Beyond Noise: Salience and Strategy in Stochastic Choice

Riccardo Manghi[*] Daniela Di Cagno[†] Giorgio Coricelli[‡]

December 16, 2025

Abstract

We study stochastic choice in experimental lotteries to distinguish cognitive noise from deliberate randomization. Building on Agranov and Ortoleva (2017), we use a design with repeated lottery pairs that allows us to measure randomization both within and across repetitions, and to identify its structural and individual determinants. We replicate the original evidence of persistent stochasticity, but show that it cannot be explained by task difficulty or cognitive load. Using lottery–level attributes, we find that randomization is systematically related to conflicts between salient attributes—such as the probabilities of extreme outcomes—rather than to expected–value differences or complexity. A high–stake orientation strongly predicts stochastic mixing, indicating a strategic rather than noisy component. We propose a salience–based interpretation of deliberate randomization in which stochastic choice reflects attention shifts across competing attributes when dominance is ambiguous.

[*]LUISS Guido Carli University, CESARE Laboratory. Email: riccardo.manghi@luiss.it.
[†]LUISS Guido Carli University, Department of Economics and Finance.
[‡]University of Southern California and University of Trento.

1 Introduction

In most standard economic models, individuals are assumed to make stable and deterministic choices. Yet, empirical evidence consistently shows that even in simple and repeated decision problems, subjects often switch between alternatives and do not display stable and consistent preferences. This behavioral variability -long treated as random error or imperfect optimization- has increasingly been recognized as potentially systematic and meaningful. Understanding the origins of such stochasticity is crucial for models of choice, welfare analysis, and mechanism design.

While classical frameworks such as expected utility theory interpret stochastic choice as noise (Hey and Orme (1994), Hey (1997), Carbone and Hey (2000)), alternative perspectives suggest that randomness itself may be a *deliberate component* of behavior. In certain environments, randomization can serve an instrumental purpose: students facing multiple-choice exams may deliberately mix among plausible options (that's why penalty points exist); investors may occasionally "test" riskier assets to explore payoffs; and individuals in repeated strategic interactions may intentionally randomize to avoid being predictable. These examples illustrate that stochastic behavior can arise not only from cognitive limitations, but also from purposeful strategy, exploration, or emotional regulation.

A key contribution by Agranov and Ortoleva (2017) provided experimental evidence that individuals sometimes deliberately randomize repeated lottery choices. Their findings challenged the view that stochastic choice merely reflects confusion or attentional noise. However, whether this phenomenon extends beyond their specific settings remains an open question. Does deliberate randomization systematically emerge in risky environments? What structural and individual factors drive such behavior? What strategy or motive/s does it serve?

We address these questions by replicating and extending the experimental logic of Agranov and Ortoleva (2017) in a controlled risky-choice setting. Our design generalizes their framework to a non-strategic environment using a new and independently validated set of lotteries originally developed by Hey and Di Cagno (1994), from which we derived a subset of ten lotteries classified according to complexity, where this classification was based on heterogeneity of choices across participants and response times, following the same procedure as in Agranov and Ortoleva (2017). In doing so, we reinforce and broaden the original evidence, distinguish between deliberate and non-deliberate components of stochasticity, and investigate the structural and individual determinants of randomization. Moreover, we test whether randomization is driven by *saliency of attributes* -that is, by the attention-grabbing nature of extreme outcomes such as the probability of zero or the highest stake- thus providing a behavioral foundation for models of local thinking and attribute salience.

Our study is related to several strands of literature. First, the study of stochastic choice itself, from random utility theory (McFadden, 1974; Loomes and Sugden, 1995) to deliberate randomization (Agranov and Ortoleva, 2017). A large part of this literature has traditionally interpreted stochasticity as cognitive noise or error (Hey, 1995; Hey and Orme, 1994; Carbone and Hey, 2000), whereas our approach emphasizes strategic mixing driven by attribute salience.

Second, the literature on cognitive foundations of behavior, including limited attention and local thinking (Gennaioli and Shleifer, 2010; Bordalo et al., 2012), and the effects of cognitive load and dual-process reasoning on economic decision-making (Kahneman, 2011; Rubinstein, 2016). Third, the literature on strategic and motivational randomization, where mixing serves as a commitment or hedging device (Aumann, 1964; Camerer, 1998).

These frameworks offer complementary but distinct interpretations. Under the cognitive view, stochasticity increases with task difficulty or attentional constraints. Under the strategic view, randomization reflects a deliberate choice to diversify, hedge regret, or sustain engagement. A salience-based perspective bridges the two: individuals may overweight extreme payoffs, alternating focus between safe and high-stake components of lotteries. This can generate systematic, yet seemingly random, switching -a pattern consistent with attribute-based local thinking.

Our objective is to replicate and strengthen the evidence of deliberate randomization using a new and independently validated set of lotteries,isolating purposeful mixing from unintentional noise, and to examine what drives such behavior, combining structural and individual-level evidence. Randomization proves unrelated to task difficulty or cognitive load, but strongly linked to the lotteries' structural primitives—particularly the probabilities of zero and of the highest payoff. These patterns support a novel interpretation: randomization as a strategic, salience-driven exploration process, where individuals occasionally "test" options weighted toward salient outcomes, bridging deliberate randomization with models of local thinking and attention-based choice.

The experimental design consists of two main phases that differ in repetition structure and information transparency. In Phase 1, each lottery pair appears four times in random order, without revealing repetitions to participants. In Phase 2, the same pairs are repeated three times consecutively, making repetition explicit. The difference in switching rates between phases decomposes observed stochasticity into non-deliberate (attention or memory noise) and deliberate (strategic) components.

Task difficulty was calibrated through a pilot experiment following the logic in Agranov and Ortoleva (2017), classifying this new set of lottery pairs as **Easy**, **Hard**, or **FOSD** based on heterogeneity and response times. The Easy–Hard gradient manipulates cognitive load: if randomization arises from limited cognition, it should rise with difficulty; if it persists even in easy tasks, it is more likely deliberate. The inclusion of FOSD pairs—where one lottery strictly dominates the other—acts as a behavioral control: stability in these choices ensures that switching elsewhere cannot be attributed to boredom or mechanical errors.

In addition, a post-experimental survey collected sociodemographics and attitudinal measures (clarity, comfort, strategy awareness, attention), allowing us to link metacognitive self-reports with observed randomization. Trial-level and participant-level data were analyzed using logistic regressions of randomization and choice probability on both behavioral and structural primitives, providing a comprehensive decomposition of stochastic behavior.

The results confirm and extend the evidence of deliberate randomization. Participants exhibit significant stochasticity even when repetition is transparent, implying intentional mixing rather than cognitive error. Randomization is not predicted by task difficulty but is strongly related to the structural primitives of the lotteries, particularly the probability of zero and of the high stake. These findings suggest that randomization is guided by salient attributes and exploratory motives, providing a behavioral foundation for a new, attribute-salience interpretation of deliberate randomization.

Our contribution is threefold: (i) we replicate and generalize deliberate randomization to risky choice; (ii) we distinguish its deliberate nature from cognitive noise; and (iii) we link stochastic behavior to the saliency of attributes, laying empirical groundwork for a model of strategic randomization driven by local thinking and attentional focus.

The remainder of the paper proceeds as follows. Section 2 describes the experimental design and sample. Section 3 presents the main results on randomization and determinants of choice. Section 4 discusses implications for models of salience and stochastic choice. Section 5 concludes.

2 Experimental Design

The experiment was designed to replicate and reinforce the evidence of *deliberate randomization* documented by Agranov and Ortoleva (2017), test its robustness in a non-strategic risky-choice environment using a new set of lotteries, and explore the structural and behavioral drivers of randomization. Specifically, we aim to distinguish whether randomization is genuinely *strategic*—and what motive it serves—or a by-product of cognitive limitations, and to identify which structural primitives of the lotteries are most predictive of such behavior.

2.1 Lotteries, Pilot, and Selection Procedure

We closely replicate the experimental structure and classification procedure of Agranov and Ortoleva (2017), while extending it to a different set of lotteries originally developed by Hey and Di Cagno (1994). The goal was to test whether the evidence of deliberate randomization generalizes beyond the original setting, while mantaining a comparable difficulty structure across lotteries.

The initial design considered sixty binary lottery pairs from Hey and Di Cagno (1994). From these, following Agranov and Ortoleva (2017)' s same procedure, we conducted a dedicated pilot experiment that recorded both between-subject heterogeneity and mean response times for each lottery pair, providing two complementary indicators of cognitive demand. Pairs with high response times and substantial heterogeneity in choices were classified as more complex, whereas pairs with uniform choices and shorter times were deemed easier.

Based on this pilot evidence, we selected seven lottery pairs for the main experiment, divided into two difficulty, EASY and HARD, to which we add three first order stochastically dominated lotteries (FOSD), for a total of ten lotteries divided in three complexity categories. This procedure preserves the conceptual framework Agranov and Ortoleva (2017) while introducing a new, independently validated set of lotteries. The whole initial lottery set and the results of the pilot are reported in the Appendix. Variation in task complexity is used to generate an exogenous gradient in cognitive load and to test whether randomization intensity depends on it. The ten lotteries for the main experiment are thus grouped as:

- **Easy (pairs 1–3):** low heterogeneity and short response times.
- **Hard (pairs 4–7):** high heterogeneity and slower decisions.
- **FOSD (pairs 8–10):** one lottery clearly dominates the other in all states. These pairs act as a behavioral control: rational agents should exhibit no randomization here.

2.2 Design and Phase Structure

Each subject faced the same ten lottery pairs for a total of seven times, across two distinct phases:

- **Phase 1 (Interleaved repetitions).** Each pair was presented four times, interleaved in random order with other pairs. Since repetitions were dispersed and unannounced, subjects may not recognize identical tasks, so randomization in this phase may include both deliberate and non-deliberate components (e.g., inattention, memory noise).
- **Phase 2 (Consecutive repetitions).** Each pair was shown three times consecutively. The repetition pattern was transparent, minimizing memory and attention costs. Randomization in this setting is thus more plausibly deliberate, as switching cannot be attributed to lack of awareness of repetition.

To prevent fatigue and sustain engagement, a short *dice investment task* was administered between Phase 1 and Phase 2. This intermission provided a neutral cognitive break while preserving incentive compatibility, ensuring that behavior in Phase 2 was not mechanically affected by repetition or boredom effects. This two-phase design allows us to look observed stochasticity from two different perspectives: within-phase variation and between-phase variation. A reduction from Phase 1 to Phase 2 randomization captures the extent of non-deliberate noise, whereas persistence of randomization in Phase 2 provides evidence of deliberate mixing.

In addition, the design incorporates two internal validity checks. The **FOSD pairs** serve as a behavioral control: since one lottery strictly dominates the other, no rational subject should ever switch or randomize across repetitions. Stability in these pairs ensures that residual stochasticity in non-dominated tasks cannot be attributed to boredom, inattentiveness, or mechanical key pressing. The contrast between **Easy** and **Hard** pairs provides a direct measure of the role of *cognitive load* in randomization. If randomization were purely driven by computational difficulty or inattention, it should rise sharply with task complexity; if instead it persists even in Easy pairs, this persistence points to the existence of a strategic, deliberate component of randomization.

2.3 Lottery Representation and Structural Primitives

All lotteries were expressed over at most three monetary outcomes and were geometrically represented within four canonical Marschak–Machina triangles: (20,10,0), (30,20,10), (30,20,0), and (30,10,0), with different probabilies across the lotteries. This representation facilitates both visual interpretation of the outcomes that shape choice. The MM representation of all the 10 lottery pairs is reported in the appendix.

For each lottery, we computed its *structural primitives*:

- Expected value (EV);
- Standard deviation of payoffs (SD);
- Probability of the lowest outcome (zero);
- Probability of the highest outcome (thirty).

These primitives are used as covariates in subsequent regressions to test which attributes systematically influence the likelihood of choosing a given lottery and the probability of randomization . The design thus try to link observed behavior to interpretable features of the choice environment.

2.4 Incentives, Interface, and Implementation

Choices were incentivized through a standard random-incentive mechanism: at the end of the experiment, one decision for each phase was randomly selected and implemented for real payment. Lotteries were displayed in explicit probability–outcome notation (e.g., "0.375×0, 0.125×20, 0.500×30") alongside a visual pie chart representation. Participants completed practice rounds before the main task to familiarize themselves with the interface and ensure understanding of probability displays. Each trial required an explicit choice; no default was permitted.

An example of how lottery choices were displayed can be observed here:

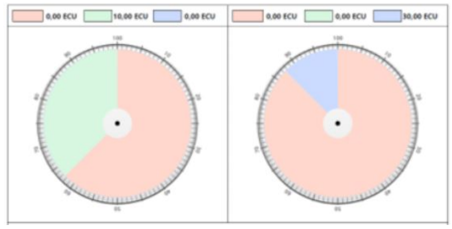

Same amounts across different lotteries have been represented with the same colour, to maximize clarity and avoid confusion.

While below all the 10 lotteries used for the main experiment are reported, with their graphical Marshak-Machina representation in the appendix:

Table 1: Structure of the 10 lottery pairs

Index	Lottery	Outcomes	Probabilities	EV	Category
1	L1	(0, 10)	(0.625, 0.375)	3.75	Easy
	L2	(0, 30)	(0.875, 0.125)	3.75	Easy
2	L1	(20, 30)	(0.75, 0.25)	22.5	Easy
	L2	(0, 20, 30)	(0.375, 0.125, 0.5)	17.5	Easy
3	L1	(0, 10, 20)	(0.375, 0.5, 0.125)	7.5	Easy
	L2	(0, 20)	(0.75, 0.25)	5.0	Easy
4	L1	(0, 10)	(0.625, 0.375)	3.75	Hard
	L2	(0, 20)	(0.875, 0.125)	2.5	Hard
5	L1	(10, 20, 30)	(0.375, 0.5, 0.125)	17.5	Hard
	L2	(10, 30)	(0.75, 0.25)	15.0	Hard
6	L1	(10, 30)	(0.75, 0.25)	15.0	Hard
	L2	(0, 10, 30)	(0.375, 0.125, 0.5)	16.25	Hard
7	L1	(20, 30)	(0.875, 0.125)	21.25	Hard
	L2	(10, 30)	(0.375, 0.625)	22.5	Hard
8	L1	(0, 30)	(0.9, 0.1)	3.0	FOSD
	L2	(0, 30)	(0.8, 0.2)	6.0	FOSD
9	L1	(10, 30)	(0.95, 0.05)	10.0	FOSD
	L2	(10, 30)	(0.7, 0.3)	16.0	FOSD
10	L1	(0, 20, 30)	(0.6, 0.3, 0.1)	9.0	FOSD
	L2	(0, 20, 30)	(0.4, 0.4, 0.2)	12.0	FOSD

Notes: Each lottery pair was shown in both experimental phases. Categories follow pilot-based difficulty classification. Expected values (EV) in experimental currency units.

2.5 Sample

The experiment was conducted at the CESARE Lab of LUISS Guido Carli University (Rome) using the `oTree` experimental platform. A total of 84 participants took part in the preliminary pilot and 109 in the main experiment. All participants were LUISS students recruited from the laboratory's subject pool and received monetary incentives based on their outcomes of the randomly selected lotteries..

The dataset combines both trial-level and participant-level information, including individual lottery choices, response times, and a comprehensive post-experimental survey. The survey collected standard sociodemographic information (age, gender, field of study, and region of origin) and a set of attitudinal measures related to the experimental task. Specifically, participants were asked to evaluate their perceived clarity and understanding of the lotteries, their comfort during decision making, whether they deliberately adopted a mixed or strategic approach, and the degree of attention they maintained throughout the session.

The following table reports summary statistics for the participants:

In summary, the experimental design replicates the methodological logic of Agranov and Ortoleva (2017) in a non-strategic risky-choice setting, using a pilot-validated subset of Hey and Di Cagno (1994) lotteries. It provides

Table 2: Summary statistics of participants

Variable	**Mean / Share**	**(SD)**
Age	23.8	(2.4)
Female	62.4%	
Faculty of Economics	68.8%	
Faculty of Political Science	21.1%	
Faculty of Law	10.1%	
Region: North	34.9%	
Region: Center	42.2%	
Region: South and Islands	22.9%	
Understood Instructions	89.9%	
Clear Presentation	91.7%	
Found Choices Easy	64.2%	
Felt Comfortable	78.0%	
Reported Strategy Use	72.5%	

Notes: The table reports mean values for continuous variables and percentages for categorical variables. Sample includes $N = 109$ participants from LUISS Guido Carli University.

a direct test of whether deliberate randomization generalizes beyond strategic environments, and offers structural evidence on the cognitive and strategic mechanisms underlying stochastic choice.

3 Empirical Analysis

This section presents the empirical analysis, structured to strenghten the framework of Agranov and Ortoleva (2017), and further explore the motives of deliberate randomization. We begin by replicating their empirical approach using our new set of lotteries, designed to preserve the same variation in dominance and difficulty while allowing for richer structural comparisons. This first step provides a direct benchmark: it verifies that the main empirical regularities of deliberate randomization—its prevalence across both interleaved and consecutive repetitions, its near absence under first-order dominance, and its stability at the subject level—also emerge in our data.

Having established this correspondence, we then move beyond replication and investigate the mechanisms underlying stochastic choice. Specifically, we decompose the determinants of randomization into two broad components: individual heterogeneity, capturing stable behavioral traits and self-reported strategies, and structural characteristics of the lotteries, capturing variation in their underlying primitives such as expected value, variance, and the probabilities of extreme outcomes. This perspective allows us to assess whether stochastic choice is primarily shaped by cognitive or attentional factors, as in standard bounded-rationality accounts, or by intrinsic features of the lotteries themselves. Then, we will collect insights from this analysis and further analyze determinants of lottery choice in a pair, in order to present a framework that allows to understand more deeply strategic randomization.

3.1 Baseline Evidence on Stochastic Choice

The first part of the analysis aims to establish whether our data reproduce the key empirical findings of stochastic choice documented in Agranov and Ortoleva (2017). To this end, we use the same indicator of stochastic choice– which we refer to as intra–lottery measure: for each participant and lottery pair, the indicator takes value one if the choices made across ripetitions in the phase differ at least once. This measure allows for a direct comparison with previous evidence and provides a consistent benchmark for assessing both the prevalence and the persistence of stochastic behavior across phases.

Figure 1 decomposes stochastic choice by experimental phase and difficulty. In the interleaved condition (Phase 1), 34.3% of "easy" and 38.4% of "hard" decisions differ from the reference choice. When the same pairs are repeated consecutively (Phase 2), the corresponding rates fall to 20.4% and 27.1%. The decline from Phase 1 to Phase 2 indicates that part of the randomization observed in the interleaved condition may reflect attentional fluctuations or cognitive load when tasks are spaced apart. Yet, the persistence of substantial randomization even when choices are made in immediate sequence-and participants are fully aware that each pair is repeated three times-suggests that a large share of stochastic choice is strategic and deliberate. This feature of the design allows us to separate noise from intentional randomization, strengthening the inference that stochastic behavior is a consistent element of preference expression rather than an artifact of inattention.

The stability of randomization across phases is formally confirmed in Table 3. The probability of randomizing in Phase 2 is strongly and positively predicted by whether the same participant randomized on that lottery pair in Phase 1. This relation indicates that randomization is highly persistent at the subject level, consistent with a stable underlying propensity rather than transitory error. This finding parallels and reinforces the evidence reported by Agranov and Ortoleva (2017), who also find that the same individuals tend to randomize systematically across temporal distances and contexts.

Table 3: Persistence of stochastic choice across phases

	All lotteries (excl. FOSD)	**Hard lotteries (4–7)**
Dependent variable:	Indicator for stochastic choice in Phase 3	
Independent variable:	Indicator for stochastic choice in Phase 1	
Constant (α)	0.188*** (0.025)	0.218*** (0.033)
Phase 1 indicator (β)	0.148*** (0.038)	0.137** (0.051)
Clustered s.e.	by subject	by subject
Observations	756	432

Notes. Each observation is a subject–lottery pair. The dependent variable equals 1 if a participant randomizes across repetitions of the same lottery in Phase 3; the independent variable equals 1 if randomization occurs in Phase 1. Standard errors are clustered by subject. Significance levels: ***$p < 0.01$, **$p < 0.05$, *$p < 0.10$.

In summary, our data reproduce and strengthen the core empirical findings on deliberate stochastic choice. Randomization is systematic, robust to task transparency, and virtually absent under dominance. The design of

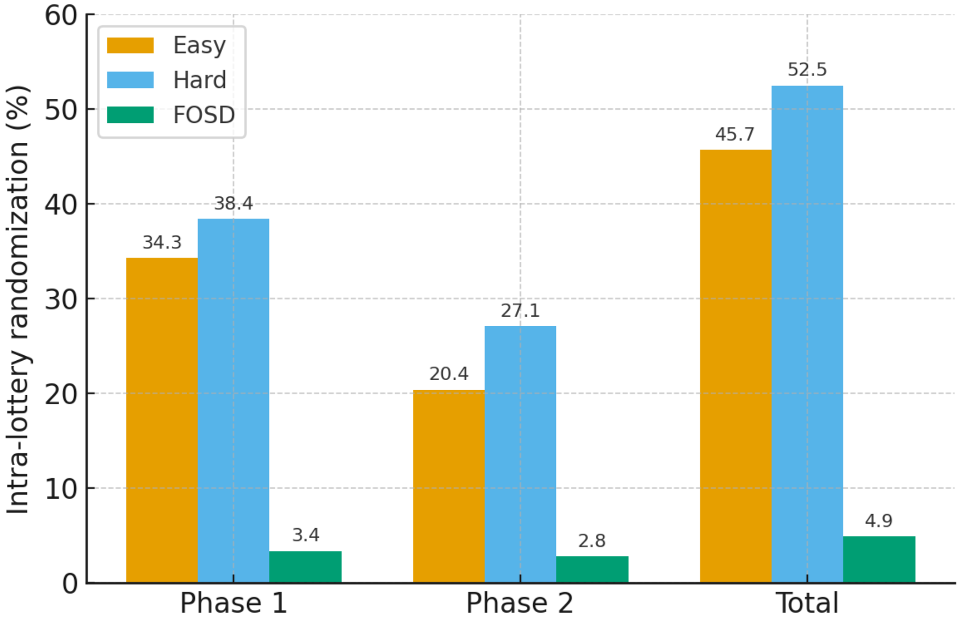

Figure 1: Intra–lottery randomization by phase and difficulty. Bars indicate the percentage of trials where the participant's choice differs from her reference choice for the same lottery pair.

the experiment-with repetitions both interleaved and consecutive-makes it possible to distinguish a small component of randomization plausibly due to cognitive load or attention, from a large and stable component consistent with deliberate strategic randomization. The next subsection examines how this deliberate component varies with both the structural properties of the lotteries and the individual characteristics of decision makers.

3.2 Determinants of Stochastic Choice

We study the determinants of randomization with a trial–level logit. The dependent variable equals 1 if, on a given repetition of pair ℓ by subject i, the choice differs from the subject's *reference* choice for that pair (the first time the pair appears within the phase). Explanatory variables include absolute differences in attributes between the two lotteries ($|\Delta EV|$, $|\Delta SD|$), an indicator for selecting the riskier option in that trial (*riskychoice*), a binary difficulty label (*difficulty*), the within–block repetition order (*rep*), response time, a Phase–2 indicator, survey covariates (clarity of presentation, use of a strategy, perceived comfort, etc.), and demographic and region controls. Standard errors are clustered by subject.

Because FOSD pairs combine very large $|\Delta EV|$ with (virtually) zero randomization, including them mechanically induces a spurious negative relation between $|\Delta EV|$ and randomization. Since the object of interest is behavior in non–dominated comparisons, the baseline specification excludes FOSD pairs.For completeness, the specification that includes FOSD pairs is reported in Appendix A.

The coefficients indicate that randomization increases with $|\Delta SD|$ and when the subject selects the riskier option in that trial. Neither $|\Delta EV|$ nor the difficulty label is predictive once FOSD pairs are excluded. Self–reported clarity of presentation is negatively associated with randomization; reported strategy use and perceived comfort are positively associated. Response time and repetition order are not significant. These results point to risk/volatility—rather than expected–value differences or nominal difficulty—as the relevant correlates of intra–lottery randomization in non–dominated comparisons. These results differs from traditional interpretations of deliberate randomization, that consider stochasticity -even when the result of an intentional strategy- still aimed at reducing or saving cognitive effort. Our results suggest indeed that complexity is not enough to explain alone all the observed randomization, that seems to arise when the riskier option is chosen at least once across the repetitions. This suggests that stochastic choice is not merely a byproduct of cognitive load, attentional limits, or error, but reflects a systematic and strategical behavioral response to environments in which salient attributes—such as high payoff probability or payoff dispersion—pull preferences in competing directions.

Taken together, the evidence points to a mechanism in which strategic mixing arises not from difficulty per se,

Table 4: Logit: intra–lottery randomization (trial level), non–FOSD pairs

	Estimate	Std. Error	t-stat	p-value
Intercept	-1.495	0.523	-2.855	0.004
$\lvert\Delta SD\rvert$	0.0337^{**}	0.0162	2.076	0.038
$\lvert\Delta EV\rvert$	0.0052	0.0280	0.186	0.852
riskychoice	0.741^{***}	0.0816	9.081	$< 10^{-18}$
difficulty	0.058	0.0810	0.721	0.471
repetition	0.071	0.0435	1.623	0.105
response_time	0.000001	0.000161	0.006	0.995
Phase 2 indicator	0.162	0.0833	1.943	0.053
Female	0.130	0.0760	1.711	0.087
Survey: instructions	0.252	0.237	1.065	0.287
Survey: presentation	-0.716^{***}	0.254	-2.818	0.005
Survey: choosing difficulty	-0.140	0.159	-0.884	0.377
Survey: easier choices	0.023	0.100	0.230	0.818
Survey: comfortable	0.542^{**}	0.232	2.342	0.019
Survey: strategy	0.225^{***}	0.0757	2.976	0.003
Survey: preference	-0.026	0.103	-0.251	0.802
Survey: consistency	-0.164	0.192	-0.852	0.394
Faculty: Economics	-0.453^{***}	0.0986	-4.596	$< 10^{-4}$
Faculty: Political Science	-0.481^{***}	0.125	-3.858	< 0.001
Region: Center	-0.181	0.129	-1.578	0.115
Region: South/Islands	0.169	0.128	1.320	0.187
Observations		4,536		

Notes. Dependent variable equals 1 if the trial choice differs from the subject's reference choice for that pair (within phase). FOSD pairs excluded. Standard errors clustered by subject. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

but from the structure of the risk trade-offs embedded in the lotteries. The next subsection studies lottery selection within a pair in order to deepen the analysis of the potential mechanisms underlying strategic randomization. We relate the probability of choosing one lottery to the signed differences in primitives (including the probabilities of extreme outcomes), providing the link to the attribute–salience interpretation developed later.

3.3 Attribute Salience and Strategic Randomization

We model the probability of selecting one lottery within a pair using a logit at the trial level. Let ΔEV, ΔSD, $\Delta p0$, and $\Delta p30$ denote the signed differences (right minus left) in expected value, standard deviation, probability of zero payoff, and probability of the highest payoff (30). The dependent variable equals 1 if the right–hand lottery is chosen. The sample excludes FOSD pairs. Standard errors are clustered by subject; phase and difficulty indicators do not alter the estimates reported below.

Table 5: Logit: choice within a pair (non–FOSD)

	Estimate	Std. Error	t-stat	p-value
Intercept	1.464	0.109	13.43	0.000
ΔEV	0.192^{***}	0.028	6.88	0.000
ΔSD	0.089^{**}	0.041	2.14	0.032
$\Delta p0$	-1.946^{***}	0.455	-4.28	0.000
$\Delta p30$	4.130^{***}	0.727	5.68	0.000
Observations		4,536		

Notes. Dependent variable: indicator for choosing the right–hand lottery. Δ variables are defined as right minus left. FOSD pairs excluded. SEs clustered by subject. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

The estimates show that selection is organized by three primitives. First, extreme–outcome probabilities are first–order predictors: $\Delta p0 < 0$ and $\Delta p30 > 0$ with large magnitudes. Second, $\Delta EV > 0$ and $\Delta SD > 0$ also increase the probability of choosing the corresponding option, with a smaller coefficient on risk. Phase indicators are small and do not affect these slopes; difficulty labels are not robust predictors once primitives enter, in line with the established evidence when FOSD are excluded.

The Marschak–Machina representation clarifies the mechanism. When attributes are aligned (e.g., higher EV and lower $p0$ on the same option), preferences are clear and intra–lottery randomization is low. When attributes conflict (e.g., the option with higher $p30$ also has higher $p0$ or lower EV), the choice margin shrinks and randomization increases.

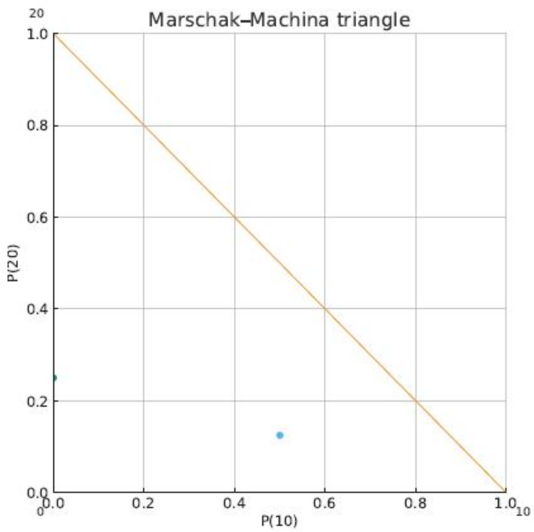

Figure 2: *
Pair 3: higher EV also implies lower $p0 \Rightarrow$ alignment, low stochasticity.

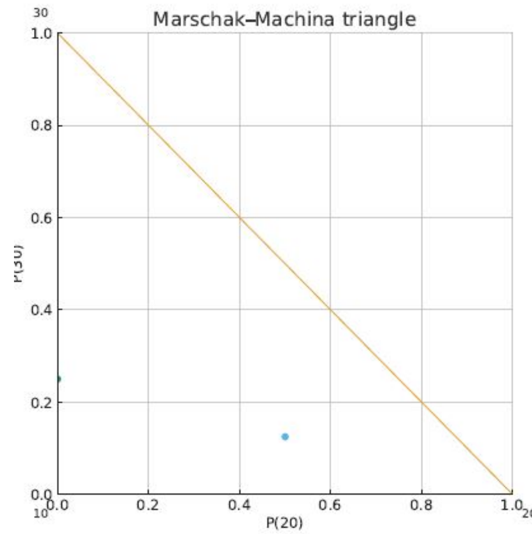

Figure 3: *
Pair 5: lower EV option has higher $p30 \Rightarrow$ conflict, higher stochasticity.

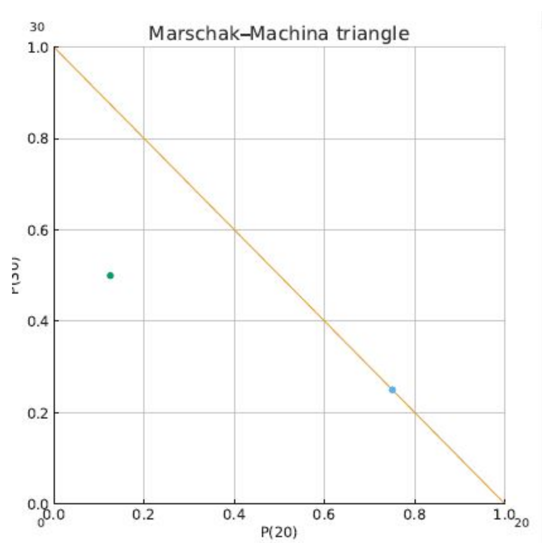

Figure 4: *
Pair 2: MM–triangle as in Pair 6.

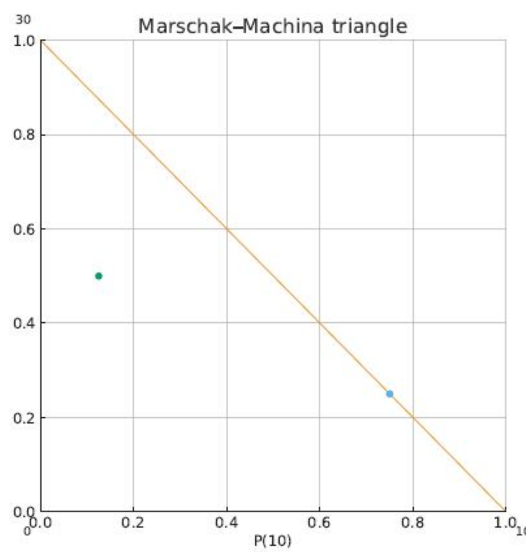

Figure 5: *
Pair 6: option with $p0 > 0$ also has higher $p30 \Rightarrow$ conflict, high stochasticity.

The Marshak-Machina representation of all lottery choices is reported in the Appendix.

Since choice probabilities load on $\Delta p0$ and $\Delta p30$ alongside ΔEV and ΔSD, and stochasticity rises precisely when these attributes point in opposite directions, the salience of extreme–outcome cues offers a natural explanation of deliberate randomization.

We now connect these observations to a "high–stakes" mechanism. When salient attributes point in different directions, no single lottery clearly dominates the comparison. In such situations, one option may be preferred on some dimensions, while the alternative assigns greater weight to the highest payoff. If this trade-off does not generate a decisive ranking, agents may rationally randomize across repetitions in order to occasionally "try" the high-stake option, using stochastic choice as a strategic device under attribute conflict.

To capture this "go for the high stake" motive when attributes are in conflict, we introduce an indicator *HighStake* that equals 1 when the chosen lottery assigns a higher probability weight to the maximum payoff (30), and an indicator of *attribute conflict* that equals 1 for lottery pairs in which expected value, the probability of the zero outcome, and the probability of the highest payoff are not aligned across options, so that no single attribute clearly favors the same lottery.

Replacing the *riskychoice* dummy with *HighStake* in the trial-level logit for stochasticity yields the specification reported in Tablee 6.

Table 6: Determinants of trial-level stochastic choice

	Coefficient	**Clustered SE**	**p-value**
Intercept	-1.224	(0.901)	0.174
$\lvert\Delta SD\rvert$	0.048^{*}	(0.028)	0.091
$\lvert\Delta EV\rvert$	-0.229^{***}	(0.048)	< 0.001
High-stake choice	0.597^{***}	(0.196)	0.002
Attribute conflict	0.855^{***}	(0.184)	< 0.001
Difficulty (Hard)	-0.190	(0.175)	0.277
Repetition	0.073^{***}	(0.028)	0.008
Response time	-0.000003	(0.000160)	0.983
Phase 2 indicator	0.142	(0.147)	0.335
Female	0.176^{*}	(0.068)	0.098
Survey: instructions	0.220	(0.158)	0.163
Survey: presentation	-0.828	(0.551)	0.133
Survey: choosing difficulty	-0.210	(0.342)	0.540
Survey: easier choices	-0.111	(0.205)	0.588
Survey: comfortable	0.528	(0.457)	0.249
Survey: strategy	0.295^{*}	(0.162)	0.068
Survey: preference	-0.118	(0.212)	0.576
Survey: consistency	-0.193	(0.448)	0.666
Faculty: Economics	-0.132	(0.164)	0.421
Faculty: Political Science	omitted		
Region: Center	-0.239	(0.149)	0.109
Region: South/Islands	omitted		
Observations		4536	
SE		Clustered by participant	

Notes. The dependent variable is a trial-level indicator equal to one if the choice differs from the subject's reference choice for the same lottery pair within the same phase. One category for faculty and region is omitted due to collinearity with the intercept. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 6 reports the results of the trial-level logit regression for stochastic choice. Stochasticity is strongly related to the structure of the choice problem rather than to nominal measures of difficulty or cognitive effort, with variables capturing attention to salient outcomes exhibit the strongest associations. The coefficient on *HighStake* is large and highly significant, showing that stochastic choice is much more likely when subjects select the lottery that places greater weight on the maximum payoff. Similarly, the indicator for *Attribute conflict* displays a strong and robust positive effect. Randomization increases sharply in lottery pairs where expected value, the probability of zero, and the probability of the highest payoff do not align, so that no single attribute clearly dominates the comparison. Together, these results point to attribute-level tension and attention to extreme outcomes as central drivers of stochastic choice.

The joint pattern of Table 6 is consistent with a salience mechanism centered on extreme–outcome attributes. To further probe our strategic interpretation of stochastic choice, we extend the baseline specification by allowing the effect of pursuing the high-payoff option to depend on the presence of attribute conflict. Specifically, we augment the trial-level logit with an interaction between *HighStake* and *Attribute conflict*. This specification tests whether randomization is especially likely when attention to extreme outcomes operates under conflicting salient attributes.

Table 7: Trial-level stochastic choice with High-Stake × Attribute Conflict

	Coefficient	**Clustered SE**	**p-value**
Intercept	-1.373	(0.941)	0.145
$\lvert\Delta SD\rvert$	0.113^{***}	(0.036)	0.002
$\lvert\Delta EV\rvert$	-0.174^{***}	(0.040)	< 0.001
High-stake choice	-0.485	(0.311)	0.118
Attribute conflict	0.056	(0.262)	0.830
High-stake × Conflict	1.521^{***}	(0.414)	< 0.001
Difficulty (Hard)	0.186	(0.178)	0.295
Repetition	0.073^{***}	(0.027)	0.008
Response time	-0.00002	(0.00015)	0.904
Phase 2 indicator	0.162^{**}	(0.067)	0.016
Female	0.125	(0.152)	0.410
Survey: instructions	0.205	(0.170)	0.228
Survey: presentation	-0.833	(0.553)	0.132
Survey: choosing difficulty	-0.237	(0.345)	0.492
Survey: easier choices	-0.144	(0.212)	0.495
Survey: comfortable	0.545	(0.445)	0.221
Survey: strategy	0.328^{**}	(0.165)	0.047
Survey: preference	-0.090	(0.216)	0.676
Survey: consistency	-0.151	(0.475)	0.750
Faculty: Economics	-0.138	(0.167)	0.408
Region: Center	-0.249^{*}	(0.151)	0.100
Observations		4536	
SE		Clustered by subject	

Notes. The dependent variable equals one if the choice differs from the subject's reference choice for the same lottery pair within the same phase. One category for faculty and region is omitted due to collinearity. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

The interaction results provide strong support for a strategic interpretation of stochastic choice. While neither *HighStake* nor *Attribute conflict* alone predicts randomization once both are included, their interaction is large, positive, and highly significant. This indicates that stochastic choice emerges precisely when the pursuit of the high-payoff option takes place in lottery pairs characterized by conflicting salient attributes. In such cases, agents appear to randomize as a way to occasionally "go for the high stake" when no single attribute provides a clear ranking of the options.

To organize the empirical findings and provide a unifying interpretation, we introduce a reduced-form salience-based representation of choice and randomization. Let $\Delta = (\Delta EV, \Delta SD, \Delta p0, \Delta p30)$ denote the signed differences in expected value, dispersion, and extreme-outcome probabilities within a lottery pair. We define a salience-weighted valuation index

$$V_i(\Delta) = w_{EV,i}\,\Delta EV - w_{0,i}\,\Delta p0 + w_{30,i}\,\Delta p30,$$

where the weights $w_{k,i}$ capture subject-specific attention to attributes.

This index governs deterministic choice whenever a clear ranking emerges. When the weight placed on the probability of the highest payoff is sufficiently large relative to the weight on downside risk, the valuation tilts toward the high-stake lottery, consistent with the empirical role of $\Delta p30$ in within-pair choice and with the strong association between high-stake selection and trial-level stochasticity.

Randomization arises when the valuation margin $|V_i(\Delta)|$ is small because salient attributes point in opposing directions—for instance, when a higher probability of the top payoff is offset by a higher probability of zero or a lower expected value. In this region, agents mix across repetitions, which we summarize by

$$\Pr(randomize_{i\ell r}) = \sigma\big(\tau_i - |V_i(\Delta_\ell)|\big),$$

where $\sigma(\cdot)$ is increasing and τ_i captures heterogeneity in tolerance for mixing. This formulation rationalizes why stochastic choice concentrates in comparisons characterized by extreme-outcome trade-offs, and why dispersion and high-stake selection predict randomization.

Taken together, this reduced-form framework interprets deliberate stochastic choice as a salience-driven response to competing extreme-outcome signals, rather than as a byproduct of expected-value indifference, cognitive load, or nominal task difficulty.

4 Conclusion

This paper revisits the foundations of stochastic choice by combining a replication of Agranov and Ortoleva (2017) with a structurally richer lottery environment, and by developing a deeper account of the mechanisms underlying deliberate randomization. Rather than treating stochasticity as the residual of cognitive noise or limited attention, we show that mixing follows systematic patterns that are tightly linked to the attribute structure of the decision environment.

Our experimental design separates genuinely stochastic behavior from mechanical variation arising within repetitions, allowing us to study both the incidence and the determinants of mixing. Across specifications, randomization does not respond to nominal difficulty or to expected-value proximity once first-order–dominance pairs are removed. Instead, it depends on volatility, extreme-outcome probabilities, and the propensity to select the high-stake option at least once within a sequence. These patterns indicate that stochastic choice emerges most strongly in environments where the attributes that drive evaluation—probability of zero, probability of the highest payoff, and expected value—pull in different directions. Attribute conflict, rather than cognitive overload, is the key condition under which behavior becomes stochastic. In these environments, randomization appears to function as a strategic device: when the attributes that guide evaluation point in opposing directions, subjects occasionally shift attention toward the option with the highest potential payoff and "try once" for the high-stake outcome. Formally, mixing increases precisely when the probability weight on the extreme outcome is sufficiently salient to offset the expected-value comparison. Stochastic choice thus reflects a deliberate attempt to balance competing evaluative cues by intermittently selecting the high-stake option, rather than an inability to discriminate between alternatives or a breakdown of cognitive resources. The evidence points toward an attribute-salience interpretation of deliberate randomization. When no single attribute provides a dominant evaluative criterion, subjects switch between attribute-based comparisons and mix across repetitions. Randomization thus appears not as an error or a demand on attentional resources, but as a strategic response to competing salient features of the choice set. Individual heterogeneity in reported strategies and in risk orientation further aligns with this view, suggesting that stochasticity reflects purpose-driven adjustments in attention and evaluation rather than involuntary noise.

By reframing deliberate randomization through the lens of attribute salience, the paper offers an alternative theoretical foundation for stochastic choice. This framework provides a tractable interpretation of when and why agents mix, and yields predictions that connect structural features of lotteries to patterns of stochastic behavior. The approach opens new avenues for modeling strategic randomization in environments with limited discrimination, conflicting incentives, or endogenous attention. Future research could extend these insights to dynamic choice, to mechanism-design contexts where mixing carries allocative consequences, or to neuroeconomic settings that study how salient attributes are formed and weighted in real time.

References

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

Aumann, R. J. (1964). Mixed and behavior strategies in infinite extensive games. *Advances in Game Theory*, pages 627–650.

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

Camerer, C. F. (1998). *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press, Princeton.

Carbone, E. and Hey, J. D. (2000). Which error story is best? a reconsideration. *Journal of Risk and Uncertainty*, 20(2):161–176.

Gennaioli, N. and Shleifer, A. (2010). What comes to mind: Cognitive accessibility and judgment. *Quarterly Journal of Economics*, 125(4):1399–1433.

Hey, J. D. (1995). Experimental investigations of errors in decision making under risk. *European Economic Review*, 39(3–4):633–640.

Hey, J. D. (1997). Which error story is best? In Hey, J. D. and Loomes, G., editors, *Experimental Economics*, pages 63–85. Wiley.

Hey, J. D. and Di Cagno, D. (1994). Sequential decision making under risk: Experiments on the role of cognitive ability and risk attitudes. *Journal of Economic Psychology*, 15(4):481–497.

Hey, J. D. and Orme, C. (1994). Investigating generalizations of expected utility theory using experimental data. *Econometrica*, 62(6):1291–1326.

Kahneman, D. (2011). *Thinking, Fast and Slow*. Farrar, Straus and Giroux, New York.

Loomes, G. and Sugden, R. (1995). Incorporating a stochastic element into decision theories. *European Economic Review*, 39(3–4):641–648.

McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. pages 105–142.

Rubinstein, A. (2016). A typology of players: Between instinctive and reflective reasoning. *Quarterly Journal of Economics*, 131(2):859–890.

Appendix A — Pilot study and selection of non-trivial lotteries

Before running the main experiment, we conducted a pilot study with the full set of 60 lotteries borrowed from Hey and Di Cagno (1994). For each lottery, we collected two diagnostic measures: (i) the mean response time, and (ii) the percentage of participants choosing option 1. These two indicators jointly capture the underlying "difficulty" or complexity of a choice: lotteries that generate more heterogeneous choices (i.e., closer to a 50/50 split) or systematically higher response times can be interpreted as involving greater cognitive conflict, whereas lotteries yielding near-unanimous choices or very fast decision times are effectively trivial.

Despite the substantial variation in payoffs and probabilities across the 60 candidate lotteries, the distribution of response times was relatively compressed. This is consistent with the structure of the task: in none of these 60 lotteries does one option first-order stochastically dominate the other, so the choice problem is never "obvious" in the normative sense. Consequently, selection had to be based on subtle differences in conflict intensity rather than on the presence of clear FOSD relationships.

For the purposes of the main experiment we selected seven lotteries: three "easy" items (indices 1, 2, and 15) and four "hard" items (indices 20, 28, 35, and 44). The easy items were chosen among those with the lowest response times that have less choice heterogeneity, while the hard items were chosen among those with the highest response times that have the greatest heterogeneity. Table 7 reports all 60 lotteries together with their mean response time in the pilot and the share of participants selecting option 1.

Idx	Side	Payoffs	Probs	Mean RT	% choosing 1	Class
1	L1	(0, 10)	(0.625, 0.375)	6.37	71.41	Easy
	L2	(0, 30)	(0.875, 0.125)			
2	L1	(20, 30)	(0.75, 0.25)	5.44	89.11	Easy
	L2	(0, 20, 30)	(0.375, 0.125, 0.5)			
3	L1	(10, 30)	(0.5, 0.5)	8.65	42.59	
	L2	(0, 30)	(0.25, 0.75)			
4	L1	(10, 20)	(0.875, 0.125)	7.22	88.89	
	L2	(0, 20)	(0.75, 0.25)			
5	L1	(0, 10)	(0.625, 0.375)	8.76	11.11	
	L2	(0, 30)	(0.75, 0.25)			
6	L1	(0, 20)	(0.125, 0.875)	6.78	68.89	
	L2	(0, 30)	(0.625, 0.375)			
7	L1	(20, 30)	(0.875, 0.125)	6.87	70.37	
	L2	(0, 30)	(0.375, 0.625)			
8	L1	(10, 20)	(0.75, 0.25)	8.28	81.48	
	L2	(0, 10, 20)	(0.375, 0.125, 0.5)			
9	L1	(10, 30)	(0.875, 0.125)	6.74	29.63	
	L2	(0, 30)	(0.375, 0.625)			
10	L1	(20)	(1)	6.96	44.44	
	L2	(10, 30)	(0.5, 0.5)			
11	L1	(10, 20)	(0.5, 0.5)	6.63	37.04	
	L2	(10, 30)	(0.75, 0.25)			
12	L1	(20, 30)	(0.875, 0.125)	8.02	85.19	
	L2	(10, 30)	(0.75, 0.25)			
13	L1	(20, 30)	(0.5, 0.5)	6.62	64.81	
	L2	(0, 30)	(0.25, 0.75)			
14	L1	(0, 10, 30)	(0.125, 0.5, 0.375)	6.89	12.96	
	L2	(0, 30)	(0.25, 0.75)			
15	L1	(0, 10, 20)	(0.375, 0.5, 0.125)	5.91	94.44	Easy
	L2	(0, 20)	(0.75, 0.25)			
16	L1	(10, 20, 30)	(0.125, 0.5, 0.375)	6.67	11.11	
	L2	(10, 30)	(0.25, 0.75)			
17	L1	(0, 20)	(0.125, 0.875)	7.67	74.07	
	L2	(0, 30)	(0.25, 0.75)			
18	L1	(0, 10)	(0.5, 0.5)	6.85	57.41	
	L2	(0, 20)	(0.75, 0.25)			
19	L1	(0, 20)	(0.5, 0.5)	7.17	79.63	
	L2	(0, 30)	(0.75, 0.25)			
20	L1	(0, 10)	(0.625, 0.375)	7.61	67.59	Hard
	L2	(0, 20)	(0.875, 0.125)			

Idx	Side	Payoffs	Probs	Mean RT	% choosing 1	Class
21	L1	(10, 30)	(0.875, 0.125)	6.98	59.26	
	L2	(0, 30)	(0.75, 0.25)			
22	L1	(10, 20)	(0.875, 0.125)	6.96	40.74	
	L2	(0, 20)	(0.375, 0.625)			
23	L1	(10)	(1)	7.19	29.63	
	L2	(0, 30)	(0.625, 0.375)			
24	L1	(10)	(1)	6.69	62.96	
	L2	(0, 20)	(0.625, 0.375)			
25	L1	(10, 20)	(0.125, 0.875)	5.67	64.81	
	L2	(10, 30)	(0.625, 0.375)			
26	L1	(20)	(1)	6.35	72.22	
	L2	(0, 30)	(0.5, 0.5)			
27	L1	(10)	(1)	6.93	90.74	
	L2	(0, 30)	(0.375, 0.625)			
28	L1	(10, 20, 30)	(0.375, 0.5, 0.125)	7.96	64.44	Hard
	L2	(10, 30)	(0.75, 0.25)			
29	L1	(10)	(1)	7.33	42.59	
	L2	(0, 30)	(0.5, 0.5)			
30	L1	(0, 10, 20)	(0.125, 0.5, 0.375)	4.46	24.07	
	L2	(0, 20)	(0.25, 0.75)			
31	L1	(0, 10)	(0.5, 0.5)	6.72	29.63	
	L2	(0, 30)	(0.75, 0.25)			
32	L1	(0, 20)	(0.375, 0.625)	6.78	64.81	
	L2	(0, 30)	(0.75, 0.25)			
33	L1	(20, 30)	(0.875, 0.125)	6.39	57.41	
	L2	(0, 30)	(0.75, 0.25)			
34	L1	(10, 20)	(0.5, 0.5)	7.28	64.81	
	L2	(0, 20)	(0.25, 0.75)			
35	L1	(10, 30)	(0.75, 0.25)	8.30	57.41	Hard
	L2	(0, 10, 30)	(0.375, 0.125, 0.5)			
36	L1	(20)	(1)	5.89	59.26	
	L2	(0, 30)	(0.375, 0.625)			
37	L1	(0, 20)	(0.375, 0.625)	7.09	53.70	
	L2	(0, 30)	(0.875, 0.125)			
38	L1	(0, 20)	(0.25, 0.75)	6.69	64.81	
	L2	(0, 20, 30)	(0.5, 0.125, 0.375)			
39	L1	(10)	(1)	5.65	61.11	
	L2	(0, 20)	(0.375, 0.625)			
40	L1	(10, 20)	(0.625, 0.375)	6.93	64.81	
	L2	(10, 30)	(0.875, 0.125)			
41	L1	(0, 10)	(0.125, 0.875)	6.54	57.41	
	L2	(0, 30)	(0.625, 0.375)			
42	L1	(0, 10)	(0.25, 0.75)	6.56	53.70	
	L2	(0, 10, 30)	(0.5, 0.125, 0.375)			
43	L1	(10)	(1)	6.17	61.11	
	L2	(0, 20)	(0.5, 0.5)			
44	L1	(20, 30)	(0.875, 0.125)	7.89	53.70	Hard
	L2	(10, 30)	(0.375, 0.625)			
45	L1	(10, 20)	(0.125, 0.875)	5.26	51.85	
	L2	(10, 30)	(0.25, 0.75)			
46	L1	(0, 10)	(0.125, 0.875)	5.13	53.70	
	L2	(0, 30)	(0.25, 0.75)			
47	L1	(10, 20)	(0.875, 0.125)	5.67	66.67	
	L2	(0, 20)	(0.625, 0.375)			
48	L1	(20)	(1)	5.30	62.96	
	L2	(10, 30)	(0.625, 0.375)			
49	L1	(0, 10, 30)	(0.375, 0.5, 0.125)	6.02	57.41	
	L2	(0, 30)	(0.75, 0.25)			
50	L1	(0, 20, 30)	(0.375, 0.5, 0.125)	5.85	51.85	
	L2	(0, 30)	(0.75, 0.25)			
51	L1	(0, 10)	(0.625, 0.375)	6.44	66.67	
	L2	(0, 20)	(0.75, 0.25)			
52	L1	(20, 30)	(0.75, 0.25)	6.67	55.56	
	L2	(10, 20, 30)	(0.375, 0.125, 0.5)			

Idx	Side	Payoffs	Probs	Mean RT	% choosing 1	Class
53	L1	(20, 30)	(0.5, 0.5)	7.44	62.96	
	L2	(10, 30)	(0.25, 0.75)			
54	L1	(0, 20, 30)	(0.125, 0.5, 0.375)	7.52	59.26	
	L2	(0, 30)	(0.25, 0.75)			
55	L1	(20)	(1)	7.87	55.56	
	L2	(10, 30)	(0.375, 0.625)			
56	L1	(10, 20)	(0.25, 0.75)	6.31	55.56	
	L2	(10, 20, 30)	(0.5, 0.125, 0.375)			
57	L1	(20)	(1)	6.78	64.81	
	L2	(0, 30)	(0.625, 0.375)			
58	L1	(0, 10)	(0.125, 0.875)	6.24	51.85	
	L2	(0, 20)	(0.25, 0.75)			
59	L1	(10, 20)	(0.625, 0.375)	4.91	51.11	
	L2	(10, 30)	(0.75, 0.25)			
60	L1	(0, 10)	(0.25, 0.75)	6.52	70.37	
	L2	(0, 10, 20)	(0.5, 0.125, 0.375)			

Table 8: Lotteries used in the pilot, with mean response times and choice frequencies.

Appendix B - Marschak–Machina lottery representations

This appendix reports, for each lottery choice used in the experiment, its corresponding Marschak–Machina (MM) representation together with summary measures of stochastic choice. For every lottery, we display (i) the extent of randomization relative to the reference choice (first choice made by participant on that couple) in Phase 1, Phase 2, and in the aggregate, (ii) the degree of intra-lottery randomization (where a lottery choice is considered a randomization if it differs at least once across repetitions), and (iii) the observed choice frequencies for L1 and L2.

We do not report MM representations for first-order stochastically dominant (FOSD) lotteries, since in those cases the ranking is trivial and no meaningful randomization structure can be inferred beyond simple errors.

Lottery index 1 (easy)

	Phase 1	Phase 2	Total
vs reference	29.0%	31.2%	30.1%
intra-lottery	48.1%	21.3%	58.3%

L1: $0.625 \times 0,\ 0.375 \times 10$ (EV = 3.75)
L2: $0.875 \times 0,\ 0.125 \times 30$ (EV = 3.75)

Observed choices: L1 = 73.6%, L2 = 26.4%

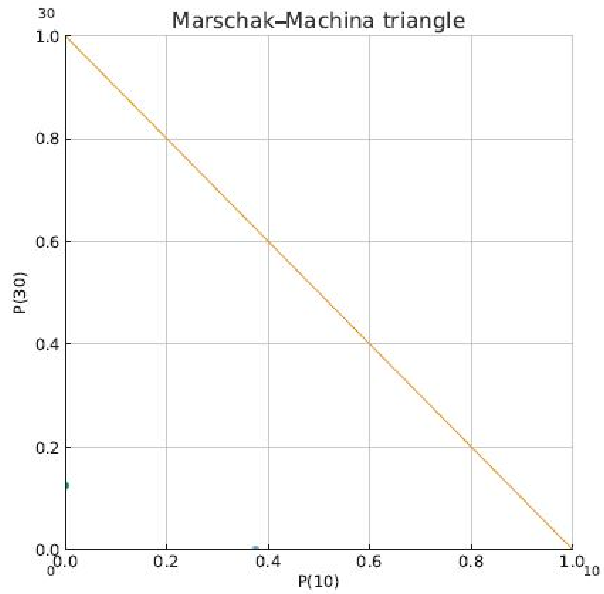

Lottery index 2 (easy)

	Phase 1	Phase 2	Total
vs reference	27.8%	31.2%	29.5%
intra-lottery	41.7%	28.7%	58.3%

L1: $0.750 \times 20,\ 0.250 \times 30$ (EV = 22.50)
L2: $0.375 \times 0,\ 0.125 \times 20,\ 0.500 \times 30$ (EV = 17.50)

Observed choices: L1 = 77.6%, L2 = 22.4%

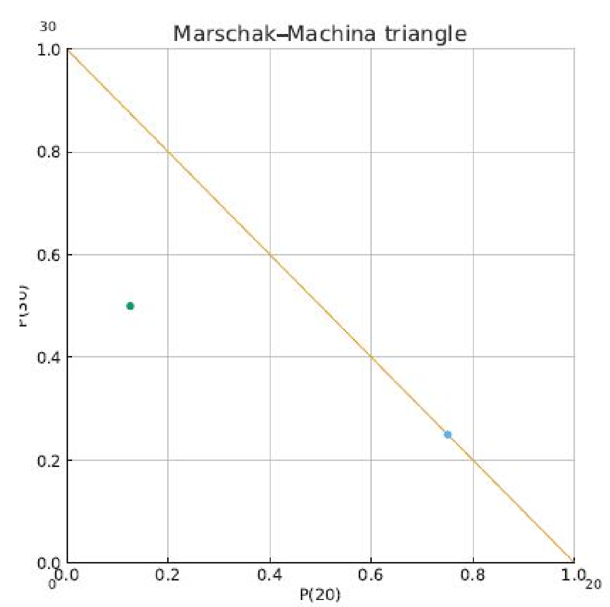

Lottery index 3 (easy)

	Phase 1	Phase 2	Total
vs reference	7.4%	10.8%	9.1%
intra-lottery	13.0%	11.1%	20.4%

L1: $0.375 \times 0,\ 0.500 \times 10,\ 0.125 \times 20$ (EV = 7.50)
L2: $0.750 \times 0,\ 0.250 \times 20$ (EV = 5.00)

Observed choices: L1 = 92.4%, L2 = 7.6%

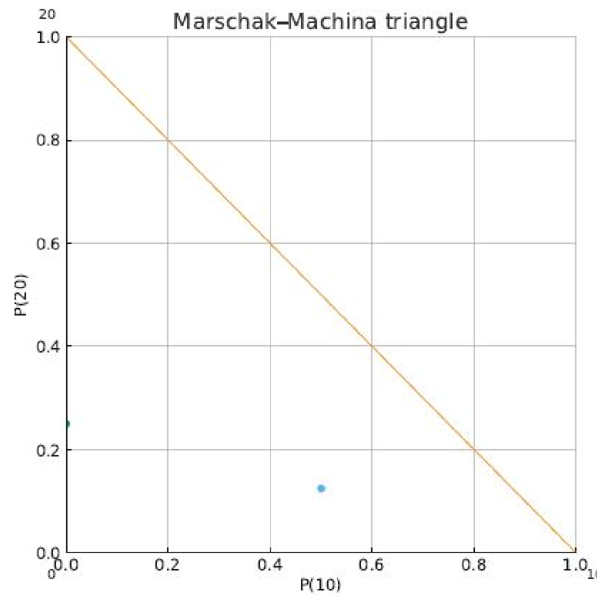

Lottery index 4 (hard)

	Phase 1	Phase 2	Total
vs reference	24.4%	25.6%	25.0%
intra-lottery	33.3%	20.4%	44.4%

L1: $0.625 \times 0,\ 0.375 \times 10$ (EV = 3.75)
L2: $0.875 \times 0,\ 0.125 \times 20$ (EV = 2.50)

Observed choices: L1 = 84.3%, L2 = 15.7%

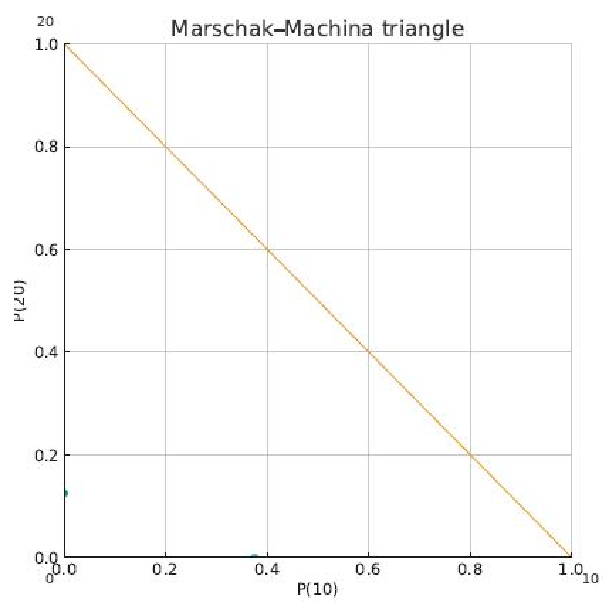

Lottery index 5 (hard)

	Phase 1	Phase 2	Total
vs reference	26.2%	29.9%	28.1%
intra-lottery	38.9%	34.3%	58.3%

L1: $0.375 \times 10,\ 0.500 \times 20,\ 0.125 \times 30$ (EV = 17.50)
L2: $0.750 \times 10,\ 0.250 \times 30$ (EV = 15.00)

Observed choices: L1 = 80.9%, L2 = 19.1%

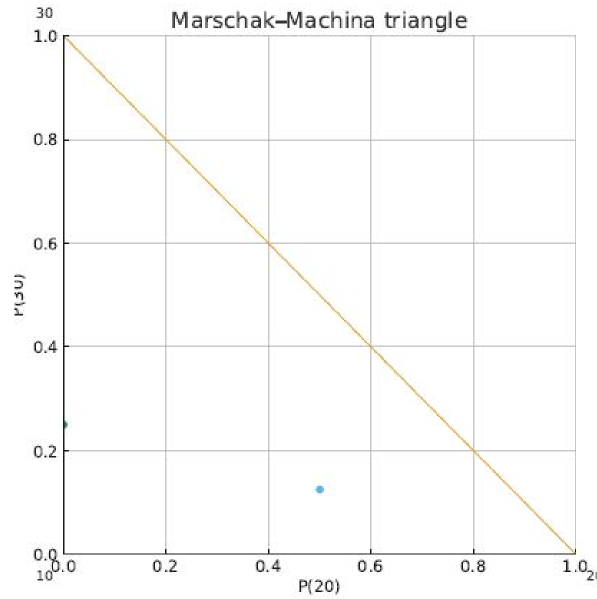

Lottery index 6 (hard)

	Phase 1	Phase 2	Total
vs reference	29.9%	27.8%	28.9%
intra-lottery	49.1%	28.7%	60.2%

L1: $0.750 \times 10,\ 0.250 \times 30$ (EV = 15.00)
L2: $0.375 \times 0,\ 0.125 \times 10,\ 0.500 \times 30$ (EV = 16.25)

Observed choices: L1 = 63.4%, L2 = 36.6%

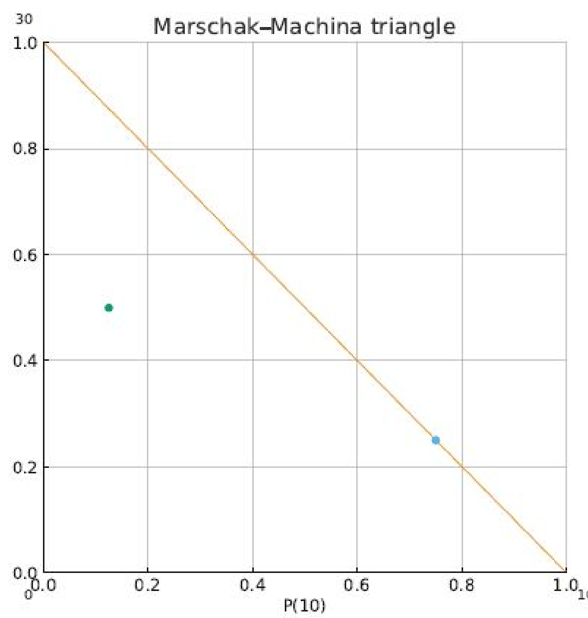

Lottery index 7 (hard)

	Phase 1	Phase 2	Total
vs reference	19.1%	21.6%	20.4%
intra-lottery	32.4%	25.0%	47.2%

L1: $0.875 \times 20,\ 0.125 \times 30$ (EV = 21.25)
L2: $0.375 \times 10,\ 0.625 \times 30$ (EV = 22.50)

Observed choices: L1 = 17.6%, L2 = 82.4%

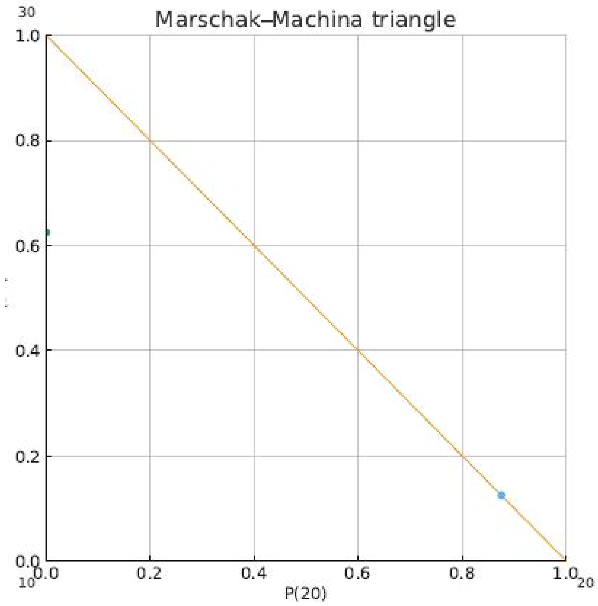

Appendix C - Description of theoretical consistency

Table 9 classifies participants according to the theoretical consistency of their behavior across the two phases. Following the same logic of the previous literature, we identify five behavioral types. A minority of subjects (6.5%) never randomize, consistent with deterministic expected–utility behavior. Roughly one third (31.5%) randomize only in the interleaved condition, which may reflect random–utility or diffusion–drift models with attention lapses. The majority (56.5%) randomize consistently across both phases except for FOSD pairs, a pattern fully consistent with the deliberate stochastic–choice (CSC–type) interpretation. The residual 3–4% exhibit inconsistent or noisy behavior. These shares are remarkably similar to those reported by Agranov and Ortoleva (2017), underscoring the robustness of the phenomenon across experimental designs and different lotteries.

Table 9: Classification of subjects according to theoretical models

Model	**Distant Repetitions**	**Repetitions in a Row**	**Fraction of Subjects**
EU	No	No	6.5%
REU* + DDM*	Yes (at least once) in non-FOSD	No	31.5%
CSC	Yes except FOSD	Yes except FOSD	56.5%
Other (1)	Yes in FOSD	Yes in EASY and/or HARD	2.8%
Other (2)	No	Yes in HARD	0.9%
Total			100%

Notes. REU* + DDM* aggregates subjects who randomize in the interleaved phase but not in the consecutive phase. A subject is classified as stochastically choosing in a part if randomization occurs in at least one lottery of that part. EASY = 1–3, HARD = 4–7, FOSD = 8–10.