we adopted a dynamic starting condition where participants were instructed to actively initiate their movement, aligning their behavior with the progressive unfolding of evidence. To ensure compliance, a corrective warning message was displayed at the end of the trial if they failed to complete the required motion within the first 500 milliseconds after pushing the "Start" button.

Given this design, we used the AUC as our primary metric of interest. The AUC captures the overall deviation of the mouse trajectory from an ideal straight-line path to the selected option. Since all movements originated from a fixed starting location, deviations reflect the dynamic competition between options during the decision process. Importantly, because participants were required to initiate movement rapidly, the resulting trajectories offer a continuous and temporally rich window into the unfolding of the decision-making process. Figure 1 provides an illustration of the AUC in the context of our experimental setting

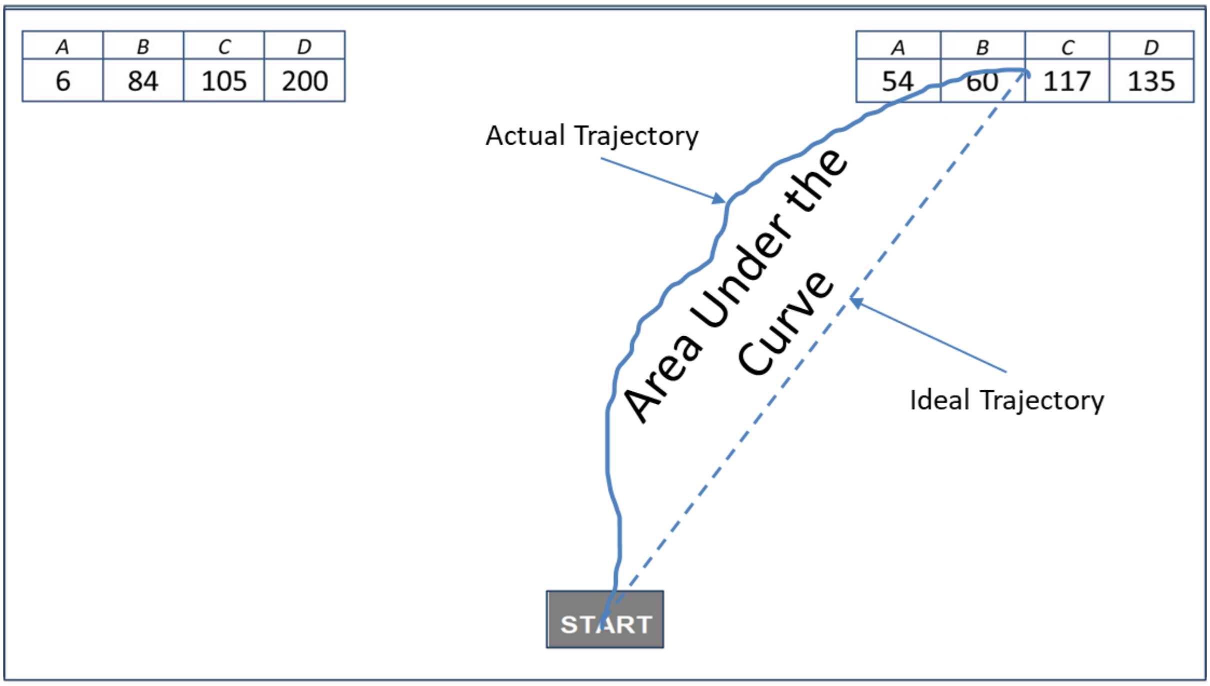

Figure 1: Illustration of the Area Under the Curve. The choice represents an actual trial from the experiment.

We selected AUC over alternative metrics, such as maximum deviation from a linear ideal path, because it more accurately reflects the continuous nature of decision dynamics. Whereas maximum deviation captures only the single most extreme point of divergence, thereby reducing the process to a discrete moment, it may overlook the overall shape and evolution of the trajectory. In contrast, AUC aggregates deviations across the entire movement path, providing a more comprehensive and temporally sensitive measure of indecision and cognitive conflict. We include average mouse trajectories in the analysis.

3.1.2 Defining Stochastic Choice

We define SC in the experimental context in two ways, one that follows the formulation of A&O and one that extends its logic on a trial-by-trial basis. The authors propose a question-wise definition of SC, assessing inconsistencies in repeated decisions within the same choice set. Under this formulation, an individual is classified as engaging in SC if their responses to the same question vary across repetitions. The authors hence provide a structured approach to capture deviations from deterministic behavior, summarizing decision variability while maintaining an aggregated perspective.

We address individual-level differences in stochastic behavior on a trial-by-trial basis, while maintaining the same underlying principle of SC computation. As in the question-wise definition, we take the first choice as the reference point. However, instead of collapsing all inconsistencies into a single indicator per question, we define switching events at the single-trial level. We exclude coin flip rounds from the analysis because the presence of the coin in the upper middle of the screen may bias the curvature of the AUC, irrespective of whether the final choice involves the coin. As a result, we can record two repetitions per question for each phase.

4. Results

4.1. Stochastic Behavior Between Phases

Figure 2 illustrates the proportion of participants displaying SC at least once in Phase I and Phase III.

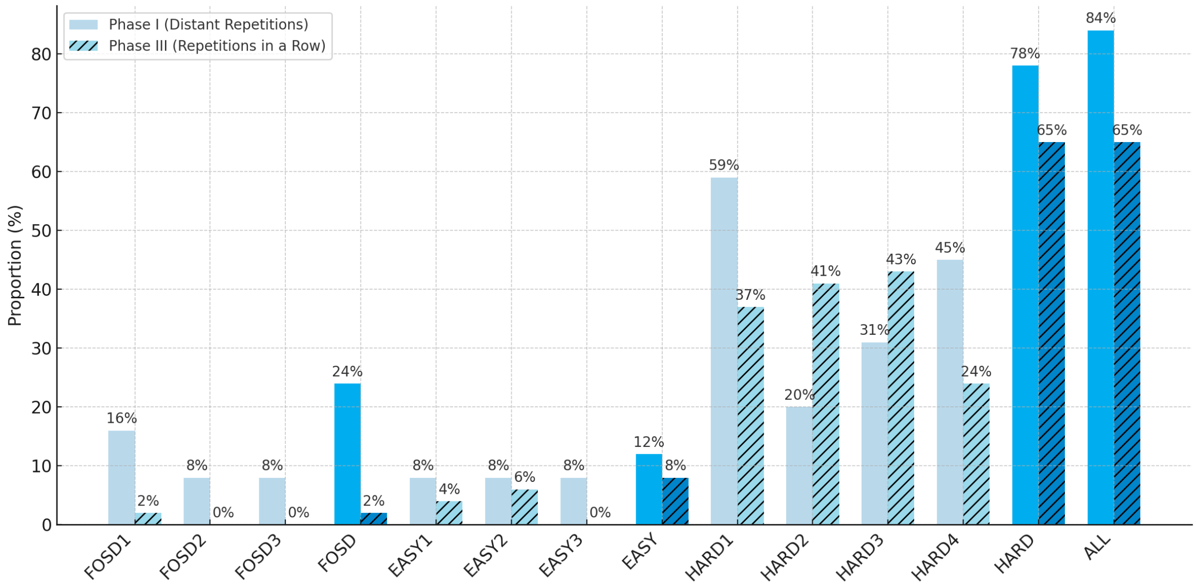

Figure 2: Fraction of Subjects Displaying Stochastic Choice at least one time, by phase.

The vast majority of subjects exhibit inconsistent behavior in phase I (84%), with many instances of SC occurring multiple times (146)(see Table I, Appendix A2). A substantial portion of participants report SC in Phase III (65%), with 99 total occurrences. All of the participants who display SC in phase III do so in phase I too. In both conditions, SC is prevalent in HARD questions (105 occurrences in phase I, 92 in phase III). However, in Phase I, 12 participants (24% of the sample) exhibit SC in questions involving a dominated option, with 25 total occurrences. This erratic behavior is no longer present in Phase III.

To measure individual-level correlation in SC across the main parts of the experiment, we employ a random effects logistic regression (Table 2[7]). While stochastic behavior between Phases I and III is significantly and positively correlated when considering all questions, this relationship is no longer significant when focusing exclusively on HARD questions[8]. Notably, for the HARD questions, we observe a different pattern of SC across the two conditions (see Fig. 2). Specifically, HARD1 and HARD4 show lower frequencies

[7]In this regression, SC is measured using a binary indicator at the question level, which equals 1 if the participant randomizes over lotteries at least once within a given question.

[8]The results are robust when excluding the 12 subjects violating stochastic dominance (Table IIa, Appendix A2).

of SC in Phase III compared to Phase I, whereas HARD2 and HARD3 exhibit higher frequencies SC in Phase III. This pattern could explain the lack of a significant correlation between the two conditions for the HARD questions.[9]

Reliance on the coin flip is strongly associated with inconsistent behavior in Phase I, only weakly in Phase III. Notably, when we run separately the behavioral analysis for the participants who violated FOSD choices in Phase I and the regular group, we find that the results gain statistical power, especially regarding the coin flip, while the relationship between SC between phases in HARD questions remains non-significant (see Appendix A2, Table IIa, except for Model 2, HARD questions). Differently, the FOSD violators show no significant correlations between SC in distant repeats and sequential repeats, no effect of the coin flip choice on SC in both conditions (Table IIb, Appendix A2).

TABLE 2: Regression Analysis of Stochastic Choice

	Stochastic Choice Distant Repeats (Part I)			Stochastic Choice Sequential Repeats (Part III)
	(1)	(2)	(3)	(4)
Stochastic choice in sequential reps	0.20*** (0.06)	0.05 (0.08)		
Flip coin			0.14*** (0.06)	0.11* (0.06)
Constant	0.22*** (0.03)	0.37*** (0.006)	0.19*** (0.02)	0.20*** (0.03)
Observations	343	196	490	343
N subjects	49	49	49	49
R^2	0.032	0.0032	0.02	0.006
Sample of questions	All	HARD	All	All

Note: ***p<0.01, **p<0.05, *p<0.1. Random-effects GLS regressions with standard errors clustered at the subject level.

Next, we address whether SC can be explained by the expected utility difference between options, using specifications based on risk-neutral, CRRA, and CARA utility functions[10]. Results, reported in Tables IIIa and IIIb of Appendix A2, indicate that although SC is more prevalent in HARD questions, question type difficulty is not fully captured by expected utility differences in either phase. This suggests that sources of

[9]HARD2 and HARD3 differ from the other two HARD questions in terms of differences in standard deviations between lotteries (i.e., risk; see Table 1). Thus, in Phase III, where SC deliberation may be stronger, these specific features of the lotteries could shape participants' SC behavior differently across the HARD questions.

[10]This result is consistent with Agranov and Ortoleva (2017).

complexity beyond simple differences in expected value, transformed under risk-averse subjective utility, may be driving SC[11].

We move to a summary of the main replication results (see also Appendix A2), highlighting both consistencies and differences relative to Agranov and Ortoleva (2017).

Result 0
Replications (R):

- (R1) Prevalence of stochastic choice (SC) in both phases, with higher rates in distant repeats.
- (R2) Strong correlation between SC in Phase I and Phase III.
- (R3) Strong correlation between SC and coin-flip selection in trials that allow randomization.
- (R4) RTs decrease following the first choice in both conditions.
- (R5) RTs remain larger in the distant-repeats phase.
- (R6) Differences in expected value do not fully account for stochastic behavior in HARD choices.

Differences (D):

- (D1) Twenty-four percent of subjects display SC in FOSD questions in Phase I, whereas A&O report virtually no randomization in either phase for these items.
- (D2) No correlation is observed between SC in Phase I and Phase III for HARD questions.
- (D3) A lower proportion of participants violate expected utility in common-consequences and common-ratio questions (8% vs. 25% in A&O).

4.2. Characterization of Decision Processes underlying repetitions structures: Patterns of Mouse Trajectories and RT

The present section analyzes sources of stochastic choice using process data. We begin by examining potential differences in the decision-making process across the two phases, by

[11] A note of caution is warranted here: the structure of expected value differences across items, particularly within HARD questions, shows limited variability, which may restrict the explanatory power of such factors in accounting for observed heterogeneity in SC.

analyzing the evolution of mean RT and AUC across choice repetitions[12]. Figure 3 and Table IV, Appendix A2, illustrate the differences in RT across phases and repetitions [13].

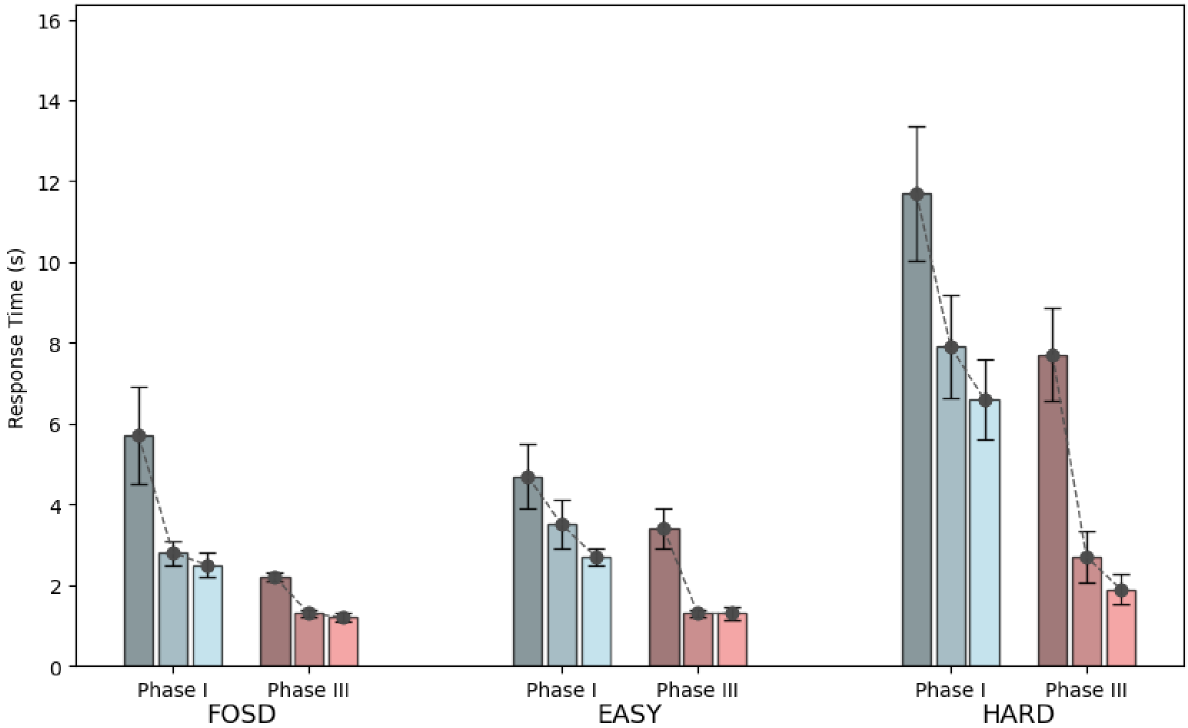

Figure 3: Evolution of Response Time across repetitions, by phase.

A pattern shared by both phases is that of decisions related to HARD question types, taking significantly longer times than those related to FOSD and EASY ones ($p < .01$ RE GLS Regression Table IVb, Appendix A2, rows 9 to 12). Also, RTs peak during the first choice and drop subsequently in both phase I and III consistently across all question types ($p < .01$ RE GLS Regression Table IVb, Appendix A2, rows 13,15,17,19,21 and 23). However, decisions in Phase I take longer than those in Phase III in all questions and order of repetition ($p < 0.001$ Mann-Whitney U-test).

Result 1. RT: Across both phases, HARD items consistently elicited longer response times than FOSD and EASY items, RTs peaked on the first trial and declined thereafter for all question types, and overall decisions were slower in Phase I than in Phase III.

The analysis of AUC distributions reveals clear differences between repetitions across the two phases (Figure 4, Table V and Table Vb, Appendix A2 for Statistical Analysis).

[12] In Phase I, we exclude from the analysis the repetition that includes the coin-flip option, consistent with the original study and with the concern that the coin, displayed in the top-center of the screen, would distort the interpretation of the AUC data.

[13] These differences were tested formally using a random-effects GLS regression, with detailed results reported in Table IVb of Appendix A2. The p-values shown correspond to rows 9 through 24 of the table.

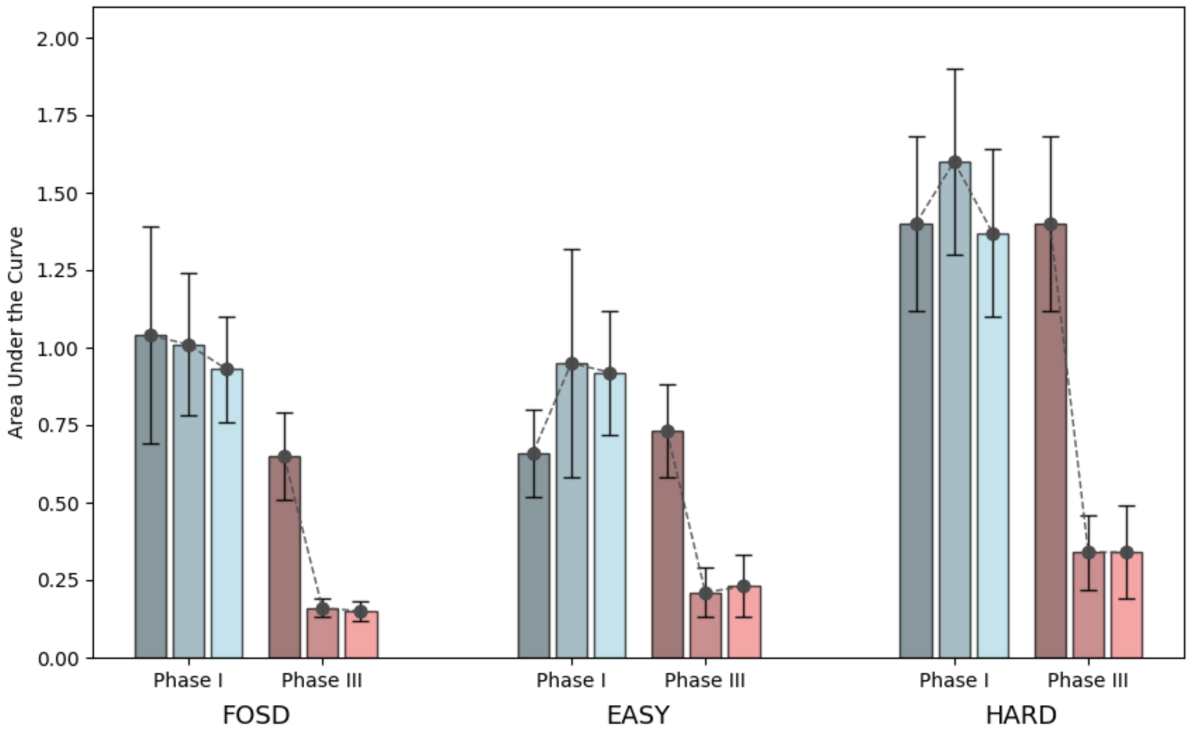

Figure 4: Evolution of Area Under the Curve across repetitions, by phase.

In phase I, the AUC shows no signs of reductions across repetitions for all question types ($p > .1$ RE GLS Regression Table Vb, rows 13-18). In Phase III, in contrast, AUC peaks at the first choice and drops sharply at subsequent repeats ($p < .001$ RE GLS Table Vb, Appendix A2), reaching its minimum throughout the first repeat and remaining stable in the second ($p >> 1$ RE GLS Table Vb, Appendix A2).

Result 2. AUC: In Phase I, AUC remains stable across repetitions for all question types, whereas in Phase III it peaks at the first choice, drops sharply thereafter, and stabilizes at its lowest level from the first repeat onward.

A common pattern emerges for RTs and AUC: choices in HARD questions exhibit larger AUC values than those in EASY and FOSD questions across both conditions (RE GLS, Table Vb, Appendix A2), and AUC is systematically higher in Phase I than in Phase III. This pattern is illustrated in Figure 5, which compares the average trajectories of Phase I with those of Phase III for the same question types[14].

[14]Figure 6 and 5 depict the average mouse trajectory along the experiment's coordinate axes. We computed the average trajectory by averaging out the subject-specific mean by question type of each of the 101 points recorded along both the x- and y-axes.

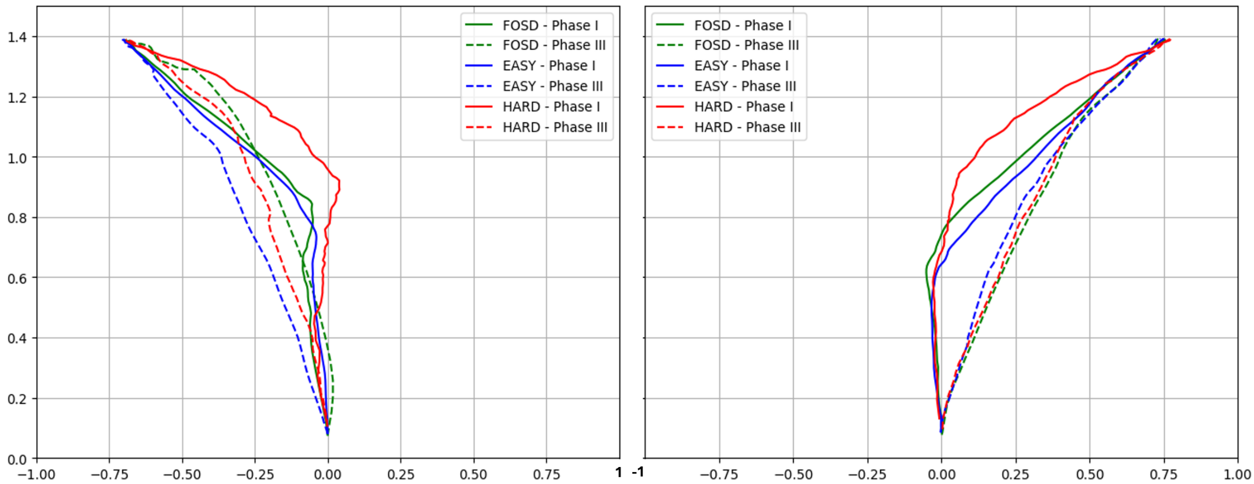

Figure 5: Average mouse trajectories of all trials, by question type (FOSD, EASY, HARD). The continuous lines represent Phase I trials; the dashed lines represent Phase III trials.

We further examine differences in AUC trajectories between the first choice and its repetitions across the two phases of the experiment. Figure 6 (and Table Vb in Appendix 2 for statistical analysis) shows no change in AUC between the first choice and its repeats in Phase I, whereas in Phase III the AUC is substantially higher on the first choice and markedly lower on subsequent repeats. This pattern is especially pronounced for HARD choices, for which the trajectories in Phase I exhibit near-perfect overlap across repetitions. In addition, comparing the repeat trajectories across phases reveals a sharp reduction in AUC and considerably more direct movement paths in Phase III.

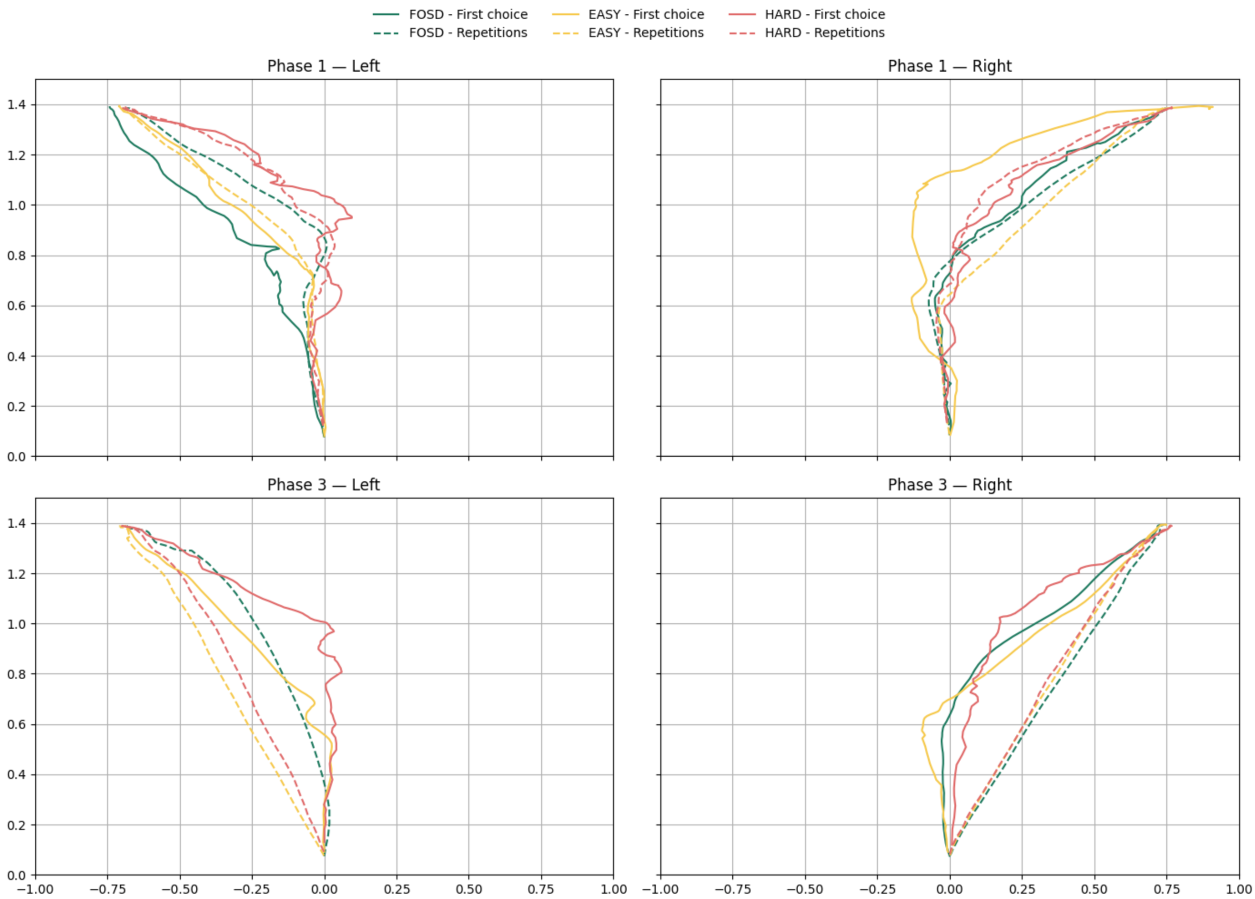

Figure 6: Average mouse trajectories, by question type (FOSD, EASY, HARD). The continuous lines represent the 1st trials, or reference choice; the dashed lines represent repeats of the same question type.

We thus observe a stable pattern of AUC signal across repetitions in Phase I, particularly for HARD questions where this stability is most pronounced (Figure 6). In contrast, the distinct AUC pattern observed in Phase III, characterized by a pronounced peak on the first choice and sharply reduced AUC on subsequent repeats, may reflect the role of planning and explicit deliberation enabled by the known sequence of three repetitions.

We now examine whether differences in the decision processes (using RT and AUC) between the two phases exist with respect to stochastic behavior. We investigate the trial-by-trial explanatory power of RT and AUC in accounting for the variation between SC trials against consistent trials. Table 3 examines the impact of RT and AUC on the probability of switching choice, using the single-trial specification of SC as a dependent variable[15].

While both AUC and RT increase with question difficulty (see Table Ia in Appendix A2), consistent with the interpretation that higher complexity increases the uncertainty in value comparison, AUC outperforms RT in predicting the likelihood of SC. In fact,

[15] Refer to Section 3.2.2 for the single-trial specification of SC

TABLE 3: Random-effects Logit for SC, Phase I and Phase III

	All		**HARD**	
SC	**Phase I**	**Phase III**	**Phase I**	**Phase III**
AUC	0.165***	0.825***	0.113*	0.611***
	(0.05)	(0.17)	(0.067)	(0.19)
RT	0.02	-0.022	0.016	-0.021
	(0.02)	(0.04)	(0.017)	(0.042)
Difficulty	0.816***	2.5***	—	—
	(0.14)	(0.42)		
Intercept	-4.135***	-9.36***	-1.47***	-1.815***
	(0.38)	(1.27)	(0.25)	(0.304)
Obs.	980	980	392	392
Groups	49	49	49	49
Prob > χ^2	0.000	0.000	0.000	0.000
Questions	All	All	HARD	HARD

Notes: Random-Effects Logistic Regression with standard error Clustered at the individual level. Here, we considered trials that not including the coin flip option and defined SC = 1 if the choice in that trial differed from the first choice.

RT shows no significant relationship with SC across either phases when controlling for difficulty level. The coefficient for AUC is instead highly significant across both phases, stronger for Phase III, while it is only weakly significant in Phase I for HARD choices. Notably, the effect of AUC on stochastic behavior in Phase I is driven exclusively by subjects displaying SC in the presence of FOSD lotteries (regression Table VIIIb, Appendix A2). Also, these participants show an incoherent stochastic behavior between conditions, randomizing substantially in Part I while reversing their behavior in Part III, where they show on average only one switch (Table 4; Wilcoxon Signed Rank test: $p < .05$).

In contrast, the significant relationship between AUC and SC in Phase III is sustained only by regular subjects (i.e., subjects who did not violate FOSD, regression Table VIIIa, Appendix A2), whose average switching behavior is consistent across phases (Table 4[16]; Wilcoxon Signed Rank test, $p \approx .90$).

Result 3. AUC vs. RT in SC: AUC outperforms RT in predicting the likelihood of stochastic choice, providing a substantially stronger and more reliable process-level indicator of when subjects will randomize.

Differences in Mouse Tracking trajectories between SC and consistent choices across

[16]To measure mean absolute frequency of SC in Phase I meaningfully, questions FOSD2, FOSD3 and EASY3 were excluded, as these were not displayed to subjects in Phase III. By including these questions in the analysis, the average difference becomes even more significant for FOSD violators (5 in Phase I), while remains consistent for the Regular Subjects (2.32 in phase III).

the two subject groups are illustrated in Figures 7 and 8. These figures compare trials in which subjects made SC with those in which they responded consistently, separately for phase I (Figure 7) and phase III (Figure 8), and visually support the patterns identified through logistic regressions. The regular group indeed shows no difference in AUC between SC and non-SC trials in Phase I (i.e. invariance), while displaying a larger AUC for SC trials in Phase III. In contrast, participants who violated FOSD in Phase I exhibit a significantly larger and more erratic mouse-tracking trajectory for SC compared to non-SC trials in Phase I, while showing no difference in Phase III, where they also exhibit very few SC trials.

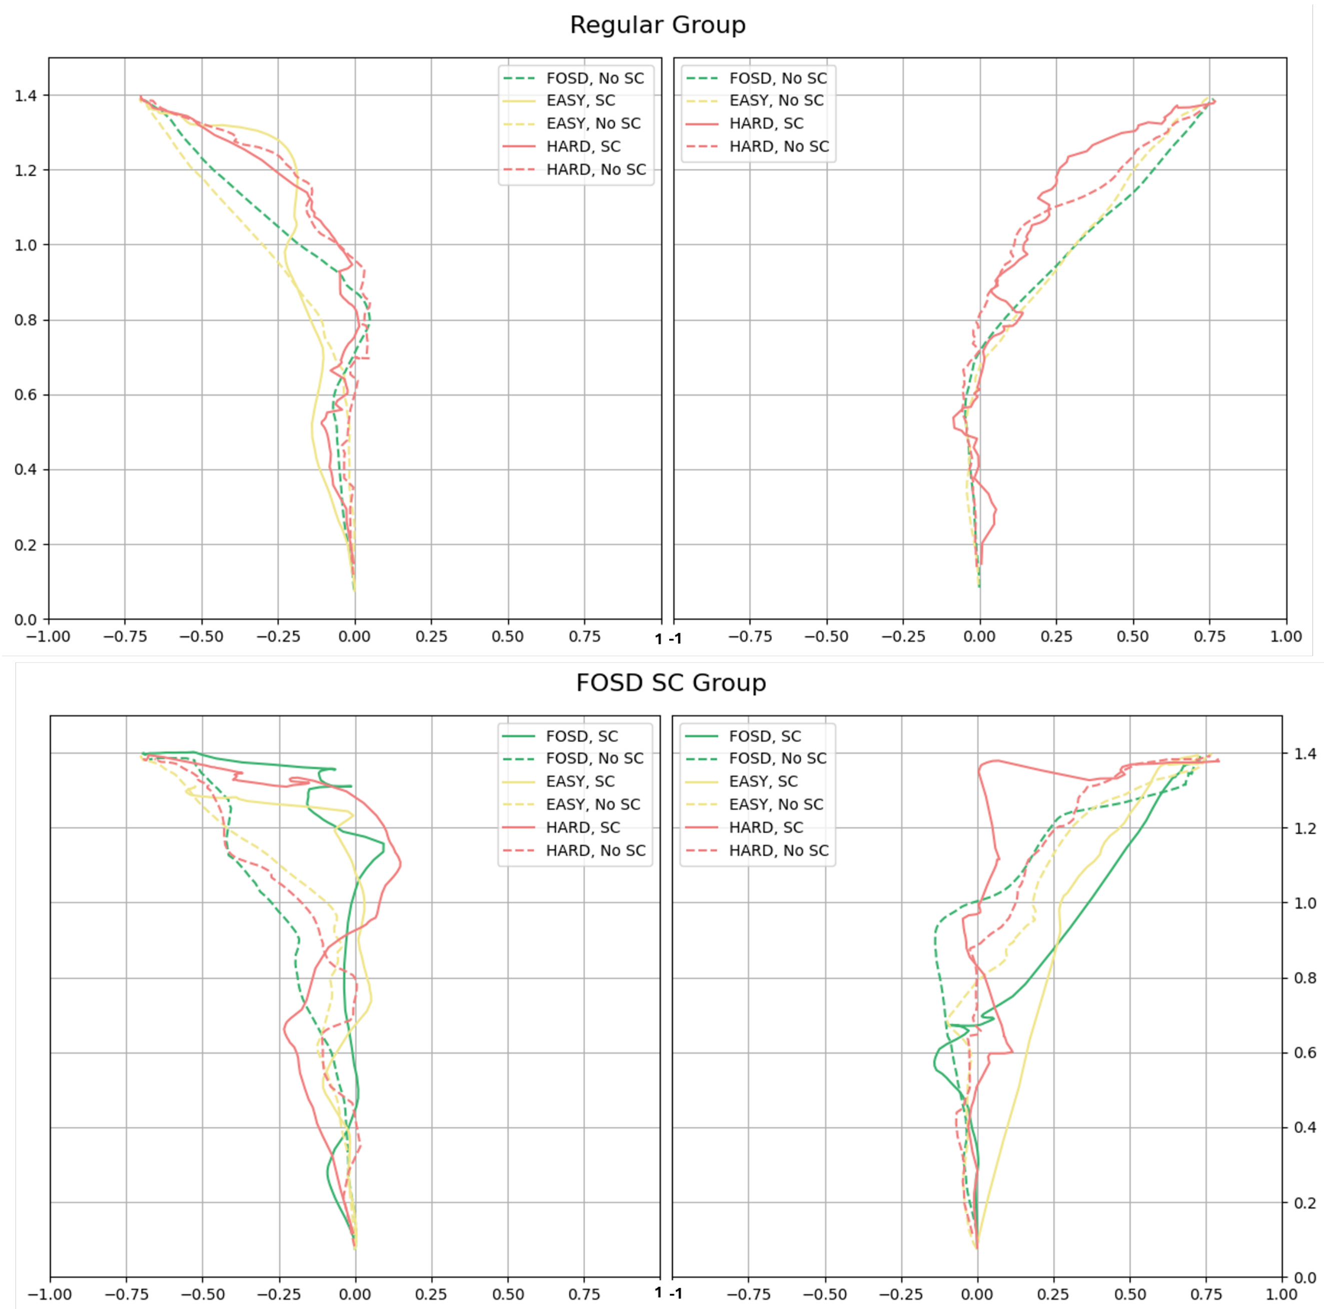

Figure 7: Average mouse trajectories of Phase I, by question type (FOSD, EASY, HARD). The continuous lines represent SC trials; the dashed lines display no-SC trials. Regular subjects are displayed above; subjects displaying SC in FOSD question types are displayed below.

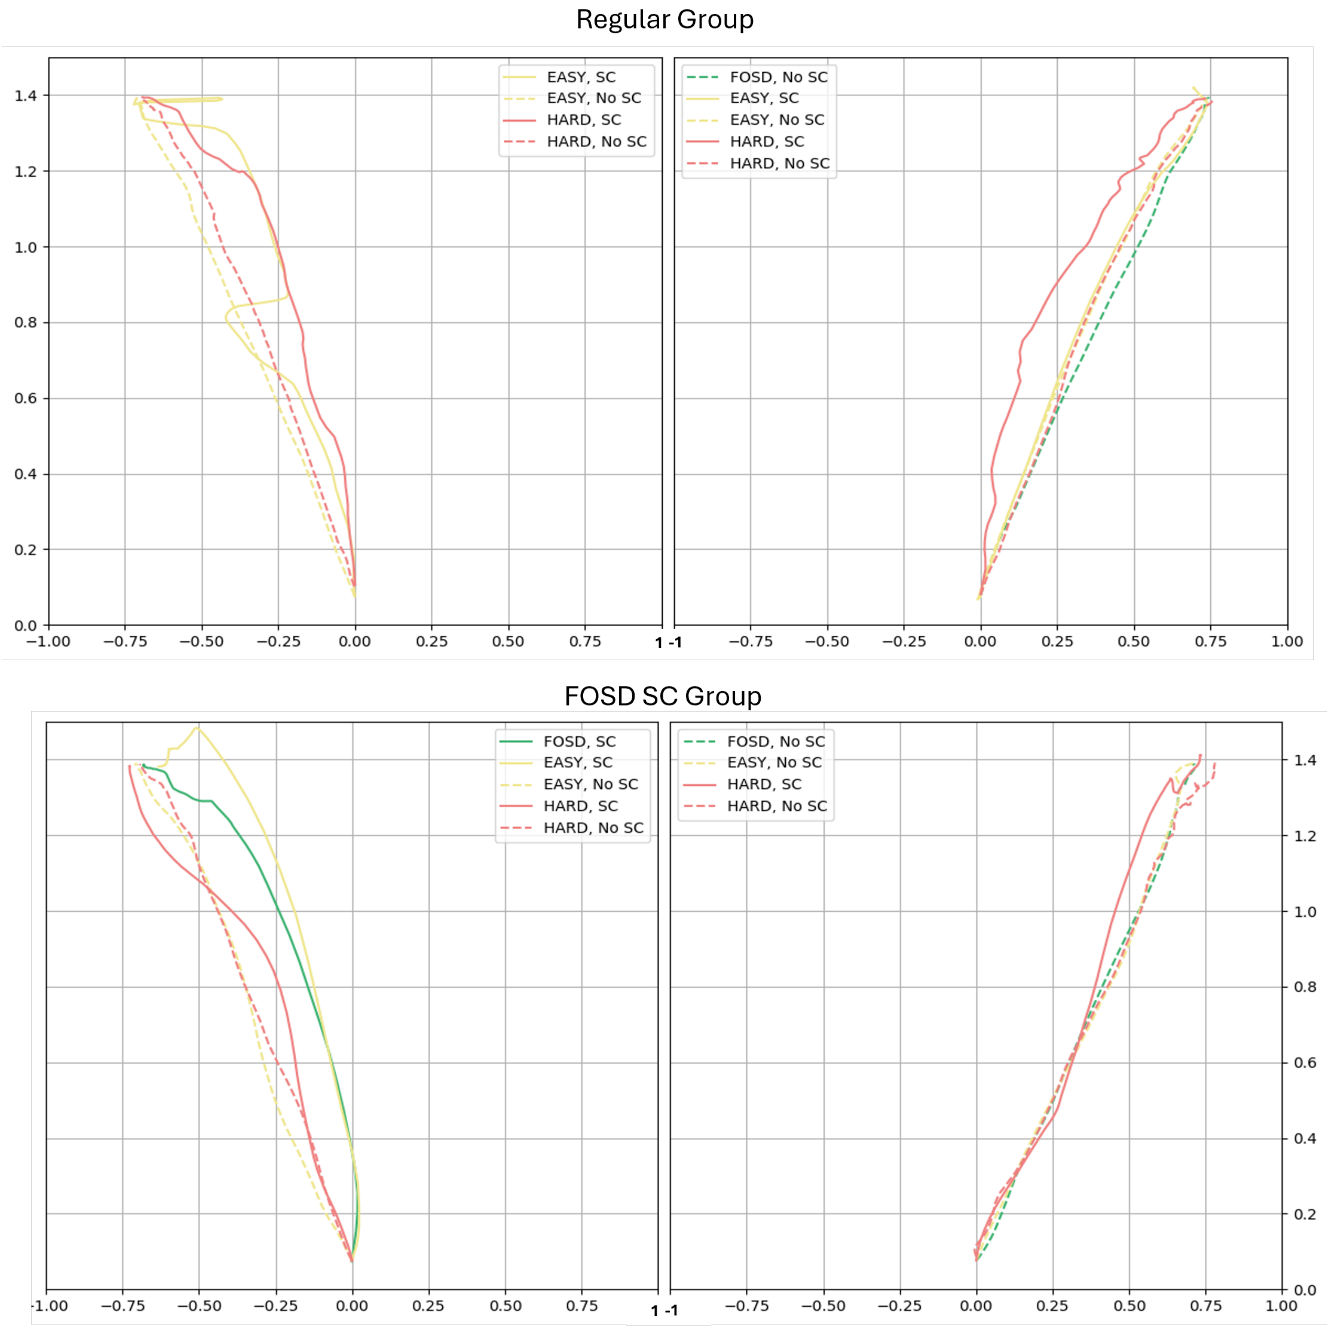

Figure 8: Average mouse trajectories of Phase III, by question type (FOSD, EASY, HARD). The continuous lines represent SC trials; the dashed lines display no-SC trials. Regular subjects are displayed above; subjects displaying SC in FOSD question types are displayed below.

TABLE 4: Mean (SD) Frequency of SC by group.

Mean SC occurrences	Phase I	Phase III
FOSD SC Group	3.59 (.08)	1.08 (0.05)
Regular Group	2.27 (.033)	2.32 (.04)

The differences between the two groups of participants in terms of AUC measured during SC trials in Phase I and the behavior in Phase III, show how the pattern of AUC in Phase I may act as a negative predictor of SC in Phase III. This is supported by a significant negative correlation between the difference in AUC across SC and non-SC trials in Phase I and the individual-level frequency of SC in the subsequent phase (Table 5). Thus, participants with larger AUC in SC compared to non-SC in Phase I show less instances of SC in Phase III for HARD choices.

TABLE 5: Correlation between the difference in AUC across SC vs. non-SC trials in Phase I and the individual-level frequency of SC in Phase III.

Difficulty	All	p-value
FOSD	-0.303	0.338
EASY	0.150	0.679
HARD	-0.314	0.038**

5. Discussion and Conclusions

In this study of stochastic behavior in choice under risk, we replicated the experimental framework of A&O, whose work is notable for being the first to compare, within subjects, a distant repeats scheme without providing participants with explicit information about the repetition structure (Phase I) and a sequential repetition scheme with full disclosure of the repetition structure (Phase III).

We leveraged mouse-tracking to extract process tracing metrics of choice, specifically analyzing AUC in addition to RT, on a trial-by-trial basis. We completed this analysis with post-experimental tests aimed at examining the relationship between preferences for convex combinations, choice precision, and memory with SC (see Appendix A3).

Our replication closely aligns with the main findings of A&O, while also revealing several noteworthy deviations. First, we confirm the prevalence and stability of SC across both phases, its strong correlation with coin-flip use, and the expected response-time patterns, namely, longer RTs for HARD items, sharp RT reductions after the first choice, and overall slower decisions in the distant-repeats condition. We also replicate that differences in expected value do not fully explain SC in HARD choices. However, our data differ from A&O in three respects: we observe substantial SC in FOSD items in Phase I, no correlation in SC across phases for HARD items, and a markedly lower incidence of expected-utility violations (Allais-type behavior). Process measures further refine this picture: AUC remains stable across repetitions in Phase I but declines sharply after the first choice in Phase III; and AUC consistently outperforms RT in predicting the likelihood of SC, offering a more sensitive process-level diagnostic of randomization behavior. Moreover, AUC clearly differentiated first from subsequent choices in both conditions, exhibiting stable trajectories in distant repeats (i.e., no learning-related decreases) and revealing planning-related processing in the sequential-repeats condition, where AUC peaked on the first choice and declined sharply thereafter.

The use of mouse-tracking within the dual-repetitions experimental settings represents our main contribution. We employed this method primarily to compare the processes underlying choice across phases, and AUC proved highly informative for this purpose, revealing notable differences in patterns between Phase I and Phase III. In Phase I, we observed values of AUC that remained stable between repetitions. This pattern indicates differences in the effectiveness of conflict resolution (Konovalov and Krajbich, 2020), suggesting that distant repetitions retain at least a more complex process of information retrieval. In contrast, during Phase III, AUC decreased sharply from the first choice compared to the following repetitions, suggesting a decision process whose main computation happens during the first choice, which might also incorporate planning of subsequent choices.

The second motivation for employing mouse-tracking was to gain insight into the mechanisms underlying stochastic behavior, distinguishing error-driven fluctuations from deliberative stochastic processes. In Phase I, the stability of AUC across repetitions and its invariance between SC and non-SC trials for the regular group point to an across-trial balancing mechanism consistent with CSC-type models. By contrast, the FOSD violators exhibit higher AUC and more erratic trajectories for SC relative to non-SC trials, indicating an error-driven source of variability. In Phase III, AUC emerges as a strong predictor of SC, an effect driven primarily by the regular group, suggesting that stochastic choices under sequential repeats involve a more planning-related form of deliberation, consistent with CSC-like accounts.

Moreover, the regular group of participants (the strong majority of our sample) showed an identical relative frequency of SC across both phases of the experiment, indicating that, even though the mechanisms at play may differ throughout the choice process, the aggregate randomization ends up being consistent. In contrast, the group of participants making FOSD errors do not display any of these regular patterns, but an erratic large SC frequency in the distant repeats and an aversion to randomize in the sequential repeats phase.

In A&O, the support for the conclusion regarding the deliberative nature of SC relied on the assumption that SC in Phase III is unambiguously deliberative (a la CSC model), and that a correlation of SC across both conditions would indicate a shared deliberative basis. Our mouse-tracking data complement this approach by showing that stochastic choices (in the regular group) in Phase I do not reflect noisier or more conflict-driven decision processes compared to consistent choices. This is reflected in the stability of AUC across repetitions and its invariance between SC and non-SC trials, which supports the interpretation of stochastic behavior as deliberative in both conditions of our experiment.

More in general, AUC provides a clear measure of choice complexity, with its strong correlation with increasing choice difficulty. Our data show how choice imprecision, reduced memory, and convex preference contribute to explaining the pattern of stochastic choice in decision making under risk.

6. Limitations of the Experimental Design

While the A&O design offers a powerful framework for studying stochastic choice, it also imposes several limitations on our ability to distinguish among competing theoretical models. First, the small number of repetitions within each choice set constrains the evaluation of dynamic stochastic-choice models, particularly diffusion-based approaches such as the DDM, which require richer within-item time series to identify trial-by-trial updating or sequential dependencies. Moreover, the design relies on a narrowly defined set of lotteries. A more comprehensive and systematically varied set, including broader probability ranges, presence/absence of zero outcomes, and more heterogeneous payoff structures, would be necessary to evaluate competing models of SC across a richer decision space. In the Appendix, we report a model-fitting exercise comparing two highly simplified frameworks: expected utility with noise and a minimally deliberative model allowing for coin-flipping at quasi-indifference. Although the EU-plus-error model performs better once differences in parameterization are taken into account, this advantage is not decisive. Given the conceptual oversimplification of both models, the analysis remains inconclusive. Future work should employ designs with many more repetitions, integrate process-tracing measures with behavioral data, and include procedures capable of distinguishing the motives driving deliberation, such as regret avoidance, risk hedging, or strategic exploration.

References

Agranov, M., Healy, P. J., and Nielsen, K. (2023). Stable randomisation. *The Economic Journal*, 133(655):2553–2579.

Agranov, M. and Ortoleva, P. (2017). Stochastic choice and preferences for randomization. *Journal of Political Economy*, 125(1):40–68.

Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école américaine. *Econometrica: journal of the Econometric Society*, pages 503–546.

Allen, R. and Rehbeck, J. (2023). Revealed stochastic choice with attributes. *Economic Theory*, 75(1):91–112.

Alós-Ferrer, C., Garagnani, M., and Hügelschäfer, S. (2016). Cognitive reflection, decision biases, and response times. *Frontiers in psychology*, 7:1402.

Arieli, A., Ben-Ami, Y., and Rubinstein, A. (2011). Tracking decision makers under uncertainty. *American Economic Journal: Microeconomics*, 3(4):68–76.

Ballinger, T. P. and Wilcox, N. T. (1997). Decisions, error and heterogeneity. *The Economic Journal*, 107(443):1090–1105.

Bault, N., Wydoodt, P., and Coricelli, G. (2016). Different attentional patterns for regret and disappointment: An eye-tracking study. *Journal of Behavioral Decision Making*, 29(2-3):194–205.

Becker, G. M., DeGroot, M. H., and Marschak, J. (1964). Measuring utility by a single-response sequential method. *Behavioral science*, 9(3):226–232.

Block, H. D. and Marschak, J. (1959). Random orderings and stochastic theories of response.

Börger, T. (2016). Are fast responses more random? testing the effect of response time on scale in an online choice experiment. *Environmental and Resource Economics*, 65(2):389–413.

Brady, R. L. and Rehbeck, J. (2016). Menu-dependent stochastic feasibility. *Econometrica*, 84(3):1203–1223.

Busemeyer, J. R. and Townsend, J. T. (1993). Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychological review*, 100(3):432.

Camerer, C. (1998). Bounded rationality in individual decision making. *Experimental economics*, 1:163–183.

Camerer, C. F. (1989). An experimental test of several generalized utility theories. *Journal of Risk and uncertainty*, 2:61–104.

Camerer, C. F., Johnson, E., Rymon, T., and Sen, S. (1993). Cognition and framing in sequential bargaining for gains and losses. *Frontiers of game theory*, 104:27–47.

Caplin, A. and Dean, M. (2011). Search, choice, and revealed preference. *Theoretical Economics*, 6(1):19–48.

Caplin, A., Dean, M., and Leahy, J. (2019). Rational inattention, optimal consideration sets, and stochastic choice. *The Review of Economic Studies*, 86(3):1061–1094.

Cerreia-Vioglio, S., Dillenberger, D., and Ortoleva, P. (2015). Cautious expected utility and the certainty effect. *Econometrica*, 83(2):693–728.

Cerreia-Vioglio, S., Dillenberger, D., Ortoleva, P., and Riella, G. (2019). Deliberately stochastic. *American Economic Review*, 109(7):2425–2445.

Cettolin, E. and Riedl, A. (2019). Revealed preferences under uncertainty: Incomplete preferences and preferences for randomization. *Journal of Economic Theory*, 181:547–585.

Coricelli, G., Polonio, L., and Vostroknutov, A. (2020). The process of choice in games. In *Handbook of experimental game theory*, pages 69–94. Edward Elgar Publishing.

Cubitt, R. P., Navarro-Martinez, D., and Starmer, C. (2015). On preference imprecision. *Journal of Risk and Uncertainty*, 50:1–34.

Dai, J. and Busemeyer, J. R. (2014). A probabilistic, dynamic, and attribute-wise model of intertemporal choice. *Journal of Experimental Psychology: General*, 143(4):1489.

Dshemuchadse, M., Scherbaum, S., and Goschke, T. (2013). How decisions emerge: action dynamics in intertemporal decision making. *Journal of Experimental Psychology: General*, 142(1):93.

Dwenger, N., Kübler, D., and Weizsäcker, G. (2012). Flipping a coin: Theory and evidence.

Dwenger, N., Kübler, D., and Weizsäcker, G. (2018). Flipping a coin: Evidence from university applications. *Journal of Public Economics*, 167:240–250.

Feldman, P. and Rehbeck, J. (2022). Revealing a preference for mixtures: An experimental study of risk. *Quantitative Economics*, 13(2):761–786.

Fiedler, S. and Glöckner, A. (2012). The dynamics of decision making in risky choice: An eye-tracking analysis. *Frontiers in psychology*, 3:335.

Freeman, J. B. and Ambady, N. (2010). Mousetracker: Software for studying real-time mental processing using a computer mouse-tracking method. *Behavior research methods*, 42(1):226–241.

Freeman, J. B., Ambady, N., Rule, N. O., and Johnson, K. L. (2008). Will a category cue attract you? motor output reveals dynamic competition across person construal. *Journal of Experimental Psychology: General*, 137(4):673.

Fudenberg, D. and Strzalecki, T. (2015). Dynamic logit with choice aversion. *Econometrica*, 83(2):651–691.

Glöckner, A. and Herbold, A.-K. (2011). An eye-tracking study on information processing in risky decisions: Evidence for compensatory strategies based on automatic processes. *Journal of Behavioral Decision Making*, 24(1):71–98.

Gneezy, U. and Potters, J. (1997). An experiment on risk taking and evaluation periods. *The quarterly journal of economics*, 112(2):631–645.

Gold, J. I. and Shadlen, M. N. (2007). The neural basis of decision making. *Annu. Rev. Neurosci.*, 30(1):535–574.

Gul, F., Natenzon, P., and Pesendorfer, W. (2014). Random choice as behavioral optimization. *Econometrica*, 82(5):1873–1912.

Gul, F. and Pesendorfer, W. (2006). Random expected utility. *Econometrica*, 74(1):121–146.

Harrison, G. W. and Swarthout, J. T. (2019). Eye-tracking and economic theories of choice under risk. *Journal of the Economic Science Association*, 5(1):26–37.

Hey, J. D. and Orme, C. (1994). Investigating generalizations of expected utility theory using experimental data. *Econometrica: Journal of the Econometric Society*, pages 1291–1326.

Heydari, P. (2024). Regret, responsibility, and randomization: A theory of stochastic choice. *Journal of Economic Theory*, 217:105824.

Johnson, J. G. and Busemeyer, J. R. (2005). A dynamic, stochastic, computational model of preference reversal phenomena. *Psychological review*, 112(4):841.

Kocher, M. G., Pahlke, J., and Trautmann, S. T. (2013). Tempus fugit: Time pressure in risky decisions. *Management Science*, 59(10):2380–2391.

Konovalov, A. and Krajbich, I. (2016). Gaze data reveal distinct choice processes underlying model-based and model-free reinforcement learning. *Nature communications*, 7(1):12438.

Konovalov, A. and Krajbich, I. (2020). Mouse tracking reveals structure knowledge in the absence of model-based choice. *Nature communications*, 11(1):1893.

Krajbich, I., Armel, C., and Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nature neuroscience*, 13(10):1292–1298.

Krajbich, I., Bartling, B., Hare, T., and Fehr, E. (2015). Rethinking fast and slow based on a critique of reaction-time reverse inference. *Nature communications*, 6(1):7455.

Krajbich, I., Lu, D., Camerer, C., and Rangel, A. (2012). The attentional drift-diffusion model extends to simple purchasing decisions. *Frontiers in psychology*, 3:193.

Krajbich, I. and Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences*, 108(33):13852–13857.

Luce, R. D. (1959). *Individual choice behavior*, volume 4. Wiley New York.

Machina, M. J. (1985). Stochastic choice functions generated from deterministic preferences over lotteries. *The economic journal*, 95(379):575–594.

McFadden, D. and Richter, M. K. (1990). Stochastic rationality and revealed stochastic preference. *Preferences, Uncertainty, and Optimality, Essays in Honor of Leo Hurwicz, Westview Press: Boulder, CO*, pages 161–186.

Mosteller, F. and Nogee, P. (1951). An experimental measurement of utility. *Journal of political economy*, 59(5):371–404.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological review*, 85(2):59.

Ratcliff, R. and McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural computation*, 20(4):873–922.

Regenwetter, M., Dana, J., and Davis-Stober, C. P. (2011). Transitivity of preferences. *Psychological review*, 118(1):42.

Rubinstein, A. (2013). Response time and decision making: An experimental study. *Judgment and Decision making*, 8(5):540–551.

Samuelson, P. A. (1938). A note on the pure theory of consumer's behaviour. In *The Foundations of Price Theory Vol 4*, pages 101–116. Routledge.

Scherbaum, S., Dshemuchadse, M., Fischer, R., and Goschke, T. (2010). How decisions evolve: The temporal dynamics of action selection. *Cognition*, 115(3):407–416.

Spivey, M. J., Grosjean, M., and Knoblich, G. (2005). Continuous attraction toward phonological competitors. *Proceedings of the National Academy of Sciences*, 102(29):10393–10398.

Stewart, N., Hermens, F., and Matthews, W. J. (2016). Eye movements in risky choice. *Journal of behavioral decision making*, 29(2-3):116–136.

Stillman, P. E., Shen, X., and Ferguson, M. J. (2018). How mouse-tracking can advance social cognitive theory. *Trends in cognitive sciences*, 22(6):531–543.

Sullivan, N., Hutcherson, C., Harris, A., and Rangel, A. (2015). Dietary self-control is related to the speed with which attributes of healthfulness and tastiness are processed. *Psychological science*, 26(2):122–134.

Thurstone, L. L. (1927). Psychophysical analysis. *The American journal of psychology*, 38(3):368–389.

Tsetsos, K., Usher, M., and Chater, N. (2010). Preference reversal in multiattribute choice. *Psychological review*, 117(4):1275.

Tversky, A. (1969). Intransitivity of preferences. *Psychological review*, 76(1):31.

Von Neumann, J. and Morgenstern, O. (1944). Theory of games and economic behavior. In *Theory of games and economic behavior*. Princeton university press.

Woodford, M. (2020). Modeling imprecision in perception, valuation, and choice. *Annual Review of Economics*, 12(1):579–601.