Journal of Economic Behavior and Organization 248 (2026) 107632

Contents lists available at ScienceDirect

Journal of Economic Behavior and Organization

journal homepage: www.elsevier.com/locate/jebo

Deliberate stochastic choice: A mouse-tracking experiment☆

Mattia Adamo [a],*, Giorgio Coricelli [b]

[a] IMT School for Advanced Studies, Lucca, Italy
[b] University of Southern California, Los Angeles, United States of America

ARTICLE INFO

Dataset link: Replication Material for Deliberate Stochastic Choice: A Mouse-Tracking Experiment (Original data)

Keywords:
Stochastic choice
Randomization
Choice under risk
Mouse tracking

ABSTRACT

Stochastic choice (SC)—the tendency to select different options across repeated presentations of the same choice set—is well documented, yet its cognitive mechanisms remain debated. We combine mouse-tracking with the experimental design of Agranov and Ortoleva (2017), comparing distant and sequential repetitions of lottery choices. We show that the Area Under the Curve (AUC) provides a more informative process-level measure of SC than response times (RT). RT decline after the first presentation in both repetition structures, but do not reliably predict trial-level switching. AUC remains stable across distant repetitions while drops sharply after the first sequential choice, suggesting that explicit repetitions allow participants to plan subsequent choices at first and execute them more directly thereafter. AUC also reveals heterogeneity hidden by aggregate choice data: among regular participants, SC does not increase AUC in distant repeats, but produces larger and more curved trajectories in sequential repeats, suggesting a signature of cautiousness under explicit randomization. Among dominance violators, SC is instead accompanied by larger and more erratic trajectories, suggesting error-driven variability. Overall, SC appears to combine a predominantly deliberate component with a smaller boundedly rational component.

1. Introduction

When decision-makers are repeatedly asked to choose from the same set of options, they often select different alternatives, a phenomenon documented as stochastic choice (SC). In economic settings, this behavior is often observed in lottery choices (Mosteller and Nogee, 1951; Luce, 1959; Hey and Orme, 1994; Regenwetter et al., 2011), with certain behavioral patterns that have challenged traditional models of rationality (Samuelson, 1938; Von Neumann and Morgenstern, 1944; Tversky, 1969; Luce, 1959; Camerer, 1989), paving the way to studies that merged multiple domains such as economics, psychology and neuroscience to explain the reasons of these inconsistencies (Cerreia-Vioglio et al., 2015; Fudenberg and Strzalecki, 2015; Caplin et al., 2019; Woodford, 2020; Allen and Rehbeck, 2023; Agranov et al., 2023). First, individuals might randomize choice because their behavior is governed by multiple utility functions (Becker et al., 1964; Gul and Pesendorfer, 2006). Here, stochasticity arises from which utility function governs the decision, with the selection process itself subject to unobservable random shocks. Second, SC may arise from internal cognitive constraints, whereby randomness emerges from noisy evidence accumulation shaped by cognitive, motor, energetic, belief-based, and contextual factors (Ratcliff, 1978; Camerer, 1998; Krajbich et al., 2010). Third, individuals might deliberately choose to randomize, either in the intent to diversify risk or avoid regretful consequences over imperfect preferences (Machina, 1985; Cerreia-Vioglio et al., 2015; Cubitt et al., 2015).

☆ This article is part of a Special issue entitled: 'Mindful Economics' published in Journal of Economic Behavior and Organization.

* Corresponding author.
E-mail addresses: mattia.adamo@imtlucca.it, mattadam92@icloud.com (M. Adamo), coricell@usc.edu (G. Coricelli).

https://doi.org/10.1016/j.jebo.2026.107632
Received 24 July 2025; Received in revised form 12 May 2026; Accepted 22 May 2026
Available online 18 June 2026
0167-2681/© 2026 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

Empirical investigations of SC under risk 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). From a process-level perspective, this setting injects potential noise due to the distance between repetitions and the fact that the same repetitions could be not noticed by the decision-maker. More recent research has analyzed SC in economic settings involving sequential repeats (Agranov and Ortoleva, 2017; Cerreia-Vioglio et al., 2019), minimizing the sources of noise due to time distance and memory effects. In both scenarios, SC was observed frequently.

Given that SC is observed under distant and sequential conditions, do competing mechanisms of SC leave distinguishable traces in the decision process that process-tracing data can detect? Building on the experimental design of Agranov and Ortoleva (2017) (A&O), we answer this question by introducing two novelties: (1) the use of mouse-tracking as a continuous measure of the decision process unfolding over time between choices, and (2) the explicit comparison of response times (RT) and the Area Under the Curve (AUC) as complementary metrics for evaluating the decision process. The rationale for this integration lies in the different information conveyed by the two metrics: RT, while capturing the overall duration of the decision process, reduces within-trial dynamics to a single summary measure, thus making it relatively insensitive to the trajectory of deliberation. By contrast, AUC traces the continuous evolution of motor output throughout the trial, capturing the extent to which competing response options simultaneously attract the decision process. In a setting with binary choice sets and symmetrically displayed lotteries such as the one involving SC studies, AUC becomes particularly suited to track distinct decision paths. Combined with choice data and RTs, AUC provides an information set rich enough to distinguish between deliberate and noise-based accounts of SC. We hence develop our hypotheses by using Drift-Diffusion Models (DDMs, Ratcliff (1978) and Ratcliff and McKoon (2008)) and the Cautious Stochastic Choice (CSC, Agranov and Ortoleva (2017) and Cerreia-Vioglio et al. (2019)) model as benchmarks for bounded rationality and deliberate randomization, respectively.

In noise-based models, such as 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, starting point or decision threshold (Ratcliff, 1978; Ratcliff and McKoon, 2008). Each choice reflects an independent draw of diffusion noise, with no systematic dependence on prior selections, irrespective of whether repetitions are distant or sequential. As a result, the source of variability in the mouse-tracking data in the distant repeats should stem mainly because of trial-level noise, task difficulty, or bounded-rationality mechanisms that are not necessarily specific to SC. The same pattern should be mirrored in sequential repetitions: conditional on the same choice problem, AUC and RT should not systematically depend on whether the same option was chosen or rejected in the previous repetition. Any systematic difference between distant and sequential repetitions would therefore suggest that subjects treat the repetition structure as part of the decision environment, violating the independence assumption of a purely noise-based account.

By contrast, deliberative SC models generate variability through a history-dependent randomization process (Machina, 1985): choice probabilities are adjusted over time so that deviations are gradually offset in later trials, implementing an approximate quota-balancing mechanism (Cerreia-Vioglio et al., 2015). This generates a different set of predictions for distant and sequential repetitions. In distant repetitions, history-dependence may be weaker or less salient, but stochasticity should still not primarily arise from within-trial noise. In sequential repetitions, instead, the repetition structure should make the balancing logic more explicit, leading to stronger dependence on previous choices. However, this need not imply more erratic or more conflictual mouse trajectories. If SC reflects deliberate randomization, then it may display relatively stable within-trial dynamics, because the source of variability lies in the across-trial allocation rule rather than in noisy evidence accumulation within each trial. This is where AUC becomes especially informative. RT captures how long the decision takes, but it cannot distinguish whether a long decision reflects genuine choice conflict, slow motor initiation, distraction, or deliberate caution. AUC, by contrast, captures how motor commitment unfolds during the trial. It can therefore reveal whether SC is accompanied by a stronger or more sustained attraction toward the competing option, as expected under conflict-based or noisy bounded-rationality accounts, or whether trajectories remain relatively stable despite switching, as expected under a more deliberate randomization account. Therefore, we hypothesize AUC being more informative than RT in predicting SC.

Our findings first replicate the main behavioral regularities documented in the stochastic choice literature. SC is substantial in both repetition structures, although more frequent under distant repeats, and is especially concentrated in HARD questions. Further, RTs follow the expected pattern: they are longer for HARD questions, peak on the first presentation, and decline thereafter in both repetition structures.

However, RTs are not diagnostic of trial-level SC once question difficulty is controlled for. AUC, by contrast, provides a stronger process-level signal. In distant repeats, AUC remains stable across repetitions, especially for HARD questions, while dropping sharply after the first choice peak of sequential repeats, consistent with the idea that the explicit sequence makes the repetition structure salient and allows participants to plan the allocation of choices at the first encounter, before executing it more directly in later trials.

AUC also reveals heterogeneity that aggregate choice frequencies and RTs cannot separate. Among regular participants, SC does not increase AUC in distant repeats, consistent with across-trial balancing. In sequential repeats, however, SC is associated with larger AUC and more curved trajectories, suggesting a signature of cautiousness when randomization is explicit and planned. By contrast, participants who violate FOSD in distant repeats display larger and more erratic trajectories during SC trials, pointing to a more error-driven source of variability. Overall, these results indicate that SC is not generated by a single mechanism: for most subjects, the evidence is consistent with deliberate randomization, while a smaller subgroup displays a noisier, boundedly rational pattern.

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 DDMs. Random Utility models interpret SC as arising from unobservable heterogeneity in preferences or from utility shocks (Thurstone, 1927; Block and Marschak, 1959; McFadden and Richter, 1990; Gul and Pesendorfer, 2006; Gul et al., 2014). DDMs offer a second major noise-based account, in which choices arise from stochastic evidence accumulation until a decision threshold is reached (Ratcliff, 1978; Ratcliff and McKoon, 2008; Krajbich and Rangel, 2011). Crucially, in such models, variability across identical trials is generated by independent within-trial fluctuations, with no systematic dependence on prior selections.

In contrast, deliberate randomization models treat stochasticity as the outcome of purposeful planning. These frameworks, including CSC, formalize SC as the product of a structured deliberation stage preceding implementation (Machina, 1985; Cerreia-Vioglio et al., 2015; Brady and Rehbeck, 2016; Cerreia-Vioglio et al., 2019). Empirical contributions show that agents may randomize to avoid regret, diversify risk, or accommodate imprecise preferences, and that randomization plans exhibit temporal stability (Agranov and Ortoleva, 2017; Dwenger et al., 2012; Agranov et al., 2023). Recent experimental evidence further documents systematic preferences for mixtures, showing that mixing behavior is correlated across strategic and non-strategic environments and that a preference for mixtures predicts the distribution of stochastic choices across repetitions (Agranov et al., 2023; Feldman and Rehbeck, 2022).

2.2. Process-tracing in economics

Process-tracing methods complement theoretical accounts of SC by directly capturing the real-time temporal and dynamic features of decision-making through RT, gaze patterns, and mouse trajectories (Coricelli et al., 2020). When combined with choice data, these methods provide empirical leverage to disentangle whether stochasticity arises from genuine deliberation or from bounded rationality.

Response Times: RTs are among the earliest process measures employed in risky choice. Longer RTs tend to arise when options are closer in subjective value, while very short RTs are more common when one option is clearly dominant (Alós-Ferrer et al., 2016; Konovalov and Krajbich, 2016; Agranov and Ortoleva, 2017). However, inference linking fast responses to intuition and slow responses to deliberation can be misleading, as non-linearities complicate the speed–accuracy trade-off (Krajbich et al., 2015). This renders RTs alone insufficient to disentangle the dynamics of choice, motivating the use of complementary process measures.

Mouse Trajectories: Continuous mouse-tracking records the dynamic unfolding of decision conflict in real time, with more curved trajectories indicating stronger competition between response options (Freeman and Ambady, 2010). Across a range of domains—from intertemporal choice (Dshemuchadse et al., 2013) to dietary decisions (Sullivan et al., 2015) and learning (Konovalov and Krajbich, 2020)—greater trajectory curvature consistently reflects heightened cognitive conflict. Critically for our purposes, mouse-tracking evidence shows that AUC serves as a sensitive predictor of individual risk preferences, outperforming RT even at the single-trial level (Stillman et al., 2018). This superiority likely arises because AUC captures the continuous evolution of motor competition throughout a trial, rather than reducing deliberation to a single response duration. A detailed review of the eye-tracking and mouse-tracking literatures is provided in Appendix A1.

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.[1] 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 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

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

Table 1
List of questions asked.

Question	Lottery 1				Lottery 2				EV	SD
	(.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.

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),[2] three EASY[3] and four HARD questions.[4] The 10 question types are displayed in Table 1.

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 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.[5]

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.

[2] 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.

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

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

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

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 = 0.38) 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 1280 × 1024 pixels (75-Hz refresh frequency). At the beginning of each trial, participants clicked a "Start" 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 −0.8 and 0.8, 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 ms 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. Fig. 1 provides an illustration of the AUC in the context of our experimental setting

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.

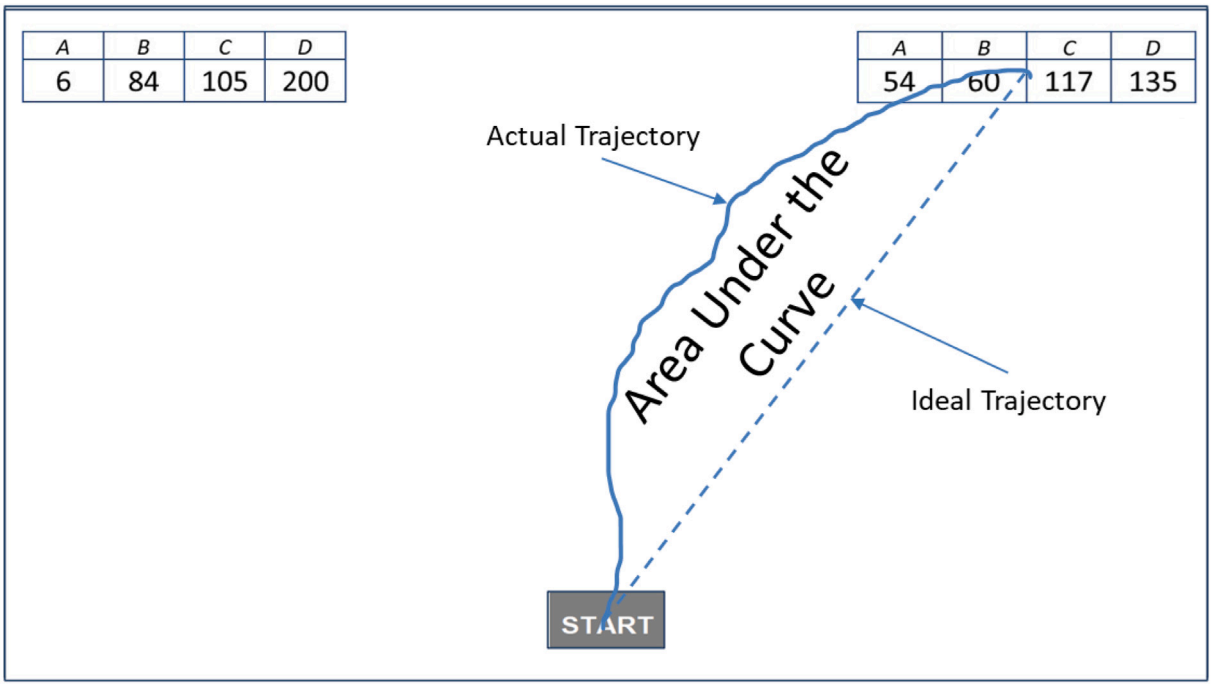

Fig. 1. Illustration of the area under the curve. The choice represents an actual trial from the experiment.

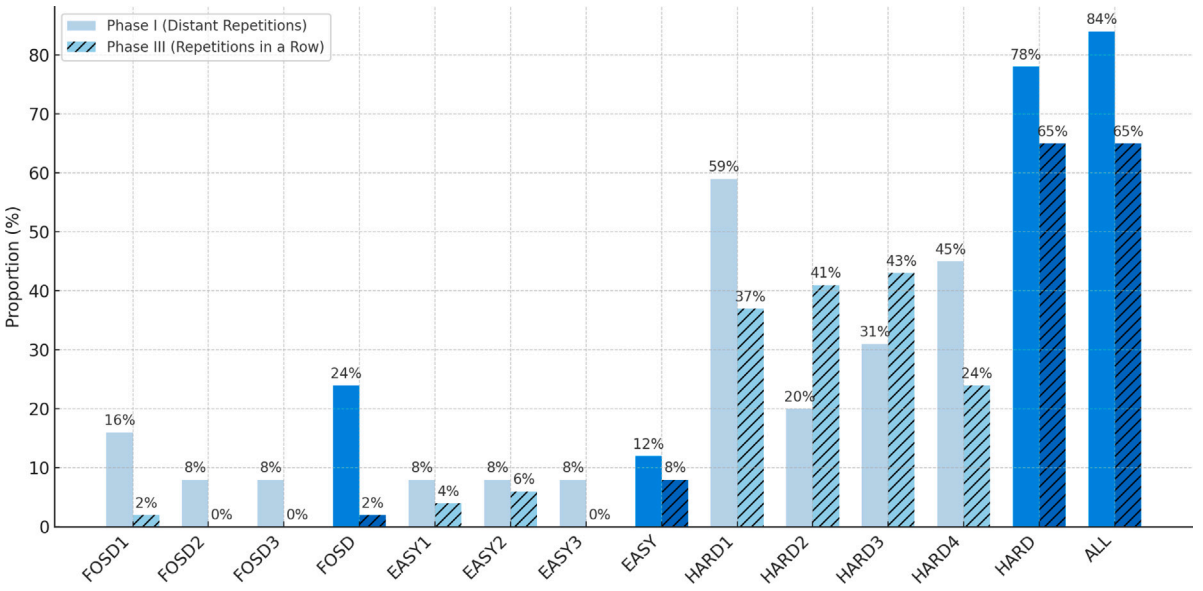

Fig. 2. Fraction of subjects displaying stochastic choice at least one time, by phase.

4. Results

4.1. Stochastic behavior between phases

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

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[6]). 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.[7] 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 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.[8]

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).

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.[9] 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 complexity beyond simple differences in expected value, transformed under risk-averse subjective utility, may be driving SC.[10]

[6] 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.

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

[8] 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.

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

[10] 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.

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***	0.05		
	(0.06)	(0.08)		
Flip coin			0.14***	0.11*
			(0.06)	(0.06)
Constant	0.22***	0.37***	0.19***	0.20***
	(0.03)	(0.006)	(0.02)	(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.

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 analyzing the evolution of mean RT and AUC across choice repetitions.[11] Fig. 3 and Table IV, Appendix A2, illustrate the differences in RT across phases and repetitions.[12]

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 (Fig. 4, Table V and Table Vb, Appendix A2 for Statistical Analysis).

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 ($p \gg 1$ RE GLS Table Vb, Appendix A2).

[11] 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.

[12] 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.

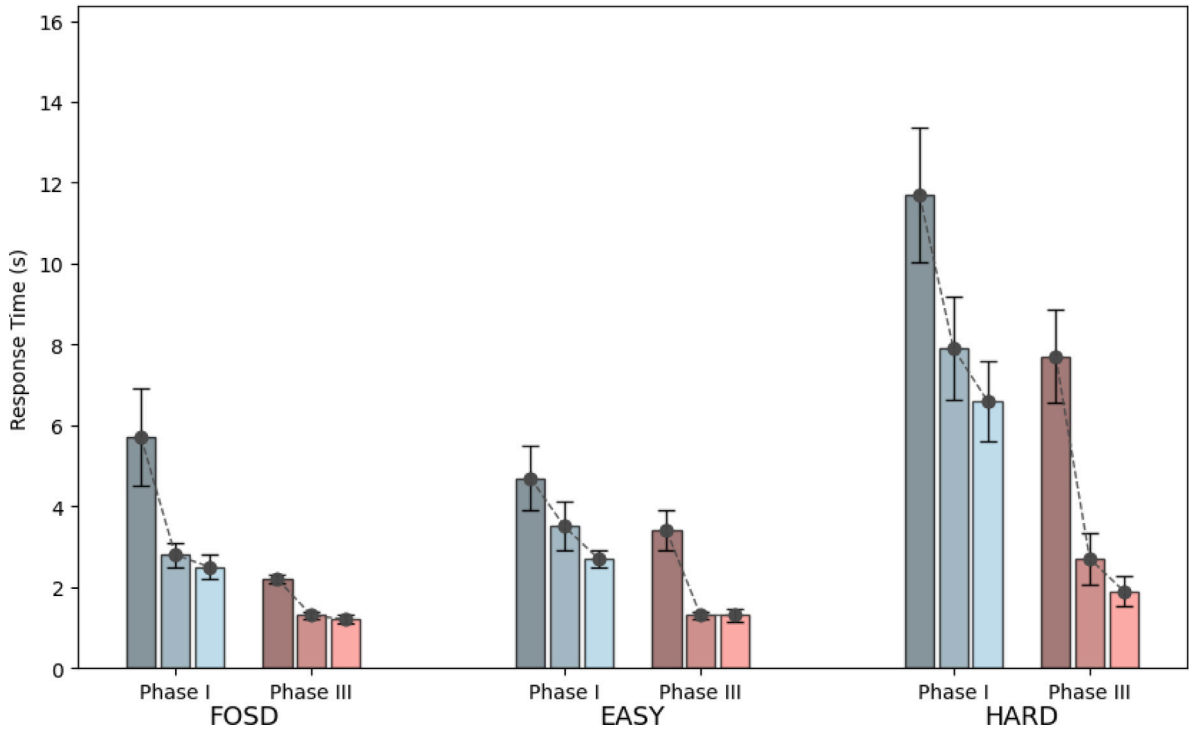

Fig. 3. Evolution of response time across repetitions, by phase.

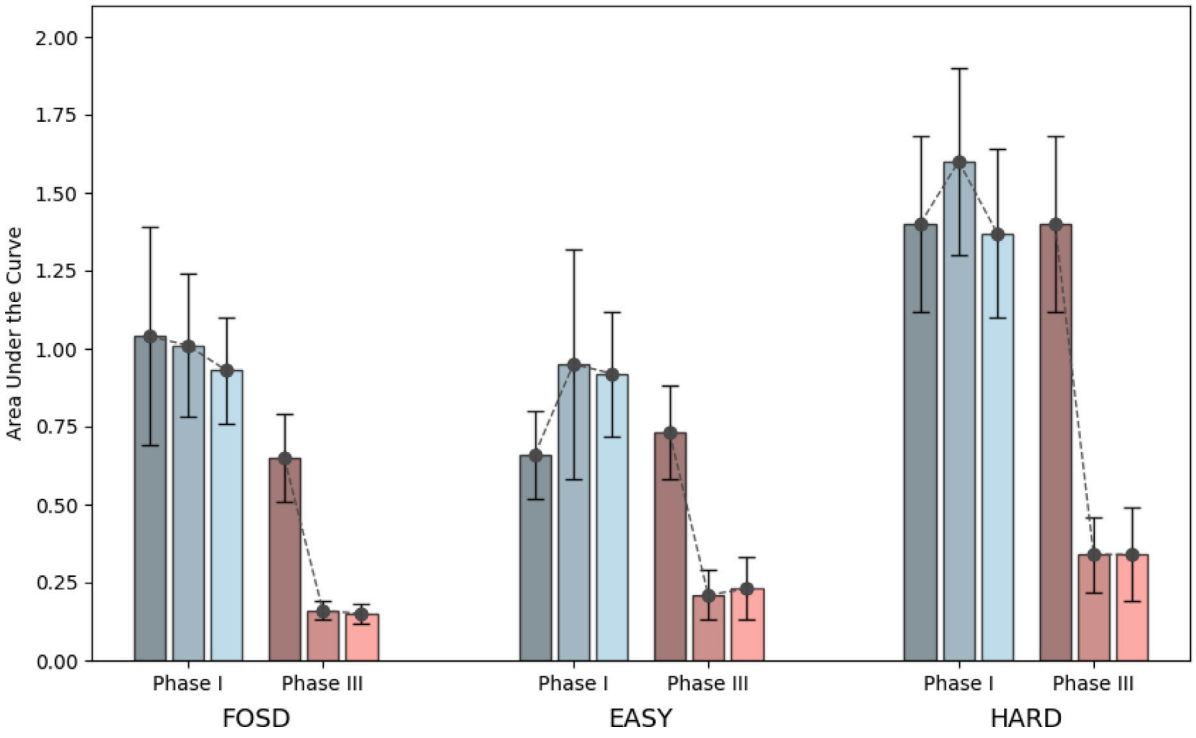

Fig. 4. Evolution of area under the curve across repetitions, by phase.

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.

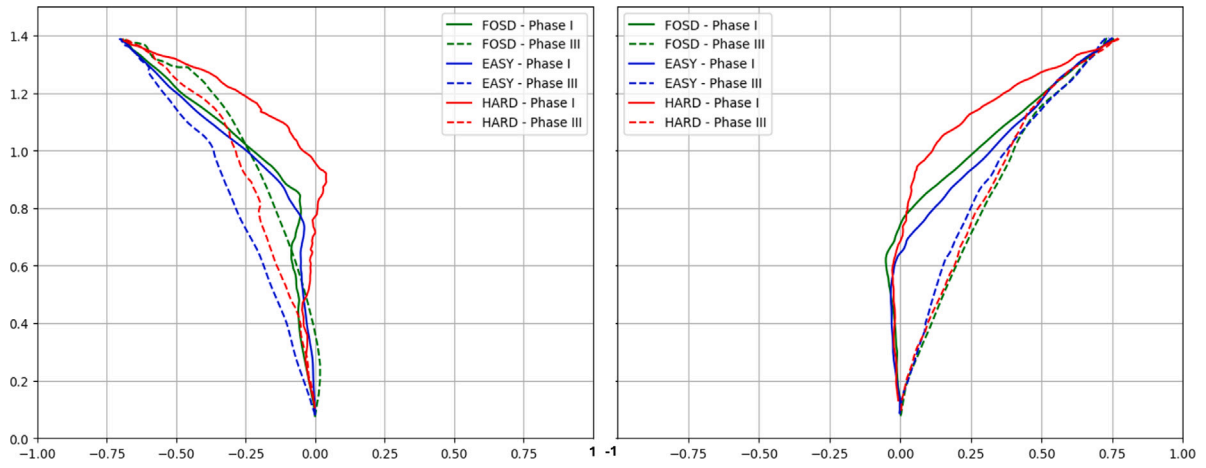

Fig. 5. Average mouse trajectories of all trials, by question type (FOSD, EASY, HARD). The continuous lines represent Phase I trials; the dashed lines represent Phase III trials.

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 Fig. 5, which compares the average trajectories of Phase I with those of Phase III for the same question types.[13]

We further examine differences in AUC trajectories between the first choice and its repetitions across the two phases of the experiment. Fig. 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.

We thus observe a stable pattern of AUC signal across repetitions in Phase I, particularly for HARD questions where this stability is most pronounced (Fig. 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.[14]

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, 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[15]; 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 the two subject groups are illustrated in Figs. 7 and 8. These figures compare trials in which subjects made SC with those in which they responded consistently, separately

[13] Figs. 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.

[14] Refer to Section 3.1.2 for the single-trial specification of SC.

[15] 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).

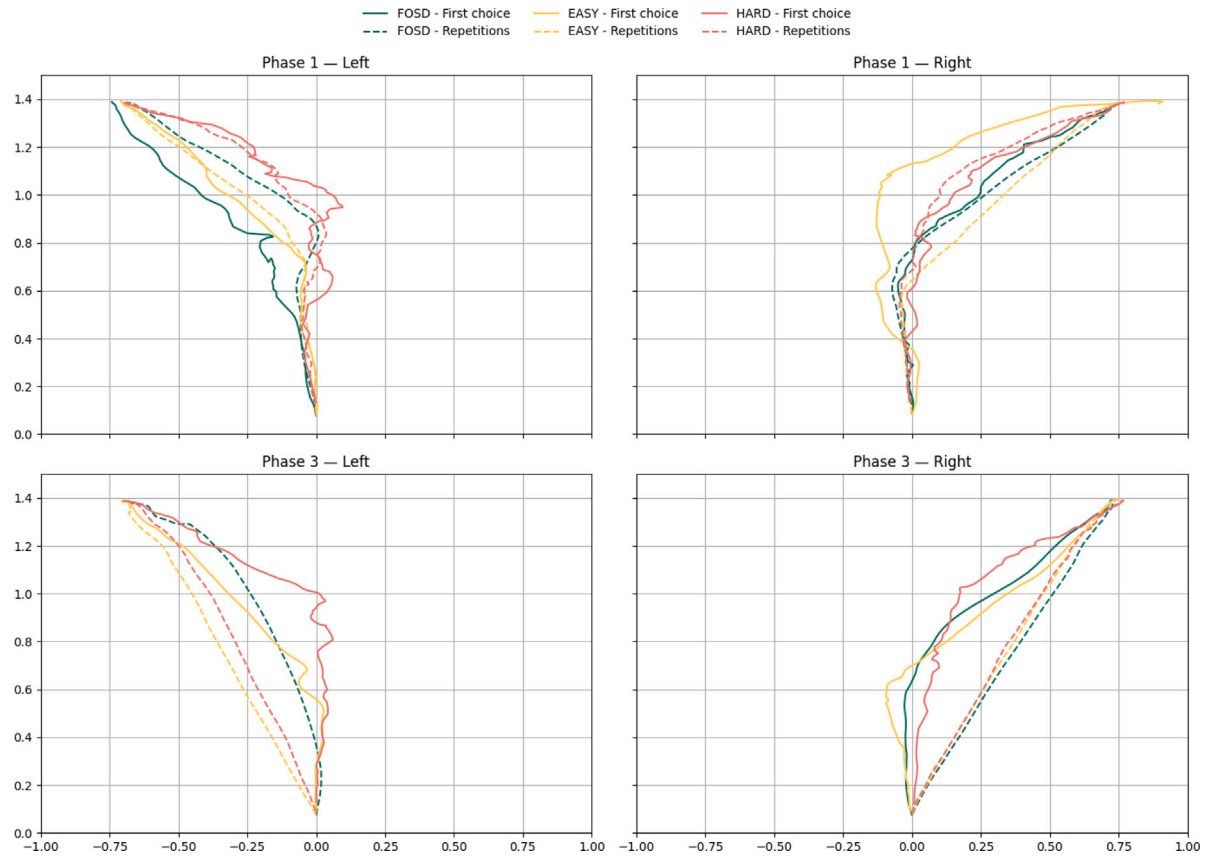

Fig. 6. Average mouse trajectories, by question type (FOSD, EASY, HARD). The continuous lines represent the 1st trials, or reference choice; the dashed lines represent repeats of the same question type.

Table 3
Random-effects logit for SC, Phase I and Phase III.

SC	All		HARD	
	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.

for phase I (Fig. 7) and phase III (Fig. 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.

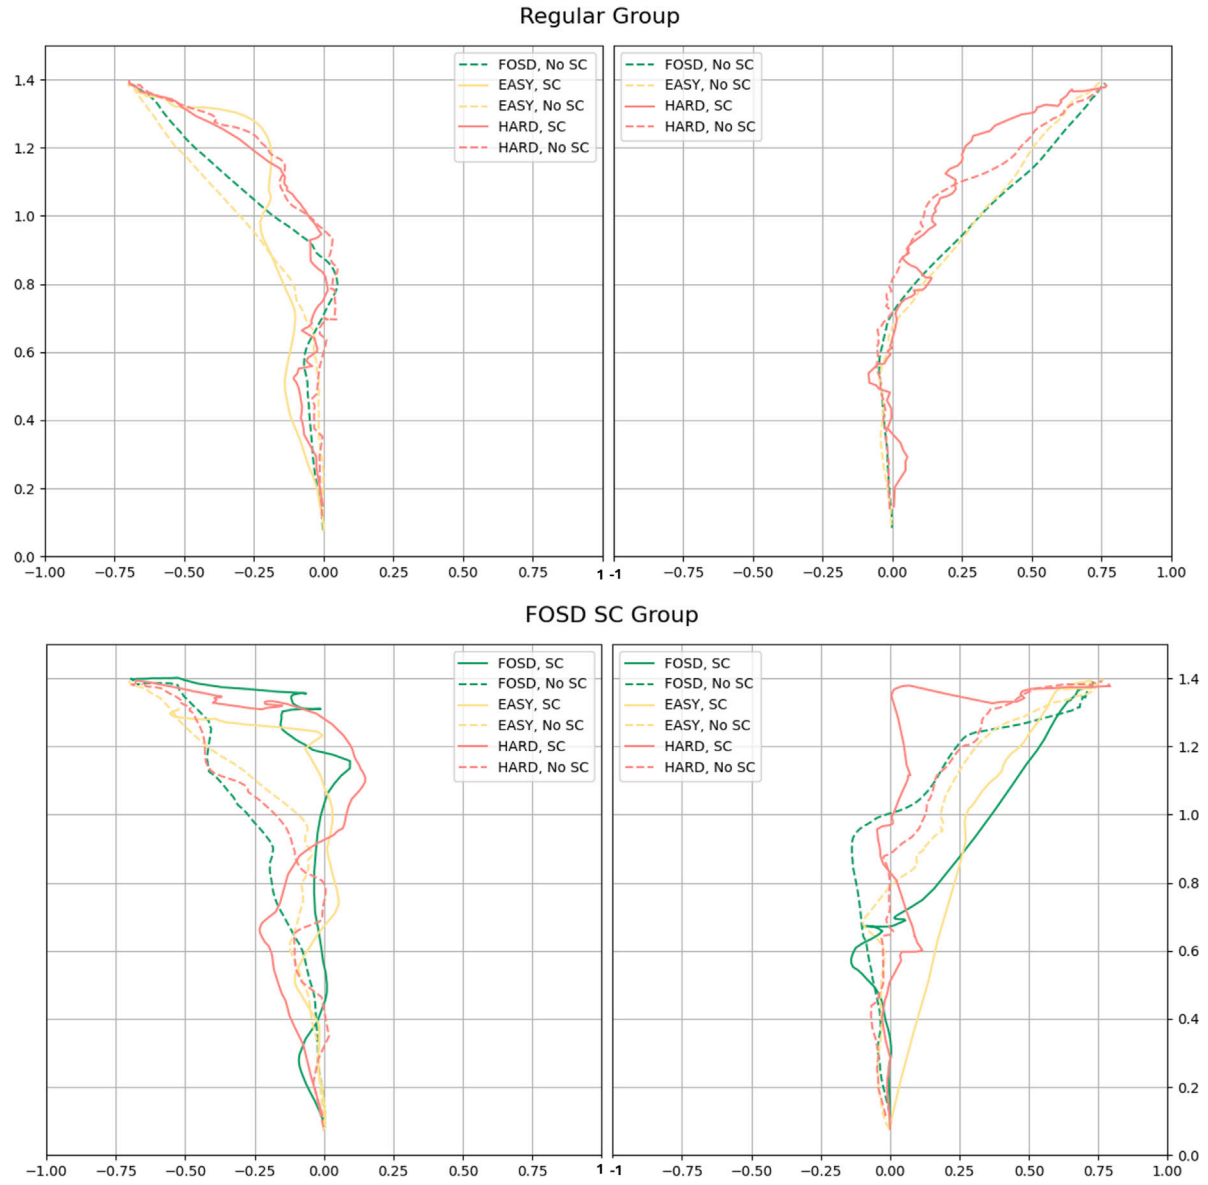

Fig. 7. Average mouse trajectories of Phase I, by question type (FOSD, EASY, HARD). The continuous lines represent SC trials; the dashed lines display no-SC trials. Regular subjects are displayed above; subjects displaying SC in FOSD question types are displayed below.

Table 4
Mean (SD) frequency of SC by group.

Mean SC occurrences	Phase I	Phase III
FOSD SC group	3.59 (0.08)	1.08 (0.05)
Regular group	2.27 (0.033)	2.32 (0.04)

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.

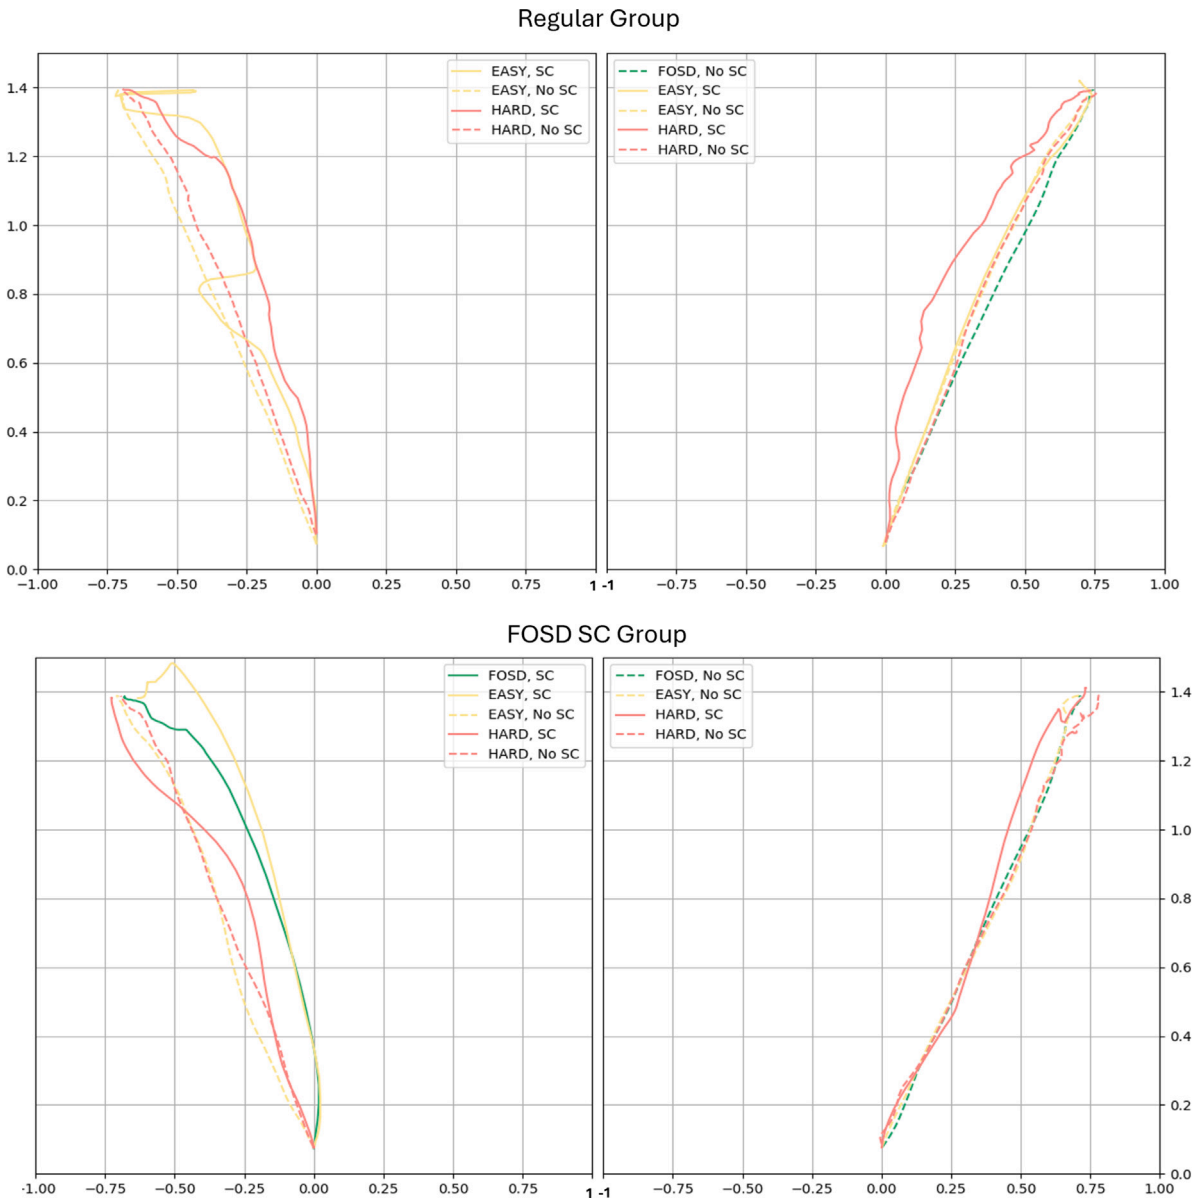

Fig. 8. Average mouse trajectories of Phase III, by question type (FOSD, EASY, HARD). The continuous lines represent SC trials; the dashed lines display no-SC trials. Regular subjects are displayed above; subjects displaying SC in FOSD question types are displayed below.

Table 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 SC under risk, we used choice-process data to distinguish between boundedly rational and deliberate mechanisms of randomization. We adopted an experimental design that compared distant and sequential repetitions of pairs of lotteries varying in complexity, adding mouse-tracking as an informative measure of the different decision processes underlying SC.

At the behavioral level, our results reproduce the main regularities of the SC literature. SC is observed in both repetition structures, is more frequent under distant repetitions, and is concentrated especially in HARD questions. RTs also follow the expected pattern: they are longer for HARD items, peak on the first presentation, and decline after repeated exposure in both phases. These results establish the behavioral baseline of the experiment. However, they also show the limits of choice frequencies and RTs as diagnostic measures: once question difficulty is taken into account, RTs do not reliably distinguish stochastic from consistent trials.

At the process level, RT and AUC revealed different aspects of the repetition structure. RT declined after the first presentation in both phases, suggesting that repeated exposure reduces decision time independently of whether repetitions are distant or sequential. AUC displayed a sharper distinction. In distant repeats, trajectories remained stable across repetitions, especially for HARD questions. In sequential repeats, instead, AUC peaked on the first choice and dropped sharply thereafter, suggesting that explicit sequential repetitions allow participants to process and possibly plan the sequence at first encounter, executing subsequent choices more directly.

On a trial-by-trial basis, AUC proved more informative than RT. Once question difficulty was taken into account, RT did not reliably predict whether a participant switched away from the initial choice. AUC, by contrast, significantly predicted SC in both repetition structures, with a stronger effect in Phase III. The effects on the two phases, however, are driven by two different groups of participants that display both behavioral and process-level differences. Among regular participants, AUC predicts SC in Phase III, but not in Phase I.

This suggests that, when repetitions are explicit, SC becomes salient in the decision process for regular subjects, and the larger curvature observed in SC trials may represent a signature of cautiousness under deliberate randomization.

The opposite pattern emerges for participants who violated FOSD in Phase I. For this group, SC in the distant-repeats phase is associated with larger and more erratic trajectories. In Phase III, however, their trajectories become much more linear and no longer distinguish SC trials from consistent trials. This process-level shift is mirrored by their choice frequencies: FOSD violators randomize frequently in the implicit distant-repeats phase, but their stochastic behavior drops sharply when repetitions become explicit and sequential.

Taken together, these results suggest that SC is not generated by a single mechanism. For regular participants, the process evidence is consistent with a predominantly deliberate form of randomization. Their AUC does not distinguish SC from consistent trials in distant repetitions, but becomes informative when repetitions are explicit and sequential, suggesting that the structure of Phase III makes SC more salient in the decision process. At the same time, their average SC frequencies remain stable across repetition structures, indicating behavioral consistency despite different process-level dynamics. For FOSD violators, instead, the dissociation between high SC with erratic trajectories in Phase I and low SC with direct trajectories in Phase III points to a boundedly rational component of SC, likely amplified by the implicit and memory-demanding structure of distant repetitions. Mouse-tracking therefore helped separate these components by showing that similar choice variability can reflect different decision processes.

5.1. Limitations of the experimental design

While the dual repetitions-structure design we drew from A&O 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.

Future studies aimed at disentangling motives of SC might consider a richer decision space, including a higher number of repetitions and a more comprehensive and systematically varied set, including broader probability ranges, presence/absence of zero outcomes, and more heterogeneous payoff structures.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was supported by the University of Southern California Internal Research Fund to GC.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jebo.2026.107632.

Data availability

I have shared the data and analysis code in OSF repository, link attached.

Replication Material for Deliberate Stochastic Choice: A Mouse-Tracking Experiment (Original data) (OSF)

References

Agranov, M., Healy, P.J., Nielsen, K., 2023. Stable randomisation. Econ. J. 133 (655), 2553–2579.
Agranov, M., Ortoleva, P., 2017. Stochastic choice and preferences for randomization. J. Political Econ. 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. Econ.: J. Econ. Soc. 503–546.
Allen, R., Rehbeck, J., 2023. Revealed stochastic choice with attributes. Econom. Theory 75 (1), 91–112.
Alós-Ferrer, C., Garagnani, M., Hügelschäfer, S., 2016. Cognitive reflection, decision biases, and response times. Front. Psychol. 7, 1402.
Ballinger, T.P., Wilcox, N.T., 1997. Decisions, error and heterogeneity. Econ. J. 107 (443), 1090–1105.
Becker, G.M., DeGroot, M.H., Marschak, J., 1964. Measuring utility by a single-response sequential method. Behav. Sci. 9 (3), 226–232.
Block, H.D., Marschak, J., 1959. Random orderings and stochastic theories of response.
Brady, R.L., Rehbeck, J., 2016. Menu-dependent stochastic feasibility. Econometrica 84 (3), 1203–1223.
Camerer, C.F., 1989. An experimental test of several generalized utility theories. J. Risk Uncertain. 2, 61–104.
Camerer, C., 1998. Bounded rationality in individual decision making. Exp. Econ. 1, 163–183.
Caplin, A., Dean, M., Leahy, J., 2019. Rational inattention, optimal consideration sets, and stochastic choice. Rev. Econ. Stud. 86 (3), 1061–1094.
Cerreia-Vioglio, S., Dillenberger, D., Ortoleva, P., 2015. Cautious expected utility and the certainty effect. Econometrica 83 (2), 693–728.
Cerreia-Vioglio, S., Dillenberger, D., Ortoleva, P., Riella, G., 2019. Deliberately stochastic. Am. Econ. Rev. 109 (7), 2425–2445.
Coricelli, G., Polonio, L., Vostroknutov, A., 2020. The process of choice in games. In: Handbook of Experimental Game Theory. Edward Elgar Publishing, pp. 69–94.
Cubitt, R.P., Navarro-Martinez, D., Starmer, C., 2015. On preference imprecision. J. Risk Uncertain. 50, 1–34.
Dshemuchadse, M., Scherbaum, S., Goschke, T., 2013. How decisions emerge: action dynamics in intertemporal decision making. J. Exp. Psychol. [Gen.] 142 (1), 93.
Dwenger, N., Kübler, D., Weizsäcker, G., 2012. Flipping a coin: Theory and evidence.
Feldman, P., Rehbeck, J., 2022. Revealing a preference for mixtures: An experimental study of risk. Quant. Econ. 13 (2), 761–786.
Freeman, J.B., Ambady, N., 2010. MouseTracker: Software for studying real-time mental processing using a computer mouse-tracking method. Behav. Res. Methods 42 (1), 226–241.
Fudenberg, D., Strzalecki, T., 2015. Dynamic logit with choice aversion. Econometrica 83 (2), 651–691.
Gneezy, U., Potters, J., 1997. An experiment on risk taking and evaluation periods. Q. J. Econ. 112 (2), 631–645.
Gul, F., Natenzon, P., Pesendorfer, W., 2014. Random choice as behavioral optimization. Econometrica 82 (5), 1873–1912.
Gul, F., Pesendorfer, W., 2006. Random expected utility. Econometrica 74 (1), 121–146.
Hey, J.D., Orme, C., 1994. Investigating generalizations of expected utility theory using experimental data. Econ.: J. Econ. Soc. 1291–1326.
Konovalov, A., Krajbich, I., 2016. Gaze data reveal distinct choice processes underlying model-based and model-free reinforcement learning. Nat. Commun. 7 (1), 12438.
Konovalov, A., Krajbich, I., 2020. Mouse tracking reveals structure knowledge in the absence of model-based choice. Nat. Commun. 11 (1), 1893.
Krajbich, I., Armel, C., Rangel, A., 2010. Visual fixations and the computation and comparison of value in simple choice. Nature Neurosci. 13 (10), 1292–1298.
Krajbich, I., Bartling, B., Hare, T., Fehr, E., 2015. Rethinking fast and slow based on a critique of reaction-time reverse inference. Nat. Commun. 6 (1), 7455.
Krajbich, I., Rangel, A., 2011. Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. Proc. Natl. Acad. Sci. 108 (33), 13852–13857.
Luce, R.D., 1959. Individual Choice Behavior, vol. 4, Wiley New York.
Machina, M.J., 1985. Stochastic choice functions generated from deterministic preferences over lotteries. Econ. J. 95 (379), 575–594.
McFadden, D., Richter, M.K., 1990. Stochastic rationality and revealed stochastic preference. In: Preferences, Uncertainty, and Optimality, Essays in Honor of Leo Hurwicz. Westview Press, Boulder, CO, pp. 161–186.
Mosteller, F., Nogee, P., 1951. An experimental measurement of utility. J. Political Econ. 59 (5), 371–404.
Ratcliff, R., 1978. A theory of memory retrieval. Psychol Rev 85 (2), 59.
Ratcliff, R., McKoon, G., 2008. The diffusion decision model: theory and data for two-choice decision tasks. Neural Comput. 20 (4), 873–922.
Regenwetter, M., Dana, J., Davis-Stober, C.P., 2011. Transitivity of preferences. Psychol Rev 118 (1), 42.
Samuelson, P.A., 1938. A note on the pure theory of consumer's behaviour. In: The Foundations of Price Theory Vol 4. Routledge, pp. 101–116.
Stillman, P.E., Shen, X., Ferguson, M.J., 2018. How mouse-tracking can advance social cognitive theory. Trends Cogn. Sci. 22 (6), 531–543.
Sullivan, N., Hutcherson, C., Harris, A., Rangel, A., 2015. Dietary self-control is related to the speed with which attributes of healthfulness and tastiness are processed. Psychol. Sci. 26 (2), 122–134.
Thurstone, L.L., 1927. Psychophysical analysis. Am. J. Psychol. 38 (3), 368–389.
Tversky, A., 1969. Intransitivity of preferences. Psychol Rev 76 (1), 31.
Von Neumann, J., 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. Annu. Rev. Econ. 12 (1), 579–601.