Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Salience, Conflict and Strategy in Stochastic Choice

Riccardo Manghi

rmanghi@luiss.it

Libera Università Internazionale degli Studi Sociali Guido Carli

Daniela Di Cagno

Libera Università Internazionale degli Studi Sociali Guido Carli

Gioele Taurino

Giorgio Coricelli

University of Southern California

Research Article

Keywords:

Posted Date: May 27th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9446633/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Salience, Conflict and Strategy in Stochastic Choice

Riccardo Manghi* Daniela Di Cagno† Gioele Taurino Giorgio Coricelli‡

April 17, 2026

Abstract

We study stochastic choice in lottery experiments to distinguish cognitive noise from deliberate randomization. Building on Agranov and Ortoleva (2017), who introduce distinct repetition conditions based on sequential and temporally distant repeats, we extend their design by introducing richer multi-attribute variation in lottery pairs. This allows us to measure randomization within and across repetition conditions and to identify its structural mechanisms and individual determinants. We replicate persistent stochastic choice but show that it is not explained by task difficulty, cognitive load, or standard measures of choice complexity. Instead, randomization arises systematically from conflicts between salient attributes, such as probabilities of extreme outcomes, and normative primitives, such as expected values. These patterns suggest that individuals shift attention across competing attributes across repetitions, leading to deliberate randomization when no option clearly dominates.

*LUISS Guido Carli University. Email: rmanghi@luiss.it.
†LUISS Guido Carli University.
‡University of Southern California.

1 Introduction

Standard economic models assume that individuals make stable and deterministic choices. Yet, empirical evidence consistently shows that even in simple decision problems, individuals frequently switch between alternatives when choices are repeated. Such stochastic behavior has traditionally been interpreted as noise or imperfect optimization. However, a growing literature suggests that variability in choice may instead reflect systematic and deliberative behavioral processes. Understanding the origins of such stochasticity is crucial for models of choice, welfare analysis, and mechanism design. In the stochastic choice literature, randomness in observed decisions is typically interpreted as reflecting noise in the choice process, arising from errors, imprecision, or unstable preferences (Davidson and Marschak, 1959; Tversky, 1969; Hey and Orme, 1994; Loomes and Sugden, 1995; Chew et al., 2022; Cubitt et al., 2015).

An alternative strand of the literature, however, challenges this interpretation by proposing that stochasticity may itself be a deliberate and systematic component of behavior, rather than mere noise (Agranov and Ortoleva, 2017; Cerreia-Vioglio et al., 2015, 2019). In some environments, randomization may serve an instrumental purpose. Individuals facing several plausible options may deliberately mix among them; investors may experiment with riskier assets to explore potential payoffs; and players in repeated strategic interactions may randomize to avoid being predictable. These examples suggest that stochastic behavior may arise not only from cognitive limitations but also from purposeful strategies such as exploration, hedging, or engagement.

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 noise in the choice process. 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 employs 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. Following Agranov and Ortoleva (2017), complexity is defined using heterogeneity of choices across participants and response times.

Beyond extending their framework, our design introduces systematic variation in key lottery attributes, such as the probability of the lowest payoff (including zero) and the probability of the highest payoff. These attributes are constructed so that they may conflict with standard primitives of the lotteries, such as differences in expected value, allowing us to examine how salient payoff components interact with normative value comparisons. To enable clean comparisons, we construct lottery pairs so that alternatives within the two main levels of difficulty remain comparable in their primitives (e.g., expected-value differences). In addition, several pairs are designed to be isomorphic across difficulty levels in their probability representation in the Marschak–Machina triangle. This structure generates controlled variation in attribute salience while holding core primitives constant, allowing us to isolate how conflicts between salient attributes and normative value metrics shape stochastic choice behavior.

Our study relates to three strands of literature. A first line of research studies stochastic choice through random utility models, where randomness is typically interpreted as cognitive noise, implementation error, or preference imprecision (McFadden, 1974; Loomes and Sugden, 1995; Hey, 1995; Hey and Orme, 1994; Carbone and Hey, 2000; Cubitt et al., 2015). A second line of research instead interprets stochasticity as purposeful behavior. In particular, Agranov and Ortoleva (2017)

show that repeated switching may reflect deliberate randomization rather than confusion, while related work studies stochastic choice as a strategic or motivational device (Aumann, 1964; Camerer, 1998; Cerreia-Vioglio et al., 2019). A third line emphasizes that inconsistency may increase when alternatives are close to indifference, so that stochastic choice varies with the strength of preference (Alós-Ferrer et al., 2025). Our results speak to this view, but add a structural interpretation: while inconsistency does tend to rise when expected-utility differences are small, proximity in expected utility alone does not account for the main switching patterns once dominance structure and salient lottery attributes are taken into account. This suggests that strength of preference captures part of the phenomenon, but that it operates through a deeper mechanism: preferences become behaviorally weak precisely in regions where salient attributes of the lotteries are in conflict. In this sense, our findings also connect to the literature on limited attention and salience in decision making (Gennaioli and Shleifer, 2010; Bordalo et al., 2012a, 2013). This result is difficult to reconcile with standard random utility models, in which stochastic choice is driven solely by the proximity of utility levels. In our data, however, similar differences in expected utility can generate markedly different levels of randomization depending on the configuration of salient attributes, suggesting that stochasticity cannot be reduced to a single-index noise around utility. Our results bridge these perspectives by showing that deliberate randomization emerges most clearly when small utility margins coincide with conflict between salient attributes.

Our contribution is threefold. First, we replicate the core evidence of deliberate randomization in a new risky-choice environment with a pilot-validated set of lotteries. Second, we show that the data are only partially consistent with a pure strength-of-preference account: inconsistency is higher when utility differences are small, but this channel alone does not explain the main switching patterns once dominance and lottery structure are taken into account. Third, we provide a structural interpretation of stochastic choice based on attribute salience and attribute conflict. Randomization is most likely when expected value, downside risk, and upside potential point in different directions, so that no single attribute yields a clear ranking of the options. More generally, our findings point to a specific behavioral mechanism underlying stochastic choice. Randomization systematically emerges when salient attributes of lotteries, such as the probabilities of extreme outcomes, conflict with normative value metrics such as expected value. In such environments, individuals appear to shift attention across competing attributes across repetitions, occasionally selecting alternatives that emphasize different payoff dimensions. This interpretation connects deliberate randomization to models of attention-based choice, in which decision weights fluctuate with the salience of attributes rather than remaining fixed.

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 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. Differences in switching rates across phases decompose observed stochasticity into non-deliberate components (e.g., attention or memory noise) and deliberate (strategic) randomization.

Task difficulty is calibrated through a pilot experiment following the approach in Agranov and Ortoleva (2017), classifying this new set of lottery pairs as **Easy** or **Hard** based on heterogeneity and response times. This manipulation allows us to test whether randomization increases with cognitive load. In addition, the inclusion of first order stochastically dominated lottery pairs (FOSD) provides a behavioral control, ensuring that switching behavior cannot be attributed to boredom or mechanical errors. Finally, post-experimental survey measures allow us to relate metacognitive reports to observed randomization 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, grouped into two difficulty levels, EASY and HARD, and added three first order stochastically dominated lotteries (FOSD), yielding a total of ten lotteries across 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.

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 explicitly announced in advanced, 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 dice investment task was administered between Phase 1 and Phase 2. In this task, participants were endowed with a fixed amount and chose how much to allocate between a safe payoff and a risky investment. The return on the risky component was determined by the outcome of a fair six-sided die, generating a known distribution over gains, losses, and neutral outcomes.

Beyond providing a neutral cognitive break while preserving incentive compatibility, this task allows us to elicit individual risk preferences. Under a CARA specification, each participant's allocation identifies a subject-specific risk aversion parameter ρ_i, which we use to compute expected utilities and certainty equivalents for the lotteries in the main task.

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 (MM) triangles: (20,10,0), (30,20,10), (30,20,0), and (30,10,0), with different probabilies across the lotteries. This representation facilitates the visual interpretation of the outcomes that shape choice. The MM representation of all the 10 lottery pairs is reported in the appendix.

The lotteries vary in expected value, standard deviation, and in the probabilities of the lowest (zero) and highest (thirty) outcomes. We treat expected value and standard deviation as normative primitives, and the probabilities of extreme outcomes as salient attributes; both sets of variables enter as covariates in subsequent regressions to identify which features systematically affect choice probabilities and the incidence of randomization. The design thus aims to link observed behavior to interpretable characteristics 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 $\times$ 0, 0.125 $\times$ 20, 0.500 $\times$ 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. Same amounts across different lotteries have been represented with the same colour, to maximize clarity and avoid confusion. An example of how lottery choices were displayed can be observed here:

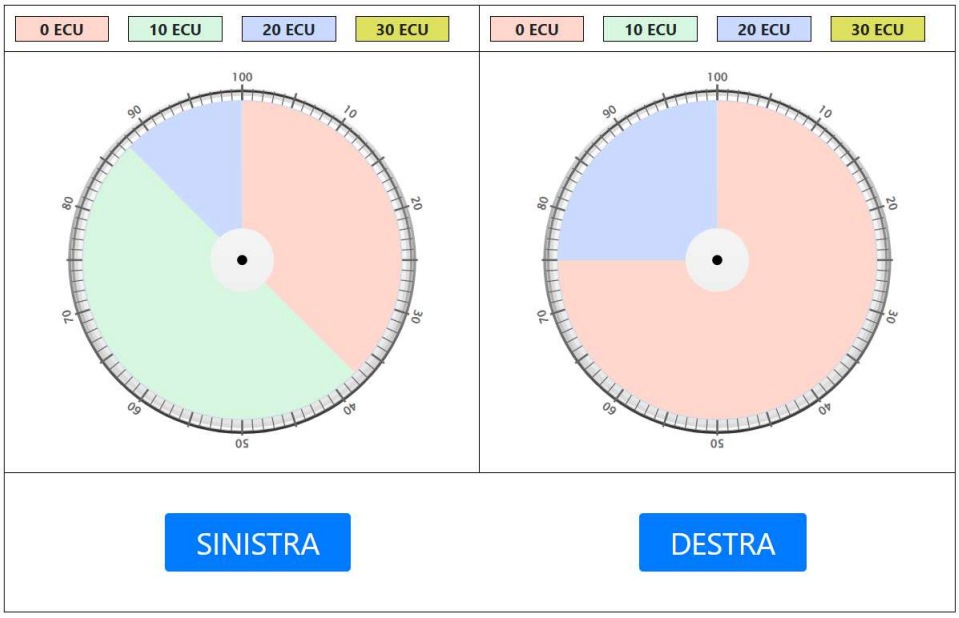


Table 1 shows all the 10 lotteries used for the main experiment are reported, with their graphical Marshak-Machina representation in the appendix.

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 107 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. No one participated to both pilot and experiment.

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.

Table 2 reports summary statistics for the participants.

Table 2: Summary statistics of participants

Variable	**Mean / Share**	**(SD)**
Age	23.8	(2.4)
Female	62.4%	
Major in Economics	68.8%	
Major in Political Science	21.1%	
Major in 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 $N = 107$ participants from LUISS Guido Carli University.

In summary, the experimental design replicates the methodological logic of Agranov and Ortoleva (2017) using a pilot-validated subset of Hey and Di Cagno (1994) lotteries. It provides a direct test of whether deliberate randomization generalizes beyond the original set of lottery choices, and offers structural evidence on the cognitive and strategic mechanisms underlying stochastic choice.

3 Empirical Analysis

This section presents the empirical analysis, structured to replicate the findings 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, (1) its prevalence across both interleaved and consecutive repetitions, (2) its near absence under first-order dominance, and (3) 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 and variance, and salient attributes, such as 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) and Alós-Ferrer et al. (2025). 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), we observed 34.3% stochastic choices in the "easy" trials and 38.4% in "hard" trials. 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.

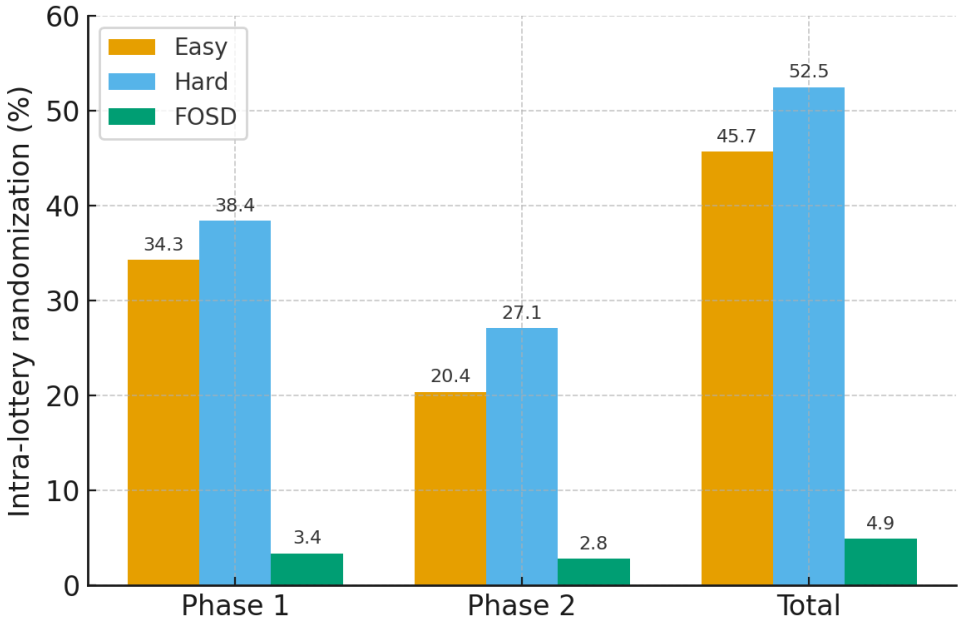


Figure 1: Intra-lottery randomization by phase and difficulty. Bars indicate the percentage of trials where the participant's choice differs from her reference choice for the same lottery pair.

Building on the replication of deliberate randomization patterns, we next examine whether our data are consistent with the strength-of-preference mechanism emphasized by Alós-Ferrer and Buckenmaier (2021). Their central insight is that choice inconsistency increases when alternatives are close to indifference: when the utility difference between options is small, preferences are weaker and behavior becomes more variable rather than fully deterministic.

To investigate this mechanism, we construct pair-level measures of consistency following the Alós-Ferrer et al framework. Specifically, we compute $IR = |\lambda_1 - 0.5|$ from the four Phase 1 repetitions and $ER = |\lambda_2 - 0.5|$ from the three consecutive Phase 2 repetitions, where λ denotes the share of choices of the lottery with weakly higher expected value. We then run a random-effects regressions of these consistency measures on the absolute difference in expected utility between the two lotteries in each pair, thereby replicating the logic of Alós-Ferrer et al. (2025)'s Table 1 within our experimental environment.

Expected utility is evaluated under a Constant Absolute Risk Aversion (CARA) specification, widely used in experimental analyses of risky choice (Conte et al., 2011). Individual risk aversion is inferred from the Dice Investment Task by recovering the coefficient ρ_i that rationalizes each participant's investment decision. Given ρ_i, lottery ℓ is evaluated using $u(x) = -\exp(-\rho_i x)$ and expected utility is computed as $EU_{i\ell} = \sum_o p_\ell(o)\, u(o)$. Table 4 reports the strength-of-preference regressions.

Across specifications, the baseline model shows a positive association between EU-distance and consistency in both phases, suggesting that when lotteries are further apart in expected utility, choices become more stable. However, this relationship vanishes once first-order stochastic dominance (FOSD) pairs are included. The FOSD indicator is large, positive, and highly significant in both phases, indicating substantially higher consistency (lower stochasticity) when one option dominates the other. This confirms that stochastic choice is not merely noise but responds to clear dominance relations in the choice environment.

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	749	428

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: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 4: Strength-of-preference regressions

Random-effects panel regressions on $IR = |\lambda_1 - 0.5|$ (models 1-4) and $ER = |\lambda_2 - 0.5|$ (models 5-8)

	IR (Phase 1)				ER (Phase 2)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
EU-distance	0.237***	0.088	0.073	0.066	0.105**	-0.003	-0.108	-0.107
	(0.065)	(0.063)	(0.088)	(0.088)	(0.051)	(0.050)	(0.071)	(0.071)
FOSD		0.104***	0.099***	0.099***		0.072***	0.057***	0.057***
		(0.010)	(0.011)	(0.011)		(0.008)	(0.009)	(0.009)
Hard			-0.010	-0.011			-0.032***	-0.032***
			(0.013)	(0.013)			(0.010)	(0.010)
Num. outcomes			-0.004	-0.004			-0.001	-0.001
			(0.008)	(0.008)			(0.006)	(0.006)
Female				0.013				-0.008
				(0.009)				(0.013)
Constant	0.389***	0.374***	0.400***	0.393***	0.428***	0.419***	0.453***	0.457***
	(0.009)	(0.008)	(0.032)	(0.033)	(0.009)	(0.008)	(0.026)	(0.026)
Observations	1070	$R^2 = 0.108$						

Notes: Each observation is a subject–lottery-pair. Dependent variables are $IR = |\lambda_1 - 0.5|$, where $\lambda_1 = k_1/4$ is the proportion of choices of the lottery with weakly higher expected value in Phase 1, and $ER = |\lambda_2 - 0.5|$, where $\lambda_2 = k_2/3$ is the corresponding proportion in Phase 2. EU-distance is the absolute difference in normalized CARA expected utility between the two lotteries. Estimates are obtained from random-effects regressions with a random intercept at the subject level. Significance: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

In Phase 2, pair difficulty (Hard) is associated with lower consistency, indicating that attribute conflict in more complex pairs increases stochastic choice when options are presented consecutively. Overall, the results provide limited support for a smooth strength-of-preference gradient.

To complement the regression analysis, we provide a graphical representation of the strength-of-preference mechanism by plotting, for each lottery type, the proportion of inconsistent choices against the difference in expected utility between Lottery 1 and Lottery 2, shown below in Figure 2:

The figure shows that inconsistency is lowest for the FOSD pairs, confirming that stochastic

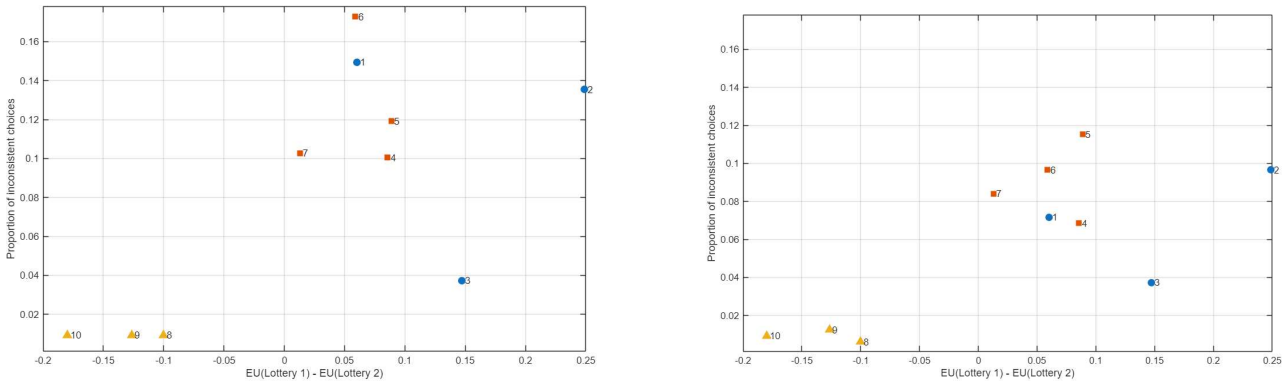


Figure 2: Inconsistency by lottery type as a function of $EU(\text{Lottery 1}) - EU(\text{Lottery 2})$ (Phase 1 left; Phase 2 right).

choice is not driven only by noise. Among the remaining pairs, inconsistency is generally higher when the expected-utility difference is small, consistent with a strength-of-preference effect, although the relationship is not strictly monotonic, underlying the need for a deeper analysis of the motives behind randomization. In summary, our data replicate the core empirical findings on deliberate stochastic choice. Randomization is systematic, robust to task transparency, and virtually absent under dominance. The design of 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. Moreover, strenght of preference do not seem to have the importance found in Alós-Ferrer et al. (2025) .

3.2 Structural Drivers of Lottery Choice

While the baseline analysis confirms the presence of substantial deliberate randomization and its persistence across phases, it does not yet identify the structural attributes of the lotteries that systematically shape subjects' preferences. To uncover these drivers, we estimate a trial-level logit model for the probability that subject i chooses the right-hand lottery in pair ℓ at trial t:

$$\Pr(\text{Choose Right}_{i\ell t}) = \Lambda(\beta_0 + \beta_1 \Delta EV_\ell + \beta_2 \Delta SD_\ell + \beta_3 \Delta p0_\ell + \beta_4 \Delta p30_\ell + \mathbf{X}'_{i\ell t}\boldsymbol{\gamma}), \qquad (1)$$

where $\Lambda(\cdot)$ denotes the logistic CDF and the regressors are signed differences (right minus left) in four key structural primitives. The choice of these four covariates is motivated by both standard economic theory and recent models of attention and salience. ΔEV and ΔSD capture the classical trade-off between expected return and risk, which are central in expected-utility and mean-variance frameworks. More crucially, $\Delta p0$ (difference in the probability of the zero outcome) and $\Delta p30$ (difference in the probability of the highest payoff) represent the most salient attributes in our visual interface (pie-chart probabilities) and in the Marschak-Machina triangles (see below). According to salience theory (Bordalo et al., 2012a, 2013) and models of local thinking (Gennaioli and Shleifer, 2010), decision makers disproportionately overweight extreme outcomes and their associated probabilities. In our lotteries, the probability of receiving nothing ($p0$) is the most visually salient downside, while $p30$ captures the upside potential. Table 5 reports the results:

The estimates show that choice behavior is systematically related to both standard primitives and salient attributes. First, extreme–outcome probabilities are first–order predictors: $\Delta p0 < 0$

Table 5: Logit: choice within a pair

	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	5249			$R^2 = 0.195$

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. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

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 larger coefficient on risk.

The Marschak–Machina (MM) representation provides a useful lens through which to interpret the structure of our lottery pairs and the resulting patterns of stochastic choice. In this framework, lotteries are mapped onto a probability simplex that preserves the distribution of probability mass across low, intermediate, and high outcomes, abstracting from the monetary scale of payoffs. Consequently, lottery pairs that differ only by a translation of outcomes occupy the same relative position in the MM triangle and are isomorphic in terms of their structural characteristics.

Pairs 3 and 5 illustrate this property. Pair 5 is obtained from Pair 3 by adding a constant to all outcomes: the probabilities attached to each outcome level are identical, while the support is shifted upward by 10. As a result, both pairs share the same geometry in the MM triangle and exhibit the same ordering of expected values and extreme–outcome probabilities. Under standard expected-utility theory, which is invariant to affine transformations of payoffs, such pairs should induce identical preference orderings and comparable choice consistency.

Empirically, however, we observe higher stochasticity in Pair 5. This divergence arises because the upward shift alters the salience of key attributes: by eliminating the zero outcome in both lotteries, it removes the downside penalty that previously disadvantaged the lower–EV option, while simultaneously raising the maximum payoff to 30, thereby strengthening the appeal of the high–payoff attribute in that same option. As a result, the lower–EV lottery in Pair 5 becomes more competitive in terms of salient attributes despite its inferior expected value, generating a conflict between expected value and extreme–outcome salience that increases randomization.

This comparison highlights that strength of preference, as measured by expected-utility differences alone, is insufficient to account for stochastic choice. Instead, the interaction between attribute salience and attribute conflict plays a central role in generating deliberate randomization.

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

Figure 3 illustrates the Marschak–Machina representation for lottery pairs 3 and 5.

To compare stochastic behavior across these structurally related lottery pairs, we report bar charts of randomization for Lottery 3 and Lottery 5, separately by phase. Figure 5 compares randomization across structurally similar pairs.

Lottery 5 exhibits substantially higher randomization than Lottery 3 in both measures and across all phases. Despite their isomorphic structure in the Marschak–Machina representation, the upward shift in outcomes increases the salience of the maximum payoff and removes the zero outcome, generating attribute conflict that leads to greater stochastic switching.

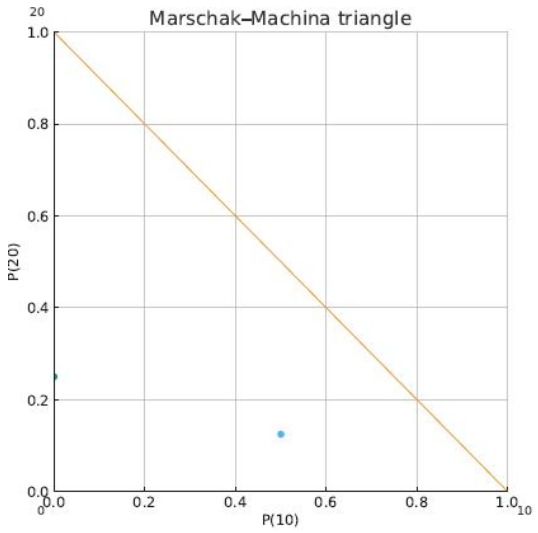


Figure 3: *
Pair 3: higher EV also implies lower $p0 \Rightarrow$ alignment, low stochasticity.

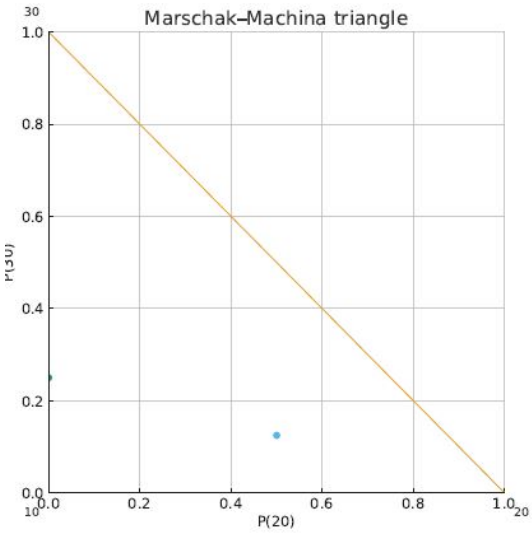


Figure 4: *
Pair 5: lower EV option has higher $p30 \Rightarrow$ conflict, higher stochasticity.

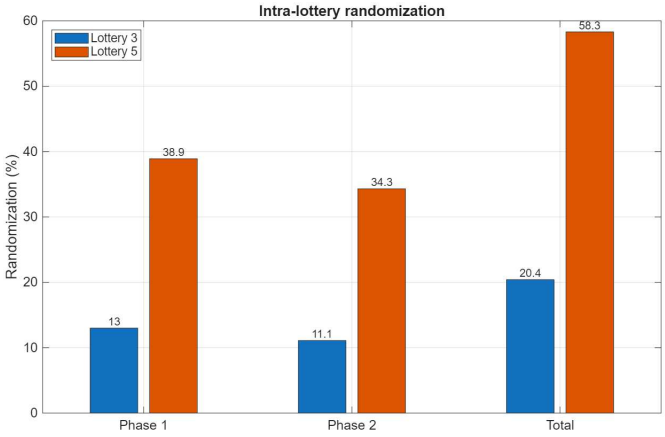


Figure 5: Intra–lottery randomization for Lottery 3 and Lottery 5 by phase.

Since choice probabilities load heavily on $\Delta p0$ and $\Delta p30$ alongside ΔEV and ΔSD, and stochasticity is systematically higher precisely when these attributes point in opposite directions, attribute salience offers a natural explanation of deliberate randomization.

To verify this implication more directly, we replicate Figure 2 using certainty equivalents instead of expected utilities. Because the two pairs share identical probability structures and differ only by a translation of payoffs, CARA certainty equivalents preserve the same ordering and yield identical differences between Lottery 1 and Lottery 2 in Pair 3 and Pair 5.

Figure 6 confirms this property: when the horizontal axis is expressed in terms of certainty equivalents, the two lottery pairs are perfectly aligned. From the perspective of standard expected–utility representations, the strength of preference between the options should therefore be identical across the two pairs.

Empirically, however, stochastic behavior differs sharply across the two pairs. Despite their

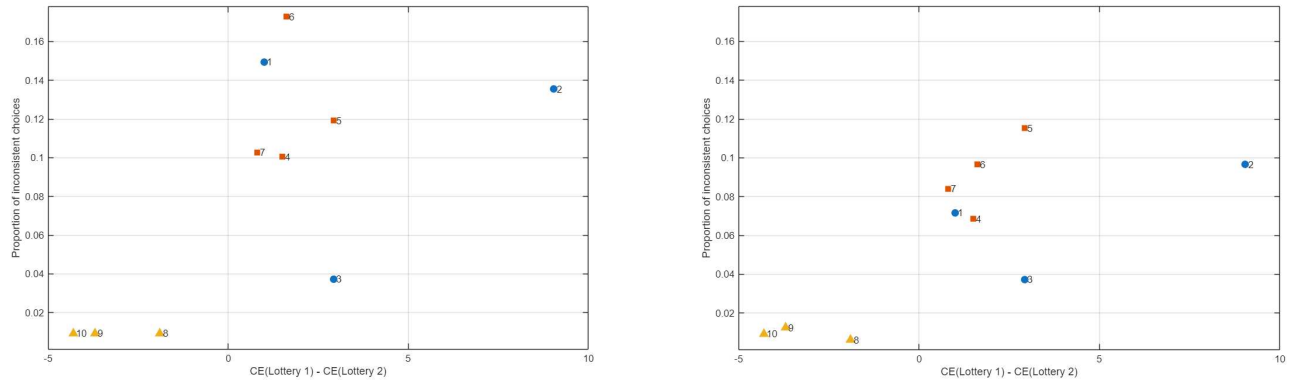


Figure 6: Replication of Figure 2 using certainty equivalents. Pairs 3 and 5 are perfectly aligned when differences are expressed in certainty equivalents.

perfect alignment in certainty–equivalent space, Pair 5 exhibits substantially higher stochasticity than Pair 3.

We now test this mechanism formally. When salient attributes conflict with each other or with expected value, no lottery is clearly dominant across all relevant dimensions. In such situations subjects may engage in deliberate mixing, occasionally "testing" the option that assigns greater weight to the highest payoff (high-stake orientation). To isolate these two channels we construct the indicators *Attribute Conflict* and *HighStake choice* and estimate a trial-level logit for stochastic choice. Results are reported in the next subsection.

3.3 Attribute Salience, Conflict and Deliberate Randomization

The estimates from the choice model (Table 5) already show that $\Delta p0$ (negative) and $\Delta p30$ (positive) exert large effects, often in tension with the ranking induced by ΔEV. This pattern suggests that stochasticity arises precisely when salient attributes conflict with each other or with expected value.

To test this hypothesis directly, we construct two theoretically motivated indicators and estimate a trial-level logit model of stochastic choice. The dependent variable is an indicator equal to one if the subject switches relative to the phase-specific reference choice for a given lottery pair, and zero otherwise.

- **Attribute Conflict**: a dummy equal to 1 when the lottery with higher expected value does *not* also dominate on the salient extreme attributes (i.e., it has lower $p30$ and/or higher $p0$). This measure captures the structural misalignment between standard economic ranking (EV) and attribute salience.

- **HighStake choice**: a dummy equal to 1 if the subject chose the lottery with the higher probability of the maximum payoff (30). This variable proxies the behavioural tendency to 'test' the salient upside option across repetitions.

HighStake is consistent with the largest coefficient on $\Delta p30$ in the choice model. Attribute Conflict is the core mechanism of our salience-based interpretation: when no option dominates across all relevant dimensions, subjects rationally randomize as a form of strategic exploration across competing attributes. Results appear in Table 6.

Table 6: Deliberate randomization and attribute conflict

	Coef.	SE	p-value
ΔEU	-1.480	(1.294)	0.253
HighStake	0.541^{***}	(0.150)	0.0003
Conflict	0.531^{***}	(0.134)	< 0.001
Difficulty (Hard)	-0.037	(0.155)	0.810
Repetition	0.087^{**}	(0.042)	0.038
Response time	0.00012	(0.00015)	0.405
Phase 2	-0.576^{***}	(0.136)	< 0.001
Female	0.126	(0.143)	0.378
Survey: instructions	0.241	(0.205)	0.239
Survey: presentation	-0.135	(0.505)	0.789
Survey: choosing diff.	-0.109	(0.294)	0.710
Survey: easier choices	-0.056	(0.198)	0.777
Survey: comfortable	0.595	(0.377)	0.114
Survey: strategy	0.388^{***}	(0.149)	0.009
Survey: preference	-0.132	(0.207)	0.524
Survey: consistency	0.080	(0.359)	0.825
Faculty: economics	-0.178	(0.147)	0.227
Region: center	-0.298^{**}	(0.147)	0.043
Constant	-2.207^{***}	(0.808)	0.006
Observations	3745	$R^2 = 0.06$	

Notes: Logit estimates with participant-clustered standard errors. The dependent variable is the indicator of switching relative to the phase-specific reference choice. One category for faculty and region is omitted due to collinearity with the intercept. Significance: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

The coefficients on both HighStake and Attribute Conflict are positive and statistically significant. Randomization therefore increases when subjects are drawn to the salient high-payoff attribute and especially when there is a structural conflict between expected value and extreme-outcome probabilities. This supports the view that stochastic choice reflects attention shifts across competing attributes rather than generic task difficulty or cognitive noise. In environments without a clearly dominant option, agents appear to use deliberate mixing as a low-cost strategy to explore the high-stake alternative. These findings support the view that stochastic choice reflects attention shifts across competing attributes rather than generic task difficulty, cognitive noise, or strenght of preference alone.

This perspective aligns naturally with an attribute-salience mechanism: when salient primitives favor different options, small shifts in which attribute is emphasized can generate systematic switching even in repeated, transparent environments.

4 Conclusion

This paper revisits the conceptual foundations of stochastic choice by providing a unified replication and extension of two influential recent approaches with a structurally richer lottery environment and by advancing a new attribute–salience account of deliberate randomization. First, Agranov and Ortoleva (2017) show that stochastic choice can reflect deliberate randomization rather than mere error, using repeated-choice designs to distinguish purposeful mixing from non-systematic noise. Second, Alós-Ferrer et al. (2025) emphasize a complementary mechanism rooted in random utility: inconsistency arises from noise around underlying preferences and increases as preference differences shrink — when alternatives approach indifference — thus interpreting stochastic choice as a graded function of preference strength rather than deliberate randomization.

Our experimental design combines repeated exposure to the same lottery pairs with a two-phase structure that varies the informational and attentional environment. In Phase 1, repetitions are interleaved, approximating the conditions under which memory and attention frictions may interact with evaluation. In Phase 2, repetitions are consecutive, reducing such frictions and allowing us to identify a deliberate component of mixing that persists even when choices are temporally clustered. This structure lets us distinguish mechanical variation from systematic switching and, in turn, study both the incidence and the determinants of randomization.

Consistent with Agranov and Ortoleva (2017), we find clear evidence of deliberate randomization: switching is not confined to obvious mistakes or dominated options and remains present when repetitions are made salient by consecutive presentation. Moreover, mixing is far from idiosyncratic. It follows systematic patterns that map tightly into the attribute structure of the lotteries. In particular, randomization responds strongly to the probability mass on extreme outcomes and to payoff dispersion, and it is closely associated with a propensity to select the high-stake option at least once within a sequence. These findings support the interpretation of stochastic choice as purposeful mixing that is guided by salient attributes rather than as residual cognitive noise.

In parallel, we replicate the qualitative prediction of Alós-Ferrer et al. (2025) that inconsistency varies with preference strength, while clarifying its scope in risky environments. Graphical and regression evidence indicate that inconsistency is lowest for first-order dominance pairs and tends to rise when expected-utility differences are small. At the same time, once dominance structure and other lottery characteristics are accounted for, proximity in expected utility alone is not sufficient to explain the main switching patterns. This suggests that strength of preference is part of the story, but that it operates through a deeper mechanism: preferences become behaviorally weak

precisely in regions where salient attributes conflict (Busemeyer and Townsend, 1993; Woodford, 2014).

The central extension of the paper is to provide this microfoundation in terms of attribute salience and attribute conflict. We show that stochastic choice emerges most strongly when the key dimensions that guide evaluation - probability of zero, probability of the maximum payoff, and expected value - pull in different directions. In these environments, attention does not settle on a single decisive attribute. Instead, subjects intermittently shift toward the option with the most salient upside and may "try once" for the high-stake outcome, even when expected value favors the alternative. Formally, mixing increases precisely when the high-payoff attribute (or the zero payoff one) is sufficiently salient to offset the expected-value comparison, and when the structure of the pair creates a genuine conflict between upside salience and downside risk (Bordalo et al., 2012b; Gennaioli and Shleifer, 2010). This mechanism is reinforced by comparisons of structurally isomorphic lottery pairs in the Marschak–Machina representation: even when probability structure is preserved, shifts that remove the zero outcome and amplify the high payoff can increase conflict and raise stochasticity, illustrating why a purely strength-of-preference account is incomplete (Butler and Loomes, 2007).

Taken together, our results suggest a reinterpretation of stochastic choice in risky environments. Randomization is not well captured as a residual error term, nor as a purely reduced-form consequence of small utility differences. Rather, it reflects systematic, salience-driven reweighting of attributes: when no single attribute dominates, individuals alternate between attribute-based comparisons and deliberately mix across repetitions. This perspective offers a tractable bridge between deliberate randomization, salience models, and empirical patterns of switching (Gul et al., 2014; Woodford, 2014). It also yields concrete predictions about when mixing should intensify - namely, in choice problems where extreme outcomes are salient and the direction of evaluation depends on which attribute is attended to.

The paper has limitations that open natural extensions. First, our inference about salience is behavioral and structural; future work could incorporate direct attention measures (eye-tracking or mouse-tracking) to test whether shifts in attention precede switching (Camerer et al., 2005; Coricelli et al., 2005; Fehr and Rangel, 2011). Second, while the two-phase design helps isolate deliberate mixing, it does not fully disentangle learning from attention dynamics; designs with explicit feedback or controlled information acquisition could separate these channels. Third, our framework is intentionally parsimonious; richer structural models could endogenize attribute weights and estimate individual-level salience parameters jointly with risk preferences (Busemeyer and Townsend, 1993). Finally, the implications of deliberate randomization for welfare and policy remain largely unexplored. In domains such as consumer finance, insurance, or health decisions - where extreme outcomes are both salient and consequential - understanding when and why agents mix could inform better choice architectures and mechanism design.

Overall, by replicating recent evidence on deliberate randomization and strength-of-preference effects, and by grounding these patterns in an attribute-salience mechanism, the paper provides a coherent theoretical foundation for stochastic choice in risky environments and a platform for future research on strategic mixing under conflicting evaluative cues.

References

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

Alós-Ferrer, C. and Buckenmaier, J. (2021). A theory of error-prone decision making. *Journal of Economic Theory*, 195:105263.

Alós-Ferrer, C., Buckenmaier, J., Garagnani, M., and Ritschel, A. (2025). Why is choice stochastic? deliberate randomization vs. strength of preference. *Management Science*. Forthcoming; version dated November 16, 2025.

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. (2012a). Salience theory of choice under risk. *The Quarterly Journal of Economics*, 127(3):1243–1285.

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

Bordalo, P., Gennaioli, N., and Shleifer, A. (2013). Salience and consumer choice. *Journal of Political Economy*, 121(5):803–843.

Busemeyer, J. R. and Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, 100(3):432–459.

Butler, D. J. and Loomes, G. (2007). Imprecision as an account of the preference reversal phenomenon. *American Economic Review*, 97(1):277–297.

Camerer, C., Loewenstein, G., and Prelec, D. (2005). Neuroeconomics: How neuroscience can inform economics. *Journal of Economic Literature*, 43(1):9–64.

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.

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

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

Chew, S. H., Miao, B., Shen, Q., and Zhong, S. (2022). Multiple-switching behavior in choice-list elicitation of risk preference. *Journal of Economic Theory*, 204:105510.

Conte, A., Hey, J. D., and Moffatt, P. G. (2011). Mixture models of choice under risk. *Journal of Econometrics*, 162(1):79–88.

Coricelli, G., Critchley, H. D., Joffily, M., O'Doherty, J. P., Sirigu, A., and Dolan, R. J. (2005). Regret and its avoidance: A neuroimaging study of choice behavior. *Nature Neuroscience*, 8(9):1255–1262.

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

Davidson, D. and Marschak, J. (1959). Experimental tests of a stochastic decision theory. In Churchman, C. W. and Ratoosh, P., editors, *Measurement: Definitions and Theories*, pages 233–269. Wiley, New York.

Fehr, E. and Rangel, A. (2011). Neuroeconomic foundations of economic choice — recent advances. *Journal of Economic Perspectives*, 25(4):3–30.

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

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

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

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.

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.

Tversky, A. (1969). Intransitivity of preferences. *Psychological Review*, 76:31–48.

Woodford, M. (2014). An optimizing neuroeconomic model of discrete choice. *NBER Working Paper*, (19897).

5 Acknowledgements

We also gratefully acknowledge the Italian Ministry of Universities and Research under Grant PRIN 2022 n. 20229LRAHK (funded by the European Union - Next Generation EU, Mission 4 Component 1, CUP B53D23012680006)

6 Disclosure Statement

The author reports there are no competing interests to declare.

7 Data Availability Statement

The data that support the findings of this study are openly available in "figshare" at https://doi.org/10.6084/m9.figshare.32040945

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)			

Idx	Side	Payoffs	Probs	Mean RT	% choosing 1	Class
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)			
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 7: 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 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.

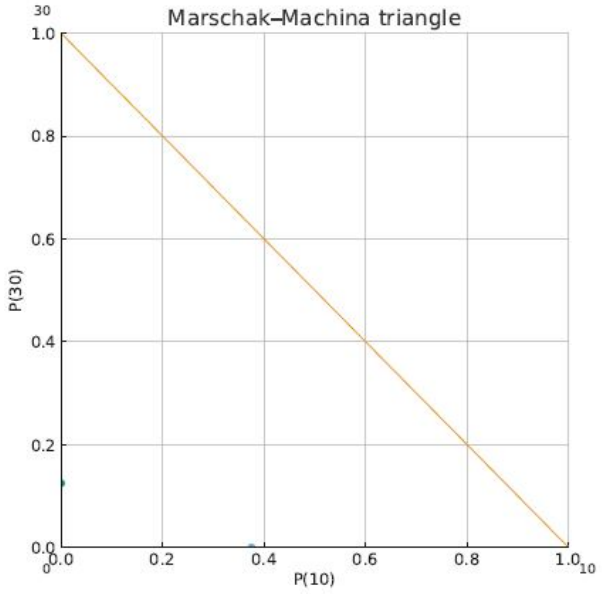
We do not report MM representations for first-order stochastically dominant (FOSD) lotteries, since in those cases the ranking is trivial and no meaningful randomization structure can be inferred beyond simple errors.

Lottery index 1 (easy)

	Phase 1	Phase 2	Total
vs reference	29.0%	13.4%	22.8%
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%

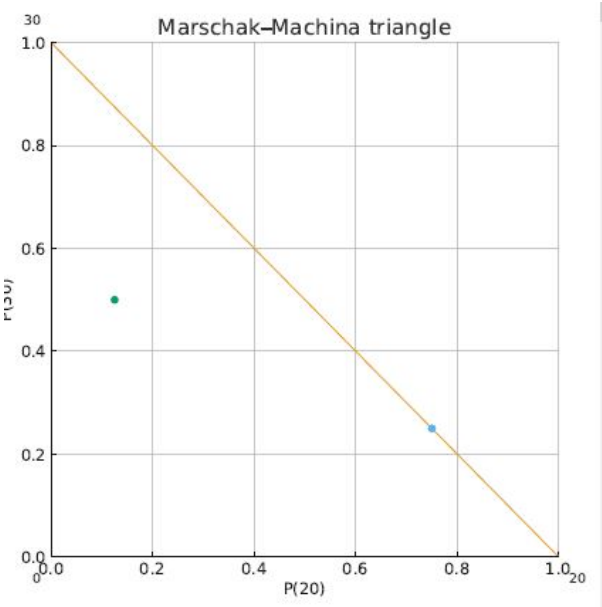


Lottery index 2 (easy)

	Phase 1	Phase 2	Total
vs reference	27.8%	17.1%	23.5%
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%

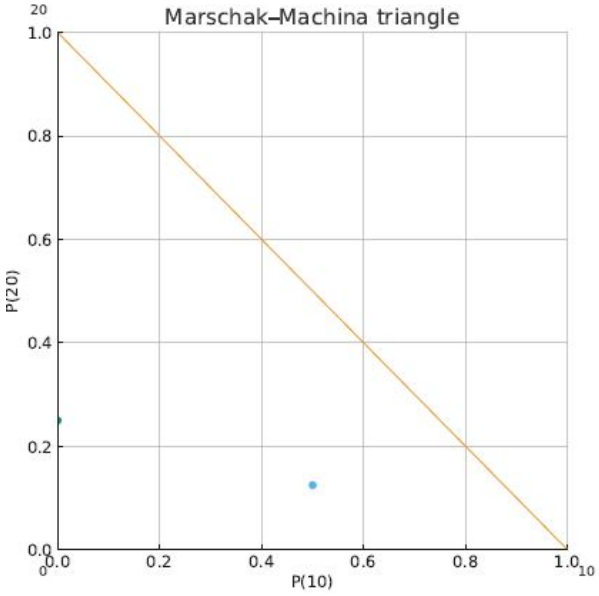


Lottery index 3 (easy)

	Phase 1	Phase 2	Total
vs reference	7.4%	6.4%	7.0%
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%

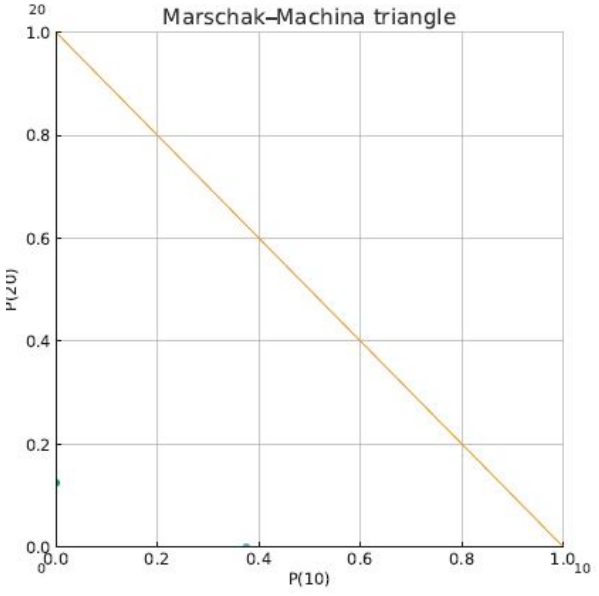


Lottery index 4 (hard)

	Phase 1	Phase 2	Total
vs reference	24.4%	12.5%	19.6%
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%

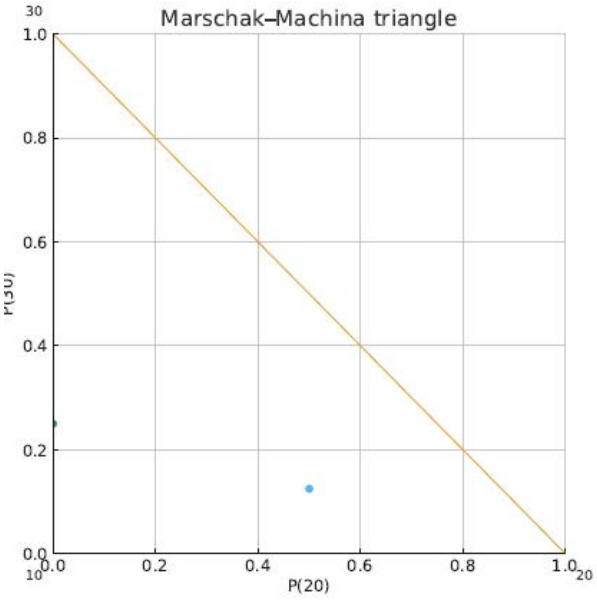


Lottery index 5 (hard)

	Phase 1	Phase 2	Total
vs reference	26.2%	18.0%	22.9%
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%

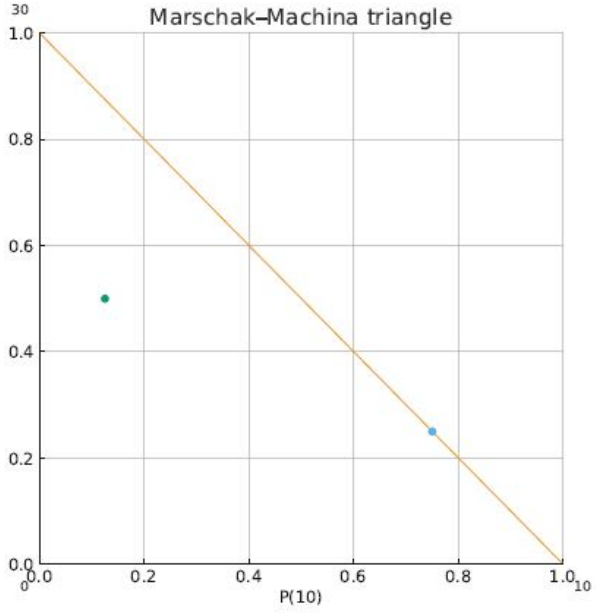


Lottery index 6 (hard)

	Phase 1	Phase 2	Total
vs reference	29.9%	17.6%	25.0%
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%	16.7%	18.1%
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%

