

Discussion Paper Series – CRC TR 224

Discussion Paper No. 687

Project A 05

# Learning to Maximize Ordinal and Expected Utility, and the Indifference Hypothesis

Thomas Dohmen<sup>1</sup>  
Georgios Gerasimou<sup>2</sup>

April 2025

<sup>1</sup>[tdohmen@uni-bonn.de](mailto:tdohmen@uni-bonn.de)  
<sup>2</sup>[georgios.gerasimou@glasgow.ac.uk](mailto:georgios.gerasimou@glasgow.ac.uk)

Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)  
through CRC TR 224 is gratefully acknowledged.

# Learning to Maximize Ordinal and Expected Utility, and the Indifference Hypothesis\*

Thomas Dohmen<sup>†</sup>  
University of Bonn

Georgios Gerasimou<sup>‡</sup>  
University of Glasgow

April 13, 2025

## Abstract

We ask if participants in a choice experiment with repeated presentation of the same menus and no feedback provision: (i) learn to behave in ways that are closer to the predictions of ordinal and expected utility theory under *strict* preferences; or (ii) exhibit overall behaviour that is consistent with utility theory under *weak* preferences. To answer these questions we designed and implemented a free-choice lab experiment with 15 distinct menus. Each menu contained two, three and four lotteries with three monetary outcomes, and was shown five times. Subjects were not forced to make an active choice at any menu but could avoid/defer doing so at a positive expected cost. Among our 308 subjects from the UK and Germany, significantly more were ordinal- and expected-utility maximizers in their last 15 than in their first 15 identical decision problems. Around a quarter and a fifth of all subjects, respectively, decided in those modes *throughout* the experiment, with nearly half revealing non-trivial indifferences. A considerable overlap is found between those consistently rational individuals and the ones who satisfied core principles of *random* utility theory. Finally, choice consistency is positively correlated with cognitive ability, while subjects who learned to maximize utility were more cognitively able than those who did not. We discuss potential implications of our study's novel set of findings.

Keywords:

Ordinal utility; expected utility; learning; indifference; avoidance/deferral; cognitive ability.

---

\*We thank Carlos Alós-Ferrer, Junaid Arshad, Andrew Caplin, Ido Erev, John List, Antonio Penta, Jakub Steiner, Jianying Qiu and audiences at Frontiers of Measurement & Formation of Skills (Jinan University), ESA 2024, Psychonomics 2024, EEA-ESEM 2024, BSE Bounded Rationality 2024, SAET 2023, Nottingham, Radboud/Nijmegen, Glasgow and St Andrews for useful comments. Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 (Project A05) and Germany's Excellence Strategy - EXC 2126/1-390838866 as well as through NSF Award 2201888 is gratefully acknowledged.

<sup>†</sup>tdohmen@uni-bonn.de

<sup>‡</sup>georgios.gerasimou@glasgow.ac.uk

# 1 Introduction

A large body of experimental work in economics and psychology shows that individuals often make different choices under risk when confronted with the same decision problems repeatedly. Such patterns are often interpreted as evidence against the hypothesis of stable, complete and transitive preferences that constitutes the rationality cornerstone of neoclassical economic analysis. Six commonly invoked explanations for the occurrence of such patterns include the possibilities of: (i) “noisy” utility maximization; (ii) systematically bounded-rational behaviour with stable or context-dependent preferences; (iii) imprecise/incomplete preferences; (iv) utility maximization with costly information acquisition; (v) limited attention; and (vi) deliberate randomization.<sup>1</sup> In this study we focus on the potential role of two basic complementary explanations by asking the following questions that have received less attention in the literature:

1. Do subjects *learn* to be more rational over the course of the experiment *without* receiving any feedback, interventions or other opportunities to acquire new information, and without being forced to always make an active choice at each menu?

This question is important for at least two reasons. First, if learning does occur in experiments featuring repeated presentation of the same decision problems without subjects receiving any new information, and in a free-choice environment where avoiding/deferring is also possible, then the targeted design and use of such experiments should be promoted further for more accurate theoretical tests and preference recovery. Second, under this learning hypothesis an analyst could be justified to focus on subjects’ behaviour at the later stages of the experiment and analyse it through the lens of expected utility. With the exception of a few studies that we discuss in Section 8, which were considerably more limited in their scope and sample sizes and did not allow subjects to avoid/defer making an active choice from the presented menus, to our knowledge there has been no systematic attempt to answer this question.

2. Can some of the observed ‘volatility’ in behaviour across different presentations of the same menus be due to subjects’ *rational indifferences* between the relevant alternatives?

This intuitive possibility is in line with conventional economic interpretations of the concept of indifference. In one of the early studies that presented results from an experiment on stochastic choice, for example, Davidson and Marschak (1959, p. 233) remarked that “*One may interpret every case of inconsistency as a case of indifference: if the subject has chosen a rather than b but soon afterwards chooses b rather than a, this is interpreted as indifference between those two objects; if he chooses a rather than b, b rather than c, and c rather than a,*

---

<sup>1</sup>See, for example, (i) Gillen et al. (2019); Apestequia and Ballester (2018); Harless and Camerer (1994); Hey and Orme (1994); (ii) Cerreia-Vioglio et al. (2015); Bordalo et al. (2012); (iii) Cubitt et al. (2015); (iv) Dean and Neligh (2023); (v) Barseghyan et al. (2021); Barseghyan and Molinari (2023); (vi) Machina (1985); Cerreia-Vioglio et al. (2019); Agranov and Ortoleva (2017); and references therein.

*this is interpreted as indifference between those three objects.*” The authors concluded with their warning that *“In empirical application, this approach would probably make indifference all-pervasive.”* Here too, however, it appears that no systematic attempt has been made to examine whether rational choice with *weak* preferences can account for some of the observed—and seemingly non-rational—choice reversals. Indeed, it is unclear to what extent Davidson and Marschak’s (1959) warning is justified once subjects’ indifferences and strict preferences are recovered from such data in a theory-guided way that uses modern computational tools.

To answer these questions we designed a targeted lab experiment, which we implemented in-person in the UK and in Germany. In addition to testing deterministic ordinal- and expected-utility maximization—with or without indifferences—as well as *random* (expected) utility maximization, our design aimed at a systematic test of the learning hypothesis from many behavioural angles and at a relatively high level of generality.

The decision environment in our experiment was structured around seven lotteries and fifteen menus derived from them. Each lottery assigned a positive probability to three monetary prizes. The fifteen menus included nine binary, four ternary and two quaternary ones. In particular, contrary to the vast majority of experimental studies in choice under risk which focus only on behaviour at binary menus and implicitly assume subjects’ compliance with ordinal utility maximization at *all* menus, this design allows testing various implications of both ordinal and expected utility theory *simultaneously*. Specifically, we aimed to assess, among others, the Transitivity, Contraction Consistency and Weak Axiom of Revealed Preference implications of ordinal utility theory, as well as the Independence, First-Order Stochastic Dominance and *Stability of Attitudes to Risk* (StAR) implications of expected utility theory. The latter implication, which we make precise, is to our knowledge new. We test it by introducing menus that feature a Second-Order Stochastic Dominance relation between the available lotteries. Risk-attitude stability in our environment reduces to the requirement that subjects always choose either the dominant (risk-averse) or dominated (risk-seeking) lottery. Additionally, our test of the Independence axiom is also new in that, although invoking Allais-style (Allais, 1953; Kahneman and Tversky, 1979) lottery mixtures, neither the prizes nor the probabilities are as extreme as in the original Allais paradox or in its “common-ratio”, “common-consequence” and “certainty-effect” variations and decompositions that have been studied extensively in the literature.<sup>2</sup> Instead, in each of the two binary menus that were designed to test Independence in our study the two lotteries had the same expected value and were unrelated by Second-Order Stochastic Dominance—hence presented relatively more difficult decisions. In addition, the expected value in these two pairs differed by a factor of two (12 and 6 pounds sterling/euros, respectively). Thus, our test of Independence differs from the standard ones in that it aims to shed light into humans’ conformity with this axiom when the two underlying decisions are not as intuitive as in those situations where it is well-known that the axiom is often violated.

---

<sup>2</sup>See McGranaghan et al. (2024) and references therein.

Importantly, moreover, incorporating a non-forced/free-choice approach into our experimental design allows us to also investigate the Decisiveness implication of rational choice theory, according to which the decision maker always chooses one of the available lotteries, in line with the Completeness/Comparability axiom of utility theory. Implicitly, Decisiveness also assumes that decision makers are not subject to fatigue or cognitive overload at any decision problem. An essential aspect of our design was granting subjects the freedom to avoid or delay, at a small expected cost, making an active choice at menus where, for any reason, they felt uncomfortable to do so. This approach allows us to test the above-mentioned other six implications of the theory while eliminating the confounds arising from the interactions between standard forced-choice experimental designs and the potential incompleteness of subject’s preferences or reluctance to engage with the decision problem at hand.<sup>3</sup>

Deploying new computational tools on the data collected from 308 subjects, we find that:

1. Nearly twice as many subjects conformed with ordinal- (57.5%) and expected-utility (40%) maximization with strict preferences at the end of the experiment than at the beginning, accompanied by significant reductions in decision times.
2. About a quarter and a fifth of all subjects consistently exhibited rational behaviour in these two respects *throughout* their 75 decisions, with about half of them revealing at least one indifference between distinct lotteries.
3. There are substantial overlaps between subjects who consistently behaved as ordinal or expected-utility maximizers and those who were potentially *random*-(expected-)utility maximizers across the five rounds.
4. The number of subjects violating either one or all of Transitivity, Contraction Consistency, Weak Axiom of Revealed Preference, Independence, First-Order Stochastic Dominance and Stability of Attitudes to Risk—a new condition that we introduce and test—decreases steadily over the course of the experiment, and significantly so between the first and last round.
5. Deferring/avoiding behaviour is generally infrequent and stable throughout, mainly occurring at menus with increased decision difficulty, i.e., those without a stochastically dominant lottery and/or where the feasible lotteries are relatively complicated.

---

<sup>3</sup>As far as testing Transitivity is concerned, early warnings to that effect appeared in Luce and Raiffa (1957) and, citing these authors, Aumann (1962). Motivated by these and also by experimental findings in psychology suggesting that hard decisions lead to choice paralysis (Tversky and Shafir, 1992; Iyengar and Lepper, 2000; Dhar, 1997; Dhar and Simonson, 2003), Gerasimou (2018) proposed models of fully consistent active choices in general non-forced decision environments. Some predictions of one of these models were subsequently tested experimentally in Costa-Gomes et al. (2022), Gerasimou (2021) and, less directly, Nielsen and Rigotti (2022). These three studies share the finding that free/non-forced choices are significantly more consistent than forced-choice ones, in line with Luce and Raiffa’s (1957) intuition that “*intransitivities often occur when a subject forces choices between inherently incomparable alternatives*”. Our study does not feature a forced-choice treatment. Instead, it allows testing, for the first time, a rich set of specific implications of expected utility theory without forcing subjects to always make active choices, thereby extending the crux of Luce and Raiffa’s insights to that domain also.

6. Cognitive ability, particularly in verbal reasoning and letter-number sequence tasks, is positively correlated with choice consistency and with “early-onset” rationality.
7. Subjects who deviated from rationality initially but learned to be rational by the end of the experiment were significantly more cognitively able than those who did not.

These findings have implications for testing existing theories of choice under risk and interpreting the results of such tests, as well as for experimental design, preference elicitation, the development of new theories, and for efforts to improve real-world decision-making under risk. We discuss these implications in the last section. Overall, we view this paper’s contribution as documenting and discussing the relevance of these novel empirical findings as the outcome of an analysis that features the following methodological innovations which, we hope, add meaningfully to the currently available toolkit in the literatures of choice under risk, revealed preference analysis, behavioural economics, as well as the emerging field of cognitive economics (Caplin, 2025):

- i. The design of a free-choice experiment that combines repetition of the distinct decision problems, with the non-provision of any kind of feedback and the possibility for subjects to avoid/defer making an active choice at any menu at a small expected cost.
- ii. A collection of 15 binary and non-binary menus from 7 lotteries with 3 outcomes, whose construction was theory-motivated and guided by the presence or absence of first/second stochastic dominance relations, allowing for targeted systematic tests of both ordinal- and expected-utility maximization in a richer environment than those typically seen in experiments on choice under risk. To our knowledge, this is the first such design that presents repeated choices from both binary *and* non-binary menus of lotteries, whether in a forced or in a non-forced choice environment. This is particularly important because it provides the basis for a very general test of these theories and allows assessing in a rather detailed way the “extra rationality steps” that separate conformity with ordinal vs expected-utility maximization.
- iii. A carefully administered presentation of the 5 repetition blocks for each subject, which combines homogeneity within and across subjects in the first and last rounds with complete randomness—also within and across subjects—in the middle three rounds, motivated by the specific goal of testing whether feedback-less learning occurs.
- iv. The analyses of (a) deterministic- as well as random-utility theories, and their juxtaposition and explanatory overlap; and (b) deterministic ordinal and expected-utility maximization through the lens of both strict and, new with this paper, *weak* preferences, and the recovery—by means of recently developed combinatorial-optimization techniques—of the best-matching such relations for each subject.

The next section introduces the relevant theoretical framework. The one following it provides a detailed description of the experimental design and its relation to the preceding

theoretical framework. Our empirical results are presented in Section 4. The penultimate section discusses the related literature and how this paper is placed in it. The final section discusses implications of our findings. Additional information about the experiment, including screenshots, is provided in Online Appendices A–E.

## 2 Theoretical Background

The data collected in our experiment can be analyzed from both deterministic and stochastic choice perspectives. In the former case, they can be additionally analysed from the point of view of (expected-)utility maximizing single-valued choices with strict risk preferences, either per individual decision round or overall, as well as from that of *multi-valued* choices with rational and weak risk preferences. Conversely, in the case of stochastic choice, they can be tested for potential conformity with *random (expected) utility* theory (Thurstone, 1927; Block and Marschak, 1960; Falmagne, 2002; Gul and Pesendorfer, 2006). Each of these two broad theoretical modes of analysis has well-known behavioural implications, many of which our experiment was specifically designed to test. We proceed with reviewing those next.

### 2.1 Deterministic Choice

We consider a general choice domain of finitely many menus containing lotteries that are defined over a finite set of monetary outcomes  $Z \subset \mathbb{R}_+$ . Denoting by  $X$  the finite set of lotteries over  $Z$  that we consider, and denoting by  $\mathcal{M}$  the collection of such menus, a decision maker’s behaviour is described by a choice correspondence  $C : \mathcal{M} \rightarrow X$ , i.e., a mapping that satisfies  $C(A) \subseteq A$  for every menu  $A$  in  $\mathcal{M}$ . Unlike a single-valued choice *function*, the value of a choice correspondence at menu  $A$  could contain multiple alternatives, which are typically interpreted as those that the decision maker *might* choose from  $A$ .<sup>4</sup> Clearly, since the empty set is a subset of every set, this basic definition allows for the possibility of  $C(A) = \emptyset$ . Because the ultimate decision outcomes in our non-forced choice experimental design are monetary amounts, with higher clearly preferred to lower, and with all non-zero amounts being desirable, the potential unattractiveness of the items in  $A$  is not a likely explanation for observing  $C(A) = \emptyset$  in our setting. Hence, this notation in our environment will be used when thinking about the decision maker opting to avoid/delay choice at  $A$  because they find it difficult to make an active choice at  $A$ .

The four basic choice axioms on the observable values of  $C$  that we list below are implied by every deterministic model of utility maximization over lotteries over  $X$  and not merely by those that belong to the expected-utility class of von Neumann and Morgenstern (1947).

#### Decisiveness

$C(A) \neq \emptyset$  for every menu  $A$ .

---

<sup>4</sup>See, for example, Chapter 1 in Mas-Colell et al. (1995) or Chapter 1 in Kreps (2012).

**Transitivity**

If  $p \in C(\{p, q\})$  and  $q \in C(\{q, r\})$ , then  $p \in C(\{p, r\})$ .

**Contraction Consistency / Independence of Irrelevant Alternatives**

If  $p \in C(A)$  and  $p \in B \subset A$ , then  $p \in C(B)$ .

**Weak Axiom of Revealed Preference (WARP)**

If  $p \in C(A)$ ,  $q \in A \setminus C(A)$  and  $q \in C(B)$ , then  $p \notin B$ .

*Decisiveness* is typically assumed in most choice-theoretic analyses as part of the definition of a choice correspondence. As was pointed out above, however, it is in fact an additional restriction that has behavioural meaning. Relaxing Decisiveness when the analyst suspects that decision makers may avoid/delay making an active choice because of decision difficulty is potentially fruitful theoretically (Hurwicz, 1986; Kreps, 1990, 2012; Gerasimou, 2018) and relevant empirically (Tversky and Shafir, 1992; Iyengar and Lepper, 2000; Dhar, 1997; Costa-Gomes et al., 2022). We stress that, because our environment is one of non-forced choices where Decisiveness is not a priori assumed to hold, in addition to ruling out *cyclic* preferences, Transitivity here also rules out acyclic but nevertheless still non-transitive preferences. For example,  $x \in C(\{x, y\})$ ,  $y \in C(\{y, z\})$ ,  $\emptyset = C(\{x, z\})$  reveal acyclic but intransitive and incomplete preferences where  $x \succsim y \succsim z$  and  $x \not\prec z \not\prec x$ . Hence, testing Transitivity with choices that have arisen from such a non-forced choice decision environment amounts to testing for transitive preferences in the absence of any potential confounds that the exogenous imposition of Decisiveness may not be able to account for (Luce and Raiffa, 1957; Aumann, 1962).

WARP is a fundamental rationality property. It requires that there be no direct choice reversals between any two lotteries. *Contraction Consistency* (Sen, 1997), also known as *Independence of Irrelevant Alternatives*, the *Chernoff axiom* (Chernoff, 1954), and *Property  $\alpha$*  (Sen, 1971), is implied by WARP under Decisiveness but not in general; for example,  $x \in B \subset A$ ,  $x \in C(A)$ ,  $C(B) = \emptyset$  violates this axiom but satisfies WARP. In the baseline case where  $C(\cdot)$  is always non-empty-valued, this axiom rules out a large class of context-dependent choice reversals that are driven by the presence or absence of irrelevant alternatives. Specifically, it requires that when an alternative is declared choosable at a menu, then removing other alternatives from that menu should not alter this status. That is, the absence of those “irrelevant” alternatives at the smaller menu should not make the agent choose something else or, in our more general environment, avoid/delay choice.

The next three revealed-preference axioms are relevant for finite choice datasets that are either obtained from general environments of choice under risk (cf *Independence*) or from more specific environments of choice over money lotteries (cf *FOSD & StAR*). All alterna-



tives  $p, q, r$  in the statements below are assumed to be of this kind. Anticipating one of these statements, we say that money lotteries  $p$  and  $q$  defined over a finite set  $Z$  have an *overlapping range* if the intervals  $[p_l, p_h]$  and  $[q_l, q_h]$  that are formed by the lowest and highest prizes –subscripted by  $l$  and  $h$ – in the supports of  $p$  and  $q$ , respectively, have a non-degenerate (i.e., non-empty and non-singleton) intersection.

## Independence

For any  $p, q, r$  and  $\alpha \in (0, 1)$ ,

$$p \in C(\{p, q\}) \implies \alpha p + (1 - \alpha)r \in C(\{\alpha p + (1 - \alpha)r, \alpha q + (1 - \alpha)r\}).$$

## First-Order Stochastic Dominance (FOSD)<sup>5</sup>

If  $p$  FOSD  $q$ , then  $C(\{p, q\}) = \{p\}$ .

## Stable Attitudes to Risk (StAR)

If  $p_1, p_2, q_1, q_2$  have an overlapping range and  $p_1$  SOSD<sup>6</sup>  $p_2$ ,  $q_1$  SOSD  $q_2$ , then

$$p_1 \in C(\{p_1, p_2\}) \iff q_1 \in C(\{q_1, q_2\}).$$

While *FOSD* and *Independence* are well-known and extensively studied implications of expected utility theory, *StAR* is less commonly encountered –in fact, we were unable to find its statement in the literature– and hence warrants some discussion. First, as is evident from its statement, the axiom requires that the agent always reveal a weak or strict preference for the same type of lottery across all pairs whose elements are ranked by second-order stochastic dominance and have overlapping ranges. That such behaviour is implied by the expected-utility model is an easy implication of well-known results, as we confirm next.

**Proposition 1.** *Every finite choice dataset that is compatible with expected-utility maximization satisfies StAR.*

*Proof.* Assume to the contrary that, for an EUM agent with a strictly increasing  $u : \mathbb{R} \rightarrow \mathbb{R}$ , we have  $p_1 \in C(\{p_1, p_2\})$  and  $q_1 \notin C(\{q_1, q_2\})$ , where  $p_1, p_2, q_1, q_2$  have an overlapping range and  $p_1, q_1$  SOSD  $p_2, q_2$ , respectively. By the EUM hypothesis,  $\emptyset \neq C(\{q_1, q_2\}) = \{q_2\}$  and

$$\sum_{x \in X} p_1(x)u(x) \geq \sum_{x \in X} p_2(x)u(x), \tag{1}$$

$$\sum_{x \in X} q_2(x)u(x) > \sum_{x \in X} q_1(x)u(x) \tag{2}$$

By the SOSD assumption and Theorems 3 and 4 in Hadar and Russell (1969),<sup>7</sup> (1) implies

<sup>5</sup>For money lotteries  $p$  and  $q$  that are defined over a finite set  $X$  and have cumulative distributions denoted by  $F_p$  and  $F_q$ ,  $p$  is said to FOSD  $q$  if  $F_p(x) \leq F_q(x)$  for all  $x \in X$ , with strict inequality at some  $x$ .


<sup>6</sup>For money lotteries  $p$  and  $q$  that are defined over a finite set  $X$  and have cumulative density functions on any interval  $[a, b]$  that contains  $X$  denoted by  $F_p$  and  $F_q$ ,  $p$  is said to second-order stochastically dominate (SOSD)  $q$  if  $\int_a^x [F_p(t) - F_q(t)]dt \leq 0$  for all  $x \in [a, b]$ , with strict inequality at some  $x$ .

<sup>7</sup>See also Hanoch and Levy (1969) and Rothschild and Stiglitz (1970).

that  $u$  is weakly concave in  $[p_l, p_h]$  while (2) implies that  $u$  is strictly convex in  $[q_l, q_h]$ . Since, by assumption,  $[p_l, p_h] \cap [q_l, q_h]$  is a non-degenerate interval, this is a contradiction.  $\square$

We note that the common-support part in the antecedent of StAR’s statement is essential. Indeed, it is an intuitive and well-known fact that safer (*SOS*Dominant) and riskier (*SOS*Dominated) lotteries might be preferred by the same expected-utility maximizer at low and high wealth levels, respectively (Friedman and Savage, 1948). With this proviso in place, StAR merely requires that the agent’s general attitude toward risk, as measured by a preference for *SOS*Dominant, *SOS*Dominated or either type of lotteries for a risk-averse, risk-seeking and risk-neutral agent, respectively, remain consistent over any fixed range of wealth levels (see Figure 1).

Figure 1: Example (pseudo-)utility functions that are (in)compatible with StAR.



To conclude this section, we note that the remarks made earlier about the generality of testing Transitivity in our non-forced-choice environment also carry over to Independence, FOSD and StAR. Namely, by not imposing forced-choice ex ante, we are allowing for testing each of these axioms independently of Completeness.

## 2.2 Stochastic Choice

Letting  $Z$ ,  $X$  and  $\mathcal{M}$  carry the same meaning as before, we now consider a general random non-forced choice environment. Specifically, a *random choice model* in our framework is a mapping  $\mathbb{P} : \mathcal{M} \times \mathcal{M} \rightarrow \mathbb{R}_+$  such that  $\mathbb{P}(\{p\}, A) \in [0, 1]$  for all  $A \in \mathcal{M}$  and all  $p \in A$ ;  $\mathbb{P}(\{p\}, A) = 0$  for all  $A \in \mathcal{M}$  and all  $p \notin A$ ; and  $\sum_{p \in A} \mathbb{P}(\{p\}, A) \leq 1$ , where  $0 \leq 1 - \sum_{p \in A} \mathbb{P}(\{p\}, A) \leq 1$  is the probability of avoiding/delaying choice at menu  $A$ . By  $\mathbb{P}(A, A)$ , finally, we denote the probability of choosing any lottery  $p$  at menu  $A$ . Clearly, in view of the above,  $\mathbb{P}(A, A) \leq 1$ .

The baseline model of stochastic rationality in the choice domain that is relevant in our experiment is *random (expected) utility* (Thurstone, 1927; Block and Marschak, 1960;

Falmagne, 2002; Gul and Pesendorfer, 2006). This posits the existence of a probability measure  $\mu$  over the set  $\mathcal{P}$  of all strict total orders over the lotteries in  $X$  such that, for every menu  $A$  and lottery  $p \in A$ ,  $\mathbb{P}(\{p\}, A) = \mu(\{\succ \in \mathcal{P} : p \succ p' \text{ for all } p' \in A \setminus \{p\}\})$ . That is, the model’s probability of choosing  $p$  at  $A$  coincides with the  $\mu$ -probability of  $p$  being the most preferred lottery at  $A$  under some strict preference ordering over  $X$ .

We state next five principles of stochastic-choice consistency that we take interest in:

### Regularity

If  $p \in B \subset A$ , then  $\mathbb{P}(\{p\}, A) \leq Pr(\{p\}, B)$ .

### Stochastic Transitivity

If  $\mathbb{P}(p, \{p, q\}) \geq 0.5$  and  $\mathbb{P}(q, \{q, r\}) \geq 0.5$ , then

$$\begin{aligned} \mathbb{P}(p, \{p, r\}) &\geq 0.5 && \textbf{(Weak)} \\ \mathbb{P}(p, \{p, r\}) &\geq \min\{\mathbb{P}(p, \{p, q\}), \mathbb{P}(q, \{q, r\})\} && \textbf{(Moderate)} \\ \mathbb{P}(p, \{p, r\}) &\geq \max\{\mathbb{P}(p, \{p, q\}), \mathbb{P}(q, \{q, r\})\} && \textbf{(Strong)} \end{aligned}$$

### Stochastic Decisiveness

$\mathbb{P}(A, A) = 1$  for all  $A$ .

The first (Block and Marschak, 1960) and fifth axioms on this list are stochastic-choice analogues of the Contraction Consistency and Decisiveness axioms of deterministic choice, respectively, and are implied by Random Utility models. Weak, Moderate and Strong Stochastic Transitivity (Marschak, 1960; He and Natenzon, 2024) on the other hand are logically nested stochastic-choice variants of Transitivity that are generally distinct from Random Utility models (Strzalecki, 2025).

## 3 Design of the Experiment

### 3.1 Lotteries and Choice Menus

We constructed 7 lotteries, each with three monetary outcomes from the set  $\{0, 9, 10, 20, 24\}$ , where the numbers denote Pounds Sterling, £, and Euros, € (Table 1; Figure B.1). Out of the 127 possible menus that are derivable from this grand choice set we selected 15 that contained either two lotteries (9 menus), three (4 menus) or four (2 menus) (Figure B.2). All menus were presented 5 times, resulting in a total of 75 decision problems. This decision environment is therefore richer than those usually seen in experiments on choice under risk where subjects are presented only with binary menus, often over two monetary outcomes.<sup>8</sup> This domain richness is relevant because it allows testing an analogously rich set of consistency principles

<sup>8</sup>That said, we note that Apestequia and Ballester (2021) also analysed an experimental dataset on choice under risk where some non-binary menus were shown. Unlike our experiment where each menu was presented five times, subjects in that experiment saw each menu once.

that pertain both to binary and larger menus.

In each decision, subjects could decide to defer the choice by choosing the option “*I’m not choosing now*”, i.e. choices were not forced. If a choice was deferred, subjects would have to make a decision at this menu at the end of the experiment if the menu was drawn for payment. We stress that no new information about any of the lotteries was given to subjects after the main part of the experiment. In particular, opting for “*I’m not choosing now*” was not associated with any informational gains.

Two of the binary menus,  $\{A1, A2\}$  and  $\{D, A2\}$ , featured a FOSD relation, hence an easy decision. Another three such menus,  $\{D, B1\}$ ,  $\{A1, B1\}$  and  $\{A1, D\}$ , featured a SOSD relation, hence an easy decision for any risk-averse (more likely) or risk-seeking expected-utility maximizing subject, who would *always* choose the SOSDominant and SOSDominated such options, respectively. Because all 7 lotteries have overlapping ranges, by Proposition 1 we can test subjects’ stability of risk attitudes (StAR) by checking whether they consistently opted for the same type of lottery at these four menus.

The remaining four binary menus, by contrast,  $\{B1, B2\}$ ,  $\{C1, C2\}$ ,  $\{A1, B2\}$  and  $\{D, B2\}$ , contained lotteries that were unrelated by SOSD, thereby resulting in potentially challenging decision problems that, as we hypothesized, could lead some subjects to opt for the costly choice avoidance/deferral option. For menu  $\{C1, C2\}$ , in particular, the absence of a dominant alternative was coupled by the complexity associated with the non-trivial probabilities in the definition of these lotteries, which included three significant decimal points. On the other hand, at menus  $\{A1, B2\}$  and  $\{D, B2\}$ , lotteries A1 and D are “almost” second-order stochastically dominating their respective counterparts, in the sense that such a relationship exists for most of the money range under consideration (see Table E.2 in Appendix E for more details). As such, deciding at these two menus is relatively easier than doing so at the other two, at least for risk-averse expected-utility maximizers. In general, these four menus, and especially  $\{B1, B2\}$ ,  $\{C1, C2\}$  invite a natural targeted test of Decisiveness via the (in)complete preferences channel.

The lotteries at the two “hard” binary menus  $\{B1, B2\}$  and  $\{C1, C2\}$ , moreover, were constructed so as to also allow for testing Independence. Indeed, letting  $R := (1, 0, 0, 0)$  be the fictitious lottery that assigns probability 1 to the zero prize and probability 0 to prizes 9, 10, 20 and 24, we have

$$C1 = \frac{1}{2}B1 + \frac{1}{2}R \quad \text{and} \quad C2 = \frac{1}{2}B2 + \frac{1}{2}R.$$

Thus, any expected-utility maximizing subject weakly prefers B1 to B2 if and only if they weakly prefer C1 to C2. This test for Independence is clearly different from existing and well-studied, Allais-type tests of that axiom such as the “common-ratio”, “common consequence” and “certainty” effects in two respects: (1) both the monetary outcomes and probabilities are less extreme here; (2) within each pair, the two lotteries have the same expected value (12 and 6, respectively) and are unrelated by SOSD. Hence, the behavioural trade-offs in

our test of Independence are different from those found in the typical tests of that axiom.

Finally, our collection of 9 binary menus was also designed to test for Transitivity at five distinct triples of lotteries:  $A1, D, A2$ ;  $A1, B2, D$ ;  $A1, B1, B2$ ;  $A1, D, B1$ ; and  $D, B2, B1$ . The 3 pairs of lotteries within each triple feature different combinations of dominance/non-dominance relations, thereby leading to varying levels of “difficulty” across triples for subjects to “pass” the Transitivity test. We discuss these in the next section.

Table 1: The 7 lotteries.

Prize Lottery	€/£ 0	€/£ 9	€/£ 10	€/£ 20	€/£ 24	Expected value
<b>A1</b>	$\frac{10}{100}$	–	$\frac{60}{100}$	$\frac{30}{100}$	–	€/£ <b>12</b>
<b>A2</b>	$\frac{20}{100}$	–	$\frac{50}{100}$	$\frac{30}{100}$	–	€/£ <b>11</b>
<b>B1</b>	$\frac{25}{100}$	–	$\frac{30}{100}$	$\frac{45}{100}$	–	€/£ <b>12</b>
<b>B2</b>	$\frac{25}{100}$	$\frac{40}{100}$	–	–	$\frac{35}{100}$	€/£ <b>12</b>
<b>C1</b>	$\frac{625}{1000}$	–	$\frac{150}{1000}$	$\frac{225}{1000}$	–	€/£ <b>6</b>
<b>C2</b>	$\frac{625}{1000}$	$\frac{200}{1000}$	–	–	$\frac{175}{1000}$	€/£ <b>6</b>
<b>D</b>	$\frac{15}{100}$	–	$\frac{50}{100}$	$\frac{35}{100}$	–	€/£ <b>12</b>

Table 2: The 15 lottery menus and some of the axioms they were designed to test.

Menu #	Lotteries in Menu	(Non-)Dominance structure	Additional remarks
1	A1 A2	A1 <i>FOSD</i> A2	
2	B1 B2	No <i>SOSD</i> dominance	Menus 2, 3 jointly test <i>Independence</i> : $C_i = \frac{1}{2}B_i + \frac{1}{2}(1, 0, 0, 0)$ , $i = 1, 2$
3	C1 C2	No <i>SOSD</i> ; ‘hard’ probabilities	
4	B1 D	D <i>SOSD</i> B1	Menus 2, 3, 5, 7, 13 feature ‘hard decisions’ and test ( <i>Stochastic</i> ) <i>Decisiveness</i> via the ( <i>in</i> ) <i>Completeness</i> channel
5	B2 D	No <i>SOSD</i> dominance	
6	A1 B1	A1 <i>SOSD</i> B1	
7	A1 B2	No <i>SOSD</i> dominance	Pairs {4, 6}, {4, 9}, {6, 9} test <i>Stable Attitudes to Risk</i> (Proposition 1)
8	A2 D	D <i>FOSD</i> A2	
9	A1 D	A1 <i>SOSD</i> D	Triples {1, 8, 9}, {2, 5, 4}, {6, 4, 9}, {7, 2, 6}, {7, 5, 9} test ( <i>Stochastic</i> ) <i>Transitivity</i>
10	A1 A2 C1	A1 is <i>FOSD</i> dominant	
11	A1 A2 C2	A1 is “nearly”* <i>FOSD</i> dominant	
12	A1 B1 B2	A1 is “nearly”** <i>SOSD</i> dominant	Pairs {1, 3} × {10}, {1, 3} × {11}, {2, 4, 5} × {13}, {2, 4, 5, 9, 12, 13} × {14}, {2, 6, 7} × {12}, {1, 3, 10, 11} × {15} test <i>Contraction Consistency</i> and <i>Regularity</i>
13	B1 B2 D	D is “nearly”** <i>SOSD</i> dominant	
14	A1 B1 B2 D	A1 is “nearly”** <i>SOSD</i> dominant	Menus 14, 15 are relatively ‘complex’ and test ( <i>Stochastic</i> ) <i>Decisiveness</i> via the <i>overload</i> channel
15	A1 A2 C1 C2	A1 is “nearly”* <i>FOSD</i> dominant	

\*A1 first-order stochastically dominates A2 & C1 and “almost” dominates C2 (see Table E.1).

\*\*A1 second-order stochastically dominates B1 & D and “almost” dominates B2;

D second-order stochastically dominates B1 and “almost” dominates B2 (see Table E.2).

We now turn to menus that contain more than two lotteries. From those with three such items, menu  $\{A1, A2, C1\}$  has a *FOSD* dominant lottery (A1) and therefore presents a relatively easy decision. A1 is “nearly” a *FOSD* dominant lottery at menu  $\{A1, A2, C2\}$  as well, dominating C2 in most of the relevant monetary range, and featuring twice as high an expected value (see Table E.1 in Appendix E for more details). On the other hand, menus  $\{A1, B1, B2\}$  and  $\{B1, B2, D\}$  do not have a dominant lottery under either the *FOSD* or *SOSD* criterion, although A1 and D, respectively, properly *SOSD* dominate B1

and “nearly” dominate B2 also, doing so for all  $x \in [0, 22.143]$  (see Table E.2 in Appendix E for more details). Finally, for these reasons, lottery A1 was “nearly” FOSDominant at menu  $\{A1, A2, C1, C2\}$  whereas no such option was available at  $\{A1, B1, B2, D\}$ , although A1 was “nearly” SOSDominant at the latter menu. Together with the nine binary menus above and their respective tests, the presence of these 3- and 4-lottery menus further allow for testing WARP and Contraction Consistency. In addition, the two quaternary menus invite tests for potential “*choice-overload*” effects (Iyengar and Lepper, 2000) whereby avoiding/deferring choice is more likely at larger menus.<sup>9</sup> In particular, they allow us to test whether any such effect is influenced by the presence or absence of a dominant alternative, as suggested, for example, by the meta-analyses on choice overload that were conducted by Scheibehenne et al. (2010) and Chernev et al. (2015), and as predicted by decision processes that are based on dominant choice with incomplete preferences (Gerasimou, 2018, Section 2). If such a mechanism is present in our data, then we should intuitively observe more violations of Decisiveness at menu  $\{A1, B1, B2, D\}$ , which lacks a FOSDominant lottery, than at menu  $\{A1, A2, C1, C2\}$ , which “nearly” does have such a clearly superior option.

### 3.2 Sequence of Choices, Tasks and Payments

After having received and having been quizzed on the experiment’s instructions, subjects were sequentially presented with 75 decisions, each on one of the 15 menus presented Table 2. Each menu was presented five times. In the set of the first 15 and in the set of the last 15 choices we presented each of the 15 menus once. The order of presentation in these two rounds of 15 choices was identical and common to all subjects, and coincides with the order that menus appear in Table 2. In the remaining 45 decision problems, i.e., from the 16th to the 60th, each menu was presented three times and the order was randomized for each participant.

Once subjects had gone through the 75 decision problems, and before the payout procedure commenced, they were asked to complete a series of questionnaires. These included questions on basic demographic characteristics as well as the ICAR-16 test of cognitive ability (Condon and Revelle, 2014). The latter contains four questions on each of the following four types of cognitive tasks: (i) letter and number series; (ii) verbal reasoning; (iii) three-dimensional rotation; (iv) matrix reasoning.<sup>10</sup>

---

<sup>9</sup>Dean et al. (2022) is a related recent study that, utilizing repeated choice data, develops a new method to test if (and to ultimately confirm that) opting for the default option—which in the authors’ experimental data on choices from arithmetic tasks was the simplest feasible option and always consisted of the same single number—was more likely in larger than in smaller menus. Similar to the many choice overload studies in psychology (see, specifically, references in the meta-analyses by Scheibehenne et al., 2010 and Chernev et al., 2015), our experimental design does not feature a default option of the same kind as the other feasible alternatives but, instead, a choice-avoidance/deferral outside option.

<sup>10</sup>A different set of items from the ICAR database of questions was also used, for example, by Chapman et al. (2023) to measure subjects’ cognitive ability.

### 3.3 Incentives

Choices in our experiment were incentivized. In particular, we informed subjects at the start of the experiment that one of the 75 decision problems would be randomly drawn at the end of the experiment and that the lottery they had chosen in that decision would be played out for them and paid out accordingly. If they had previously selected “*I’m not choosing now*” at that decision problem, they would be asked to choose a lottery from that menu at that point, and this would then be played out for them. Subjects received the lottery’s prize minus a fee of €/£ 0.5 for having deferred the decision. All subjects also received an additional €/£ 5 flat monetary fee.

Motivated by intuition and previous research (Tversky and Shafir, 1992; Danan and Ziegelmeyer, 2006; Costa-Gomes et al., 2022), we hypothesized that the option of not choosing could be chosen if subjects found a decision problem to be hard enough that they would be willing to risk the possibility of a small deduction (in our case, up to 10%) from their total monetary earnings in order to avoid/delay making an active choice there, either because they did not have a most preferred lottery at the relevant menu or because they considered the task of finding their most preferred lottery to be too cognitively costly. Although Tversky and Shafir (1992) did not use this terminology, the indecisiveness-based motivation for allowing choice avoidance/deferral follows their work. Making such deferral costly to subjects on the other hand –and hence embedding it in the design’s incentivization– follows Danan and Ziegelmeyer (2006) and Costa-Gomes et al. (2022). Unlike the design of this paper, however, the one in the latter study further allowed subjects to switch their active choice at their randomly selected menu at an even higher cost than the cost associated with avoidance/deferral. No such reversal was possible here. Furthermore, unlike the design in the working paper of Danan and Ziegelmeyer (2006), ours allows for binary as well as non-binary menus, does not frame the decisions as choices between menus of lotteries, and does not involve a week’s delay between when deferrals were made and when subjects were asked to make an active choice at their randomly selected menu.

### 3.4 Implementation and Procedural Details

The experiment was conducted in two locations: (i) the University of St Andrews Experimental Economics Lab on 17-18th January 2022 ( $N = 100$ ) and on 8-9th May 2023 ( $N = 115$ ); (ii) the University of Bonn Laboratory for Experimental Economics (BonnEcon-Lab) on 20th December 2022/11-12th January 2023 ( $N = 107$ ). Subject recruitment was done with ORSEE (Greiner, 2015) in St Andrews and hroot (Bock et al., 2014) in Bonn. The experimental interface was programmed in Qualtrics. All instructions were translated from English into German using that platform’s built-in translation tool, with manual adjustments made when necessary.

In all sessions the image describing each lottery was identical: its description was in English and the rewards were expressed in Euro (Figure B.2). St Andrews subjects were told

that the Euro amounts in the lotteries would be converted to Pounds Sterling at parity (one-for-one). After the choice part of the experiment –and before administering the additional questionnaires– the 107 subjects from Bonn and the 115 subjects from the St Andrews May ’23 sessions were told that they would receive an additional 2 Euro/Pounds to respond to a few more questions. This extra payment was first introduced in the Bonn sessions to bring the total expected hourly payment of every subject in line with that lab’s guidelines for the hourly rate in Euro. Compared to the ’22 St Andrews sessions, those conducted in ’23 in both locations also contained the following improvements to the experimental interface: (i) a fixed data-recording bug which had led to a few missing choice and questionnaire responses from 14 subjects in the ’22 sessions (we discarded those participants’ datasets); (ii) inclusion of the instructions that were missing from the 4 cognitive-ability questions that pertained to 3-dimensional rotation tasks. The implementation of our design across all sessions was identical in all other respects.

Upon entering the lab, subjects were asked to keep their phones switched off and be silent throughout the experiment. As soon as subjects finished all tasks, their randomly selected menu showed up on their screens, together with the reminder of the decision they had made at this menu. As an additional incentive for subjects to make deliberated and non-rushed decisions, they were told from the beginning that no participant would be able to receive their rewards and leave the lab in the first 60 minutes of the session. After such time, an experimenter went to the desk of each subject who had finished, had their chosen lottery played out for them using the random-number generating website <https://random.org> (St Andrews) or using an urn with pieces of paper numbered from 1 to 1000 (Bonn). A three-question understanding quiz was administered at the beginning of all sessions. Subjects could not proceed until they answered all questions correctly.

## 4 Deterministic Utility, Indifference, and Random Utility

### 4.1 Deterministically Rational Behaviour with Weak Preferences

We start by investigating the extent to which subjects behaved as if they were ordinal or expected-utility maximizers across *all* 75 decisions. More specifically, before entertaining the possibility that any choice reversals at different occurrences of the same menu is due to subjects’ “noisy” utility maximization or systematically boundedly-rational behaviour, we first examine the possibility that such reversals are instead rational manifestations of subjects’ indifference between lotteries, and that their overall behaviour is compatible with utility maximization when such alternating choices are viewed as the rational outcome of subjects’ indifference.

To answer this question, we first sliced every subject’s data into five regions, each corresponding to one “round” where the 15 distinct menus displayed in Table 2 were presented. In the following, we refer to the “ $i$ -th round” as the grouping of the set of decisions at those



menus subjects made the  $i$ -th time they saw them.<sup>11</sup> Following that, we accounted for the possibility that subjects are indifferent between distinct lotteries by *merging* their decisions at each menu across the 5 rounds, thereby creating a choice correspondence for each of them. Although this intuitive approach is endorsed by, among others, Mas-Colell et al. (1995, p.10) and has been supported by relevant free and open-source computational tools (Gerasimou and Tejiščák, 2018, [Prest](#)), apparently it has not been followed in experimental studies where subjects were repeatedly presented with the same menus.<sup>12</sup>

More specifically, letting  $C_i^n(\{A1, A2\})$  denote the (possibly empty) choice at menu  $\{A1, A2\}$  that subject  $n$  made, for example, the  $i$ -th time they saw that menu, for  $i \leq 5$ , this merging process is illustrated with the following hypothetical situation:

$$\left. \begin{array}{l} \overbrace{C_1^n(\{A1, A2\}) = \emptyset}^{\text{round-1 choice}} \\ C_2^n(\{A1, A2\}) = \{A2\} \\ C_3^n(\{A1, A2\}) = \{A1\} \\ C_4^n(\{A1, A2\}) = \{A2\} \\ \underbrace{C_5^n(\{A1, A2\}) = \{A1\}}_{\text{round-5 choice}} \end{array} \right\} \implies \underbrace{C^n(\{A1, A2\}) = \{A1, A2\}}_{\text{merged choice}}$$

The correspondence  $C^n$  thus defined satisfies

$$\emptyset \subseteq C^n(A) \subseteq A \quad \text{for every menu } A,$$

with  $C^n(A) = \emptyset$  iff  $C_i^n(A) = \emptyset$  for all  $i \leq 5$ .

We tested for conformity with ordinal and expected-utility maximization using the choice correspondence that was constructed in this way for each subject. For ordinal utility maximization, we used the Houtman-Maks (1985) (HM) method to find how close subjects' choices are to this mode of behaviour. While this is a standard—though often computationally demanding—revealed-preference test, we apply it to our data using a relatively new computational technique that also allows for testing a subject's behavioural proximity to maximization of a preference relation with indifferences.<sup>13</sup> To illustrate the idea of

<sup>11</sup>Recall that the order of choice menus was identical and common to all subjects in the first 15 and last 15 decisions, but that the 15 menus appeared three times in a subject-specific random order in the 45 decisions in between. Hence, in the sequence of decision problems numbered 16 to 60 in the Qualtrics survey, the same menu might be displayed consecutively in those 45 decisions. Furthermore, a decision problem whose Qualtrics-survey number was between 46 and 60 could be presented before a menu numbered between 16 and 30 or 31 and 45.

<sup>12</sup>That said, we remark that Balakrishnan et al. (2021) is a recent theoretical study that refines this approach by introducing a choice-probability threshold rule into the merging process, extending Fishburn (1978). The authors apply their method on the binary forced-choice data from Tversky (1969) to construct choice correspondences, and find that more than half of these are transitive under certain threshold values. Unlike that primarily theoretical study, here we do not impose a threshold in the analysis of our data and do not require the primitive or merged choices to be non-empty. Furthermore, we apply this model-free choice-merging approach to our richer and novel experimental dataset to carry out a more extensive test of subjects' conformity with rational choice under risk.

<sup>13</sup>Gerasimou (2021) is the only other existing paper that we are aware of which applies the same technique

this indifference-permitting HM test with an example, suppose  $C^n(\{A1, A2\}) = \{A1, A2\} = C^n(\{A1, A2, C2\})$  and  $C^n(\{A1, A2, C1, C2\}) = \{A1\}$ . The first two choices are consistent with maximization of a preference relation where subject  $n$  is indifferent between A1 and A2, and prefers either of these to C2. The third choice, however, contradicts this by suggesting that A1 is preferred to A2. Thus, treating the absence of A2 from the optimal choices at that menu as a mistake and “dropping” it from the dataset leads to an HM score of 1 in this example.

For our expected-utility analysis, furthermore, we construct a binary variable that indicates whether a subject’s behaviour is compatible with this model or not. Clearly, since every expected-utility subject is also an ordinal-utility maximizer, this group will be a subset of those subjects with an HM score of zero. Given the structure of the 7 lotteries in our experiment and of the 15 distinct menus presented to subjects, the additional tests that must be carried out on those subjects’ choices pertain to FOSD, Independence and StAR. In line with the discussion of Section 3.1, FOSD is satisfied if both  $C^n(\{A1, A2\}) = \{A1\}$  and  $C^n(\{A2, D\}) = \{D\}$  are true. Additionally, Independence is satisfied if and only if one of the following is true: (i)  $C^n(\{B1, B2\}) = \emptyset = C^n(\{C1, C2\})$ ; (ii)  $C^n(\{B1, B2\}) = \{B1\}$  and  $C^n(\{C1, C2\}) = \{C1\}$ ; (iii)  $C^n(\{B1, B2\}) = \{B2\}$  and  $C^n(\{C1, C2\}) = \{C2\}$ ; or (iv)  $C^n(\{B1, B2\}) = \{B1, B2\}$  and  $C^n(\{C1, C2\}) = \{C1, C2\}$ . Stability of risk attitudes, finally, is satisfied if and only if one of the following holds at *every* pair of menus of lotteries  $\{P, Q\}$  and  $\{P', Q'\}$  where  $P$  SOSD  $Q$  and  $P'$  SOSD  $Q'$ : (i)  $C^n(\{P, Q\}) = \emptyset = C^n(\{P', Q'\})$ ; (ii)  $C^n(\{P, Q\}) = \{P\}$  and  $C^n(\{P', Q'\}) = \{P'\}$ ; (iii)  $C^n(\{P, Q\}) = \{Q\}$  and  $C^n(\{P', Q'\}) = \{Q'\}$ ; (iv)  $C^n(\{P, Q\}) = \{P, Q\}$  and  $C^n(\{P', Q'\}) = \{P', Q'\}$ . Assuming that the other requirements of expected utility are satisfied, the choice patterns in cases (ii), (iii) and (iv), respectively, are necessary and sufficient for the agent to be revealed risk-averse, risk-seeking and risk-neutral, respectively.

The main results from this analysis are summarized in Table 3. Perhaps surprisingly, considering the relatively large number and difficulty of the experiment’s decision environment, approximately 26% of all subjects behaved as if they consistently maximized a stable, complete and transitive preference relation over the 7 lotteries across *all* 75 decisions. Importantly, moreover, for more than half of those subjects this conclusion could be reached only because we specifically tested for the possibility that subjects’ different choices at the same menus across distinct appearances of these menus could be due to subjects’ rational indifference rather than due to other factors.

Focusing on subjects’ behaviour at the specific subset of 45 decisions that correspond to the 5 appearances of the 9 binary menus, we further find that approximately 21% of all subjects behaved as expected utility maximizers when making those decisions. Furthermore, just less than half of those binary-menu expected-utility maximizers revealed at least one non-trivial indifference between distinct lotteries. Finally, the intersection of these two groups of

---

on possibly multi-valued choice correspondences. The experimental design, (riskless) choice domain, focus, research questions and results in that paper, however, are very different from those in the present study.

Table 3: Subjects who were rational across all 75 decisions, with or without strict preferences.

	Ordinal-Utility maximizers	Expected-Utility maximizers at binary menus	Expected-Utility maximizers
Revealing strict preferences	34 (11%)	36 (12%)	32 (10%)
Revealing strict preferences and indifferences	46 (15%)	33 (11%)	21 (7%)
<b>Total</b>	80 (26%)	69 (22%)	53 (17%)

subjects that comprise ordinal utility maximizers on the one hand and binary-menu expected-utility maximizers on the other corresponds –in our experiment– to the expected-utility maximizers at all 75 decisions. Fifty-three of the 308 subjects (17%) achieved this status, with 21 revealing a stable weak order with some indifferences and 32 revealing a stable strict preference relation that belongs to this class. Perhaps surprisingly, all but one subject in this category were revealed risk-averse. We summarise this information with the following:

**Highlight 1.** *Twenty-six percent of all subjects were perfect ordinal-utility maximizers throughout the experiment, with more than half revealing some indifferences. Moreover, 65% of those subjects were perfect expected-utility maximizers throughout, with nearly 40% of them revealing some indifferences. Fifty two of the 53 expected-utility maximizers were risk-averse and one was risk-neutral.*

Retaining our focus on the subjects’ overall behaviour across their 75 decisions through the resulting merged choices at the 15 distinct menus, we proceed next to an analysis of the main factors behind subjects’ deviations from indifference-permitting rational choice in the ordinal and/or expected-utility sense. A summary of this analysis is presented in Table 4. This clarifies that more than half of all subjects’ merged choices violated Contraction Consistency, Transitivity and StAR, while over 70% exhibited choice reversals, in violation of WARP. Among those violating Transitivity, however, no subject exhibited strict binary choice cycles of the form  $\{p\} = C(\{p, q\})$ ,  $\{q\} = C(\{q, r\})$  and  $\{r\} = C(\{p, r\})$ . Similarly, among the 8% of subjects who violated FOSD, none did so strictly in the sense of always choosing the dominated lottery. Twenty-two subjects (7%), moreover, violated Decisiveness by consistently avoiding/delaying making an active choice in at least one of the 15 distinct menus (we discuss later some patterns in those violations). Thirty-seven percent of subjects, finally, deviated from Independence in this analysis. While the proportion here is lower than those corresponding to some of the other consistency principles, we recall that –unlike the latter– there was only one pair of menus here where Independence could have been violated. By contrast, there were 3 pairs of menus where StAR could be violated, 5 triples for Transitivity, 20 pairs for Contraction Consistency, and even more for WARP. Cast in this light, our finding here that 37% of subjects violated Independence in this merged-choice analysis cannot by itself be interpreted as evidence suggesting that it is easier or harder to comply with this axiom than, say, Transitivity or StAR.

Table 4: Subjects whose 15 indifference-permitting merged choices comply with predictions of deterministic ordinal/expected utility theory.

<b>Decisiveness</b>	286 (93%)
<b>First-Order Stochastic Dominance</b>	282 (91.5%)
<b>Independence</b>	194 (63%)
<b>Stability of Attitudes to Risk</b>	178 (58%)
<b>Contraction Consistency</b>	153 (50%)
<b>Transitivity</b>	156 (49%)
<b>Weak Axiom of Revealed Preference</b>	88 (28.5%)

Figure 2: The most frequent weak revealed preferences that are compatible with expected utility based on each subject’s merged choices.

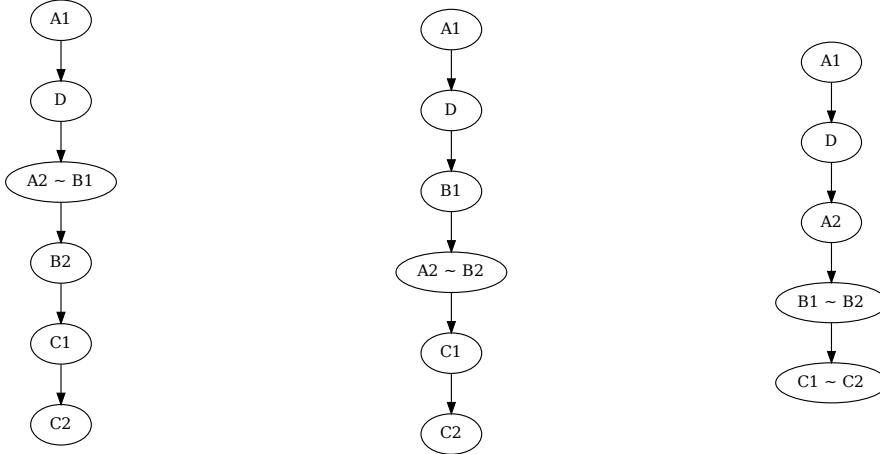


Figure 2 presents the three most frequent revealed weak preference orderings that are recoverable from each subject’s merged choices and, in addition, are compatible with expected-utility maximization. The first two, in particular, are compatible with the choices made by 29 of the 53 such subjects, while the last one is so with 20 of the remaining 24 subjects.<sup>14</sup> Recalling that  $A1 \text{ SOSD } D \text{ SOSD } B1$ , the fact that  $A1 \succ D \succ B1$  in all three orderings is a reflection of the strict risk aversion revealed by 52 of these 53 subjects’ overall behaviour.

## 4.2 Stochastically Rational Behaviour and the Role of Response Times

We now turn our attention to the hypothesis of *random* utility maximization in subjects’ overall behaviour across the five rounds of 75 total decisions. To this end, in the upper panel of Table 5 we report on subjects who conform with the two implications of this theory that were stated in Section 2.2, namely Regularity and Stochastic Decisiveness. Interestingly, of the 68 subjects (22%) who comply with these implications, and who are therefore potential random-utility maximizers, 65 and 42 also belong to the classes of 80 and 53 *deterministic*

<sup>14</sup>To arrive at these preference relations, we first used `Prest` to find each subject’s set of weak orders over the 7 lotteries that were compatible with that person’s 15 merged choices. The cardinality of this set ranged from 3 to 173 (mean/median: 92/63). This reflects the fact that only 15 out of the 120 non-singleton menus that are derivable from a set of 7 alternatives were shown to subjects, implying that the revealed preferences that are compatible with ordinal and/or expected utility maximization are not uniquely identified in general. Following this, we pinned down the EU-compatible orders with the highest and second-highest frequency. The displayed directed graphs of the `Prest`-computed optimal orderings were produced with a GraphViz (Gansner and North, 2000) add-on.

ordinal- and expected-utility maximizers, respectively, possibly exhibiting non-trivial indifference, that were identified in Table 3. The lower panel of Table 5, moreover, clarifies that 50-90% of all subjects satisfy Strong, Moderate and Weak Stochastic Transitivity, while 26% additionally satisfy Regularity. Finally, all 65 subjects who are potential random-utility maximizers also conform with all three stochastic-transitivity principles.

Table 5: Subjects whose 75 decisions comply with predictions of random utility theory and/or stochastic binary-choice consistency.

<b>Stochastic Decisiveness</b>	226	(73%)
<b>Regularity</b>	83	(27%)
Both axioms	68	(22%)
<b>Weak Stochastic Transitivity</b>	278	(90%)
<b>Moderate Stochastic Transitivity</b>	229	(74%)
<b>Strong Stochastic Transitivity</b>	155	(50%)
All five axioms	66	(21%)

These facts point to the possibility of strong—and heretofore undocumented—interpretational substitutabilities between random and deterministic but *indifference-permitting* utility maximization. The existence of such substitutabilities is important because the two modelling approaches have different behavioural foundations, and therefore their welfare implications are generally different, too. Although welcome at some level, the multiplicity of compatible explanations that can be found in the economist’s textbook toolkit presents a challenge to the analyst: which of them is more appropriate, and for which decision maker/subject? The question is certainly not new when it comes to explaining choices that deviate from utility maximization, in which case the problem is to find which one of possibly multiple “behavioural” models is better. The introductory quotation by Davidson and Marschak (1959) notwithstanding, however, this question does appear to be new insofar as the analyst must decide between a random-utility and a weak-preference deterministic utility explanation of the same data. This, in our view, merits additional exploration in future studies.

Towards delineating the two explanations using the data collected in our experiment, we follow a general insight from the recent literature and take a closer look on subjects’ response times. More specifically, this literature (see, for example, Alós-Ferrer et al., 2021) suggests that it is more likely for faster decisions at *binary forced-choice tasks* to be associated with easier choices—which are more reflective of clear preferences—than it is for slower decisions.

Although our experiment includes both binary and non-binary free-choice tasks—making it unsuitable for a direct test of the hypothesis—we can still explore a related prediction within our setup. Specifically, we consider whether the 15 out of 80 subjects who comply with deterministic utility maximization but not with random utility maximization tend to decide more quickly than the 65 subjects who comply with both models. The reasoning is that deterministically rational individuals might violate the Regularity axiom of stochastic choice yet still exhibit quicker decisions. Despite their deviations, these individuals reveal

stable, complete, and transitive weak preferences, assuming that each distinct choice from the same menu reflects a potentially optimal decision for them.

Our data supports this interpretation: the 15 subjects in this group showed faster decision-making, with a mean/median response time of 10.3/7.22 seconds, compared to 12.6/8.57 seconds for the group of 65 ( $p < 0.001$ ; two-sided Mann-Whitney test).<sup>15</sup>

Building on this finding, we examine the 65 dually explainable subjects to estimate how many might be better characterized as either “more likely deterministic” or “more likely random” utility maximizers. We start by calculating the average response time for each subject across 75 decisions. The median of these averages among the 15 deterministically rational subjects is 10.67 seconds. Using this value as a threshold, we classify the 65 subjects: those with average response times below 10.67 seconds are labeled “more likely deterministic” ordinal utility maximizers (40 subjects), and those above it as “more likely random” ordinal utility maximizers (25 subjects).

We apply the same procedure to the subgroup of 11 individuals who conform to expected utility but not random utility theory. From this, we estimate that among the 42 subjects who comply with both models, 28 are “more likely deterministic” expected utility maximizers and 14 are “more likely random” expected utility maximizers.<sup>16</sup>

We summarise this information as follows:

**Highlight 2.** *Among all subjects who are potential random-utility maximizers, 95.5% and 62%, respectively, are also deterministic ordinal- and expected-utility maximizers. Conversely, up to 81% and 79% of all deterministic ordinal- and expected-utility maximizers, respectively, are also random-utility maximizers. For subjects who comply with both the deterministic and random models, a response-times based analysis suggests that approximately 35% might be better explainable by the latter model.*

## 5 Learning to Be Rational, One Round at a Time

We now turn to the main question of the paper: *Do subjects come closer to maximizing utility with strict preferences as they make choices at the same decision problems repeatedly, without receiving any new information in the process?* Table 6 summarizes the relevant findings from this investigation that is based on the round-per-round behaviour of every subject according to the following criteria:

1. How many subjects’ decisions in each round are perfectly compatible with ordinal and expected-utility maximization with strict preferences, and how many are in violation of the seven axioms of rational choice under risk that were discussed in Section 2?

---

<sup>15</sup>The results are similar for the comparison between the 42 and 11 expected-utility maximizing subjects who are/are not potential random utility maximizers.

<sup>16</sup>Independent of our study, Alós-Ferrer et al. (2024) use response times in a different way to separate random and (non-expected-utility) deterministic preferences from repeated forced choices from binary menus of lotteries.

2. How many total and *active-choice* (i.e., excluding deferrals) decisions would have to be changed on average for each subject in each round to make decisions consistent with utility maximization with strict preferences?
3. How long does it take to make a decision on average?

With the exception of Decisiveness violators, whose proportion stayed in the 15%-16% range throughout, our findings unambiguously suggest that subjects learned to be more rational in all the above respects between the first and fifth rounds. Furthermore, in virtually all of these rationality criteria such learning occurred in a strictly monotonic way; that is, for almost every aspect of rationality that we consider, there is strictly higher conformity in the sample as we move from one round to the next.

Before discussing those findings in more detail it is worth pointing out that the random and subject-specific order of appearance of the middle 45 decision problems in our design alleviates potential concerns that such learning might be driven by the particular order of presentation. At the same time, the commonality of the presentation order between the first and fifth round and between subjects allows us to conduct a like-for-like comparison of behaviour at the beginning and at the end of the experiment and hence a targeted test of our learning hypothesis.

**Highlight 3.** *By the last round, 57.5% of all subjects' behaviour converges to utility maximization with strict preferences. For 69.5% of those subjects such convergence is to expected utility maximization. Response times fall, on average, by a factor of 2.5 by the last round.*

Indeed, the relevant proportions nearly doubled from 34% to 57.5% and from 24% to 40%, respectively, between the first and fifth rounds ( $p < 0.001$  in both cases). Notably, moreover, the proportion of strict-preference ordinal utility maximizers who are also expected-utility maximizers is relatively stable across rounds, and in the range of 61% – 70%. This suggests that subjects' ability to learn to comply with the general principles of rational choice is positively associated with their ability to do so for the more specialised principles of rational choice under risk. In addition, the proportions of *approximate* ordinal utility maximizers, defined as those who are at most one decision away from perfect conformity with that model (i.e., with an HM score less than or equal to one),<sup>17</sup> are relatively high and also increasing

---

<sup>17</sup>Echenique et al. (2023) have recently suggested an alternative approximation criterion that is distinct from Houtman-Maks and applies to data from budget-constrained demand under uncertainty. Another alternative to HM is the Swaps index proposed and axiomatically characterised in Apesteguia and Ballester (2015), which could be thought of as further penalizing how many positions away each mistaken choice is from what the rational choice would be at the respective menu under the postulated optimal preference ordering. This is an interesting additional measure that was computed, for example, on the active-choice experimental data in Costa-Gomes et al. (2022). We do not use it here because we view it as being more suitable for datasets where the menu collection is (approximately) symmetric in the sense that menus of different sizes appear an equal number of times. This is not so in our experiment, as there are 9 binary, 4 ternary and 2 quaternary menus, potentially biasing the Swaps index downwards relative to the symmetric ideal. A more detailed analysis of this issue is outside the scope of this paper.

Table 6: Subjects are significantly faster and more consistent with the maximization of stable and strict risk preferences in the last 15 than in the first 15 (identical) decision problems.

	1st	2nd	3rd	4th	5th	First vs Last 15 2-sided test <i>p</i> -values
	round of 15 decisions*					
<b>Utility Maximizers**</b>	106 (34%)	129 (42%)	143 (46.5%)	161 (52%)	177 (57.5%)	< 0.001 (Fisher's exact)
<b>Approximate Utility Maximizers***</b>	173 (56%)	202 (65.5%)	211 (68.5%)	225 (73%)	242 (78.5%)	< 0.001 (Fisher's exact)
<b>Active-Choice Utility Maximizers</b>	121 (39%)	156 (50.5%)	171 (55.5%)	186 (60%)	211 (68.5%)	< 0.001 (Fisher's exact)
<b>Average/median decisions away from Utility Maximization (HM)</b>	1.45 / 1	1.19 / 1	1.1 / 1	0.97 / 0	0.86 / 0	< 0.001 (Mann-Whitney)
<b>Average/median active decisions away from Utility Maximization</b>	1.13 / 1	0.88 / 0	0.79 / 0	0.65 / 0	0.55 / 0	< 0.001 (Mann-Whitney)
<b>Expected-Utility Maximizers at binary menus</b>	104 (34%)	113 (36.5%)	123 (40%)	123 (40%)	139 (45%)	0.005 (Fisher's exact)
<b>Expected-Utility Maximizers at all menus</b>	74 (24%)	85 (27.5%)	99 (32%)	111 (34%)	123 (40%)	< 0.001 (Fisher's exact)
<b>Average/median response time (in seconds)</b>	20.7 / 16.2	13.3 / 10.3	10.7 / 8.2	9.4 / 7.2	8.1 / 6.3	< 0.001 (Mann-Whitney)
<b>Violating FOSD</b>	21 (7%)	12 (4%)	12 (4%)	8 (2.5%)	9 (3%)	0.041 (Fisher's exact)
<b>Violating Independence</b>	98 (32%)	84 (27%)	87 (28%)	85 (27.5%)	76 (25%)	0.060 (Fisher's exact)
<b>Violating Stability of Attitudes to Risk</b>	112 (36%)	106 (34%)	96 (31%)	101 (33%)	80 (26%)	0.025 (Fisher's exact)
<b>Violating WARP</b>	186 (60%)	151 (49%)	137 (44.5%)	119 (38.5%)	97 (31.5%)	< 0.001 (Fisher's exact)
<b>Violating Contraction Consistency</b>	190 (62%)	153 (50%)	146 (47%)	122 (39.5%)	101 (33%)	< 0.001 (Fisher's exact)
<b>Violating Transitivity</b>	77 (25%)	67 (22%)	64 (21%)	56 (18%)	36 (12%)	< 0.001 (Fisher's exact)
<b>Violating Decisiveness</b>	50 (16%)	46 (15%)	47 (15%)	48 (15.5%)	46 (15%)	0.739 (Fisher's exact)

\*The order of menu presentation was identical and common across subjects in rounds 1, 5 and subject-specific in rounds 2, 3, 4. The reported statistics in round  $n \in \{2, 3, 4\}$  account for this and pertain to the  $n$ -th appearance of each of the 15 distinct menus for each subject. \*\*Unless the "active-choice" qualification is present, both active-choice and deferral decisions are accounted for and, where relevant, penalized. \*\*\*Up to one decision away from Utility Maximization.



throughout, from 56% initially to 78.5% finally ( $p < 0.001$ ).<sup>18</sup> Furthermore, the distribution of subjects' HM scores is also shifted significantly to the left in the last 15 compared to the first 15 decisions, down from 1.45 to 0.86 decisions away from rationality, on average ( $p < 0.001$ ).

## 5.1 Axiom Violations

We now turn our focus on the evolution of subjects' conformity (or lack thereof) with each of the deterministic choice axioms that were presented in Section 2.1. These findings can be summarized thus:

**Highlight 4.** *The proportions of subjects violating each of Transitivity, Contraction Consistency, WARP, Independence, FOSD and Stability of Risk Attitudes in the last round are significantly lower than in the first. The proportion of those violating Decisiveness is stable.*

This significant decreasing trend notwithstanding, the most persistent violations were those of Contraction Consistency (typically, but –in our free-choice environment– not always, associated with WARP violations too; see Section 2.1 for more details), with 101 subjects (33%) still deviating from this consistency principle in their last 15 decisions. This, despite the fact that Contraction Consistency and WARP are the two principles with the largest gains in compliance (29 percentage points). At the same time, the smallest gains in compliance were seen with the Independence and StAR axioms, where the proportions of violators fell from 32% to 25% ( $p = 0.060$ ) and from 36% to 26% ( $p = 0.025$ ), respectively). The fact that there is no SOSD ranking in the two menus involved in the Independence test and the ensuing relatively high degree of decision difficulty may partly explain the slower pace of learning with respect to that axiom. FOSD on the other hand is violated by very few subjects (down from 7% to 3% by the fifth round). This is in line with findings in Levy (2008) where, as in this study, subjects were shown lotteries without any budget-constrained environments. It is, however, in stark contrast to the findings in Dembo et al. (2021) where FOSD is routinely violated by budget-constrained subjects choosing Arrow-Debreu securities under uncertainty. We hypothesize that the very different decision environments and ways in which decision problems are presented to subjects are largely responsible for this discrepancy.

As far as Decisiveness is concerned, finally, the proportion of subjects who violated this principle by deferring in at least one decision problem remained stable and in the 15%–16% range in all five rounds (Figure 3). While these proportions themselves are in line with those seen in deferral-permitting studies with no repeated choices, the fact that they remained constant is a novel and, in our view, interesting finding. First, it suggests that subjects who are willing to incur a monetary cost in order to avoid making an active choice at a difficult decision problem are less likely than one might have thought to change this attitude with

---

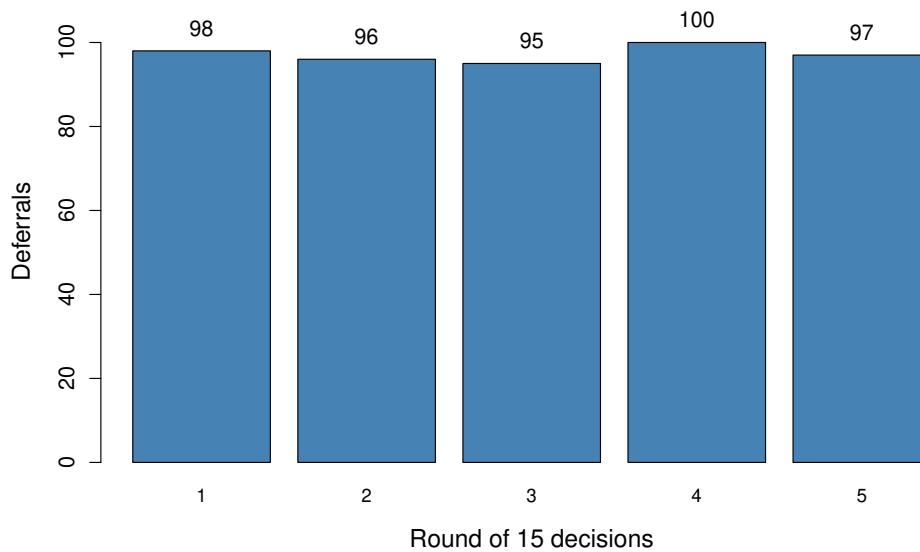
<sup>18</sup>In simulations with uniform-random behaving subjects on this collection of menus, the 2.5th percentile in the HM score distribution is 2 decisions. This suggests that our approximation threshold of one decision is unlikely to have been reached by human subjects who behaved randomly.

Table 7: Frequencies of Decisiveness violations at different menus.

Menu		Deferrals in 75 decisions	Deferrals in merged 15 decisions	1st-round presentation order	Average 1st-round response time	Average response time overall		
A1	A2	13	0	1	30.20	12.01		
<b>B1</b>	<b>B2</b>	<b>59</b>	<b>2</b>	<b>2</b>	<b>29.85</b>	<b>15.10</b>		
<b>C1</b>	<b>C2</b>	<b>192</b>	<b>20</b>	<b>3</b>	<b>30.23</b>	<b>16.69</b>		
B1	D	19	1	4	21.76	11.71		
<b>B2</b>	<b>D</b>	<b>26</b>	<b>1</b>	<b>5</b>	<b>16.00</b>	<b>10.22</b>		
A1	B1	16	1	6	16.33	10.22		
<b>A1</b>	<b>B2</b>	<b>16</b>	<b>1</b>	<b>7</b>	<b>13.01</b>	<b>8.95</b>		
A2	D	19	1	9	14.20	9.57		
A1	D	13	1	8	13.53	10.05		
A1	A2	C1	14	1	10	24.80	12.98	
A1	A2	C2	16	1	11	15.80	10.94	
<b>A1</b>	<b>B1</b>	<b>B2</b>	<b>22</b>	<b>1</b>	<b>12</b>	<b>21.08</b>	<b>14.27</b>	
<b>B1</b>	<b>B2</b>	<b>D</b>	<b>23</b>	<b>1</b>	<b>13</b>	<b>19.05</b>	<b>14.50</b>	
<b>A1</b>	<b>B1</b>	<b>B2</b>	<b>D</b>	<b>25</b>	<b>3</b>	<b>14</b>	<b>26.64</b>	
A1	A2	C1	C2	13	2	15	17.29	11.98

Note: in bold are “hard” menus that were hypothesised ex ante to have a higher deferral frequency (cf. Table 2). First-round and overall mean/median response times between “hard” & “non-hard” menus are, respectively, 22.3/16.6 vs 19.2/15.9 ( $p = 0.002$ ; 2-sided Mann-Whitney test) and 12.93/9.13 vs 12.02/8.89 ( $p = 0.004$ ).

Figure 3: The overall avoiding/deferring behaviour does not change over time.



more presentations of that decision problem. This could be thought of as evidence in support of what Sen (1997, pp. 763-764) referred to as “assertive” rather than “tentative” incomplete preferences: “It is useful to consider the distinction between: tentative incompleteness, when some pairs of alternatives are not yet ranked (though they may all get ranked with more deliberation or information), and assertive incompleteness, when some pair of alternatives is asserted to be “non-rankable”. Assertive incompleteness is the claim that the failure of completeness is not provisional—waiting to be resolved with, say, more information, or more penetrating examination. The partial ranking, or the inexhaustive partitioning, may simply not be “completable”, and affirming that some  $x$  may not be rankable vis-à-vis some  $y$  may be the right answer in these cases.”

This is corroborated by the findings shown in Table 7, which reports the deferral frequencies per menu. In line with our hypothesis (Section 3.1), three of the four binary menus featuring no SOSD relationship between the feasible lotteries have the highest absolute deferral frequencies: 192 for  $\{C1, C2\}$ , 59 for  $\{B1, B2\}$  and 26 for  $\{B2, D\}$ . Strikingly, moreover, the table also clarifies that 20 subjects always deferred at menu  $\{C1, C2\}$ , which, in addition to featuring no dominance relation, also had the most complex-looking lotteries and the lowest expected values. It is possible therefore that the decision to defer at this particular menu was driven by some convex combination of decision difficulty and aversion to incur the cognitive effort given the relatively higher complexity and lower stakes involved. Also broadly in line with our hypothesis, next in the list of deferral-inducing menus are  $\{A1, B1, B2, D\}$  (25),  $\{B1, B2, D\}$  (23) and  $\{A1, B1, B2\}$  (22). These menus feature the “nearly” SOSD dominant lotteries A1, D and A1, respectively, (see Section 3.1), but no FOSD dominance relation. Moreover, despite the relatively low deferral frequencies at the two menus with four lotteries,  $\{A1, B1, B2, D\}$  (25) and  $\{A1, A2, C1, C2\}$  (13), we note the following fact that lends support to the dominance channel in the occurrence and alleviation of choice overload that was discussed in Section 3.1 (Scheibehenne et al., 2010; Chernev et al., 2015):

**Highlight 5.** *The deferral rate is lower at the four-element menu with a (“nearly”) FOSD dominant lottery than at the one without (0.8% vs 1.6%;  $p = 0.071$ , 2-sided Fisher’s exact).*

In addition, as is also clarified in Table 7 (footnote), subjects’ response times, both in the first round and overall, are significantly longer at the 4 “hard” binary menus that lacked a SOSD dominant lottery and the 3 non-binary menus that lacked a perfectly or approximately FOSD dominant lottery than at the remaining 8 menus. It is noteworthy that such a significant difference—of almost 3 seconds, on average—is present in first-round response times despite the alleviating effect brought about by the fact that the menu ( $\{A1, A2\}$ ) that was seen first by all subjects in that round featured an “easy” decision, which was nevertheless associated with a long response time because it gave subjects their very first exposure to the experiment’s lotteries and computer interface. Taken together, these facts indicate that: (i) Decisiveness violations are more likely at “hard” binary-menu decisions than at larger menus, suggesting a bigger role of incomparability/incomplete preferences than choice

overload in our context; (ii) other things equal, deferral is more likely at large menus that do not have an obviously dominant lottery than at those that do; (iii) deferring behaviour is largely unchanged over the course of the experiment and, in general, is associated with longer response times.

We now take a closer look at violations of Transitivity. The proportion of subjects violating this axiom in at least one of the five possible triples goes down from 25% initially to 12% eventually ( $p < 0.001$ ) in this free-/non-forced-choice environment. Table 8 groups the total violations across the subjects' 75 decisions by associating each with the relevant triple where it occurred, and clarifies the (F)(S)OSD dominance structure, if any, within each of the three pairs in the triple. At the two extremes lie triples A1-D-A2 and A1-D-B1. The former features one second-order and two first-order dominance comparisons, the highest such comparisons among all five triples. As such, one would intuitively expect few violations here, which is indeed what we find (19; 1.2%). The latter triple instead features three second-order dominance pairwise relations. By Proposition 1 therefore, any expected-utility maximizing subject with risk-averse or risk-seeking strict preferences would satisfy Transitivity at this triple. Contrary to this prediction, we find that violations are actually highest here (136; 8.8%), despite the fact that triples D-B2-B1 and A1-B1-B2 each included a pair without any dominance relation (thereby increasing the cumulative decision difficulty within the triple) and one with an "approximate" SOSD relation (denoted here by  $\approx$  SOSD), while A1-D-B2 either featured proper or approximate SOSD relations in each of the three pairs. The violations-induced ordering between those three triples are broadly in line with this intuition, however, with the first (99; 6.4%) followed by the second (97; 6.3%), which in turn is followed by the third (79; 5.1%), even though differences between consecutive triples in this ranking are not significant. Although the high incidence of intransitivities at triple A1-D-B1 is somewhat puzzling, a possible explanation for it is the postulated presence of risk-neutral subjects who, by definition, are indifferent between any two lotteries in the triple and might therefore reveal single-valued choices in violation of Transitivity (such cases are picked up by our indifference-inclusive analysis of Section 4.1). Another potential explanation, finally, is that it emerges as a by-product of the significantly declining yet persistently high proportion of subjects violating StAR.

Table 8: Frequencies of violations of Transitivity at the relevant triples of binary menus.

<b>Lottery triple</b>	A1 D A2	A1 D B2	A1 B1 B2	D B2 B1	A1 D B1
<b>FOSD or SOSD pairwise relations within the triple</b>	A1 SOSD D D FOSD A2 A1 FOSD A2	A1 SOSD D D $\approx$ SOSD B2 A1 $\approx$ SOSD B2	A1 SOSD B1 B1 no-SOSD B2 A1 $\approx$ SOSD B2	D $\approx$ SOSD B2 B2 no-SOSD B1 D SOSD B1	A1 SOSD D D SOSD B1 A1 SOSD B1
<b>Intransitivities</b>	19	79	97	99	136
<i>p</i> -value from 2-sided Fisher's exact test	$p < 0.001$		$p = 0.186$	$p = 0.941$	$p = 0.014$

We conclude this analysis with the following finding of a different nature:

**Highlight 6.** *Subjects’ decisions in the last round are more than twice as fast as in the first.*

Indeed, the above-documented learning effect is accompanied by a significant shift to the left in the subjects’ distributions of mean response times between the first and last 15 decisions (20.7 vs 8.1 seconds per problem;  $p < 0.001$ ). This is an important finding because, considering also the growing literature in support of the argument that decisions are faster when preference comparisons are easier,<sup>19</sup> it suggests that subjects who learned to maximize (expected) utility in this experiment have done so while in the process of discovering or constructing their (stable, complete and transitive) *preferences*. But the finding is also important for a distinct reason: it suggests that the so-called *chronometric function* that quantifies this postulated relationship between response times and the relative easiness/difficulty of decisions is not stationary but changes with the agent’s experience. This in turn invites the analyst to exercise caution when defining and interpreting the chronometric function in a given setting.

## 6 Stability of Learning and Revealed-Preference Convergence

### 6.1 Learning and Strict Preferences

The findings in the last section invite a natural follow-up question: *To what extent is learning stable from one round to the next?* More specifically, do subjects whose decisions in one round are UM- or EUM-rational—after accounting for possible indifferences in their overall behaviour—continue to be so in the next round? The results presented in Table 9 point to a positive answer. In particular, for both ordinal and expected-utility maximization, and for each of the four possible round transitions (i.e., from the first to the second etc.), the majority of subjects exhibited *stable learning*. Moreover, the proportions of such stable learners are increasing at similar rates from the first to the fourth transition, from 80.5% to 90% for ordinal and from 85% to 90% for expected-utility maximization.

Table 9: Subjects who comply with ordinal and expected utility maximization in one round and continue to do so in the next round\* (in parenthesis are subjects who did so with *identical strict preferences*).

	Utility Maximization		Expected Utility Maximization	
Round 1 to 2	91	80.5% (42/81)	63	85% (38/45)
Round 2 to 3	118	89% (67/112)	73	86% (54/56)
Round 3 to 4	128	86.5% (90/121)	91	92% (71/74)
Round 4 to 5	149	90% (99/141)	100	90% (82/86)

\*Subjects who were (E)UM-rational throughout their 75 decisions under *weak* preferences have also been included.

A closely related question is whether the transition from the first to the fifth round, and the reported improvement in subjects’ conformity with rational choice, is also associated with learning their *preferences*. To assess this, we included an end-of-experiment question that only showed up in subjects who had chosen different lotteries at the same binary menu in the

<sup>19</sup>See Alós-Ferrer et al. (2021) and references therein

1st and 5th times they were asked to choose from it, and for every menu where such a reversal occurred. Importantly, in order to prevent the occurrence of experimenter demand effects these questions and the corresponding menus were phrased and presented neutrally. More specifically, the question stated the following: “At two different times when you saw this menu previously, you chose different lotteries. Which one of the two lotteries do you think you prefer more now, if any?” Similar to how they did in the main part of the experiment, subjects’ responded to these questions either by selecting one of the two displayed lotteries or by opting for “*I don’t know*” (instead of “*I’m not choosing now*”). No indication was given by experimenters on when the choice discrepancies at those menus appeared or on which lottery was chosen earlier and which one later.

Table 10: The ex-post stated preferences of subjects who chose different lotteries at the same binary menus in the 1st and 5th rounds.

	<b>A1,A2</b> (FOSD)	<b>A2,D</b> (FOSD)	<b>A1,B1</b> (SOSD)	<b>A1,D</b> (SOSD)	<b>B1,D</b> (SOSD)
Lottery of first choice	13.3%	50.0%	18.8%	37.5%	24.0%
Lottery of last choice	86.7%	50.0%	79.7%	62.5%	74.7%
<i>“I don’t know which one I prefer”</i>	0.00%	0.0%	1.5%	0.0%	1.3%
Total 1st/5th-round reversals at menu	15	2	69	72	75

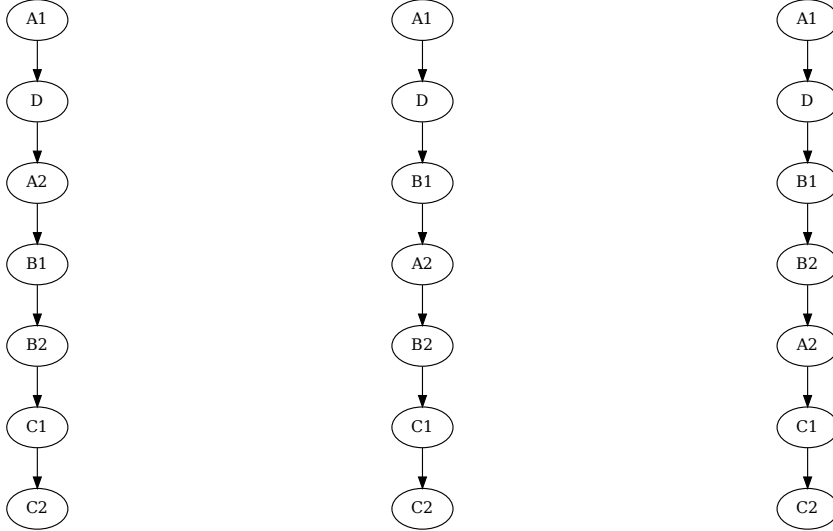
  

	<b>A1,B2</b> (“near SOSD”)	<b>B2,D</b> (“near SOSD”)	<b>B1,B2</b> (SOSD-unranked)	<b>C1,C2</b> (SOSD-unranked)
Lottery of first choice	26.3%	22.7%	29.0%	17.2%
Lottery of last choice	73.7%	72.7%	66.3%	79.6%
<i>“I don’t know which one I prefer”</i>	0.0%	4.5%	4.7%	3.2%
Total 1st/5th-round reversals at menu	38	44	86	93

The summary of subjects’ responses to these questions is presented in Table 10, and separated into three groups, depending on whether a FOSD, SOSD or neither of these dominance relations exists between the respective lotteries. As expected, the lowest proportion of reversals occurred at the two menus of the first type (17 in total), with the majority (14) of affected subjects stating that they preferred the (dominant) lottery corresponding to their last choice (recall that they were not told which was which). For each of the 7 menus in the remaining 2 groups, the majority of subjects (66 - 80%) again stated their last-chosen lottery as their preferred one at that menu. In line with our hypothesis, the highest rates of “I don’t know” responses were seen in 3 of the 4 menus where no stochastic-dominance relation exists (3 - 5%). Two of these menus, moreover, namely  $\{B1, B2\}$  and  $\{C1, C2\}$ , were also associated with the largest numbers of choice reversals between those two rounds. Across the three groups, the frequency-weighted means for first-choice, last-choice or agnostic stated-preference responses were 25.5%, 71% and 3.5%. These findings are generally in the same direction as those in the previous section and indicate that, in the aggregate, the noted improvements in subjects’ rationality during the experiment were associated with them learning their preferences along the way.

We finally turn to the most frequent preference orderings over lotteries that were compati-

Figure 4: The most frequent strict revealed preferences that are compatible with expected utility in the last 15 decisions.



ble with the behaviour of those 123 subjects who complied with expected utility maximization in their last 15 decisions. To do this we followed the same process that was explained in Section 4.1, focusing here on strict preferences instead. Figure 4 shows the three most common such orderings, compatible with 69 of 123 EU-compliant subjects. The three orders differ only in how they rank the middle three lotteries. The first and last are the most and least risk-averse among them, respectively, ranking A2—whose expected value is £11—above and below B1, B2—both of which have expected value £12.

## 6.2 Learning and Weak Preferences

In Section 4 we showed that accounting for the possibility of rational indifferences by merging a subject’s distinct choices at different presentations of the same menu contributes meaningfully to preference elicitation and explanation of observable behaviour. But, as the results in Section 5 and 6.1 show, it is clear that some subjects’ distinct choices, e.g. at FOSD-ranked menus, were errors that were subsequently corrected. This is manifested, for example, as a choice of A2 from menu  $\{A1, A2\}$  in the first round that is followed by a choice of A1 in rounds 2–4. At the same time, we know from Section 2 and Proposition 1 that if a subject violates StAR in any individual round of 15 decisions, where a single lottery could have been chosen from any menu, then this subject cannot be a strictly risk-averse or strictly risk-seeking expected-utility maximizer. Yet, this by itself does not rule out the possibility of *risk-neutral* behaviour under this model because, in our experiment, such a subject would be indifferent between any two lotteries that have the same expected value, including the three lotteries A1, D and B1 that are involved in testing StAR. One may therefore hypothesize about the potential simultaneous presence of learning and risk neutrality in our experiment, which would call for a refinement of the indifference-permitting analysis that was presented in Section 5.

To test this hypothesis, we compare subjects’ behaviour in their first and last 30 decisions

and look for differences in the preferences revealed by their respective choice correspondences. In view of the strict-preference learning results of Table 6, we are interested in: (i) subjects who cannot be regarded as (E)UM-rational when their behaviour in rounds 4 and 5 is analysed in isolation but *could* be regarded as such when their choices there are merged; and (ii) in testing whether any such subjects are relatively more numerous than those for whom round-1 and round-2 data are analysed instead. For this analysis, we continue to regard the second and fourth rounds as we did in Section 5.1. Building on this, we now also define each subject’s “early” and “late” choice correspondence in the way just described. We observe that either of these assigns a multi-valued choice at some menu if and only if it assigns exactly two “indifference-candidate” lotteries at that menu, each of which has a 50% choice frequency in the respective restricted domain.

Table 11: Comparing the (E)UM rationality of “early” and “late” choice correspondences.

	<b>First 2 rounds</b>	<b>Last 2 rounds</b>	<i>p</i> -value
<b>UM subjects</b>	86 (28%)	135 (44%)	< 0.001
<b>EUM subjects</b>	52 (17%)	83 (27%)	0.003
<b>UM with indifferences</b>	29	27	
<b>EUM with indifferences</b>	1	2	
<b>Risk-neutral EUM</b>	0	0	

Table 11 summarizes the results of this analysis. Specifically, the one and two expected-utility maximizers with at least one indifference in their first and last 30 decisions—at SOSD-unranked menus  $\{B2, D\}$  and  $\{A1, B2\}$ , respectively—always chose the SOSD-dominant lottery at all those menus where such a lottery was available. Hence, no potentially latent risk-neutral subjects are revealed by this analysis, leaving this status to the single individual who was identified as such in Section 4 where all five rounds of decisions were accounted for.

## 7 The Role of Cognitive Ability

The recent literature has documented a positive link between decision makers’ cognitive ability and their patience, risk tolerance and proximity to rational behaviour, both in non-strategic (Dohmen et al., 2010; Becker et al., 2012; Dohmen et al., 2018; Chapman et al., 2023; Echenique et al., 2023) and strategic environments (Proto et al., 2019, 2022; Gill and Prowse, 2016; Gill and Rosokha, 2024). Extending the investigation of such links, we were interested to assess the potential role of cognitive ability in our experimental subjects’ choice consistency and, additionally, their learning –or lack thereof– to maximize (expected) utility. We state from the outset that our analysis here is exploratory.

To do so, and as was noted previously, after the main part of the experiment we invited subjects to complete the ICAR-16 cognitive-ability questionnaire due to Condon and Revelle (2014). Specifically, a cognitive ability score between 0 and 1 was constructed for every subject, coinciding with the proportion of their correct answers. In addition to the all-inclusive ICAR-16 score, we also constructed in this way a variety of other scores that



featured one or more of the 4 blocks of questions from the Letter-Numeric sequence (LN), Verbal Reasoning (VR), Matrix Reasoning (MR) and 3-Dimensional Rotation (3DR) items.

Figure 5: Associations between cognitive ability and choice consistency in subjects' merged decisions.

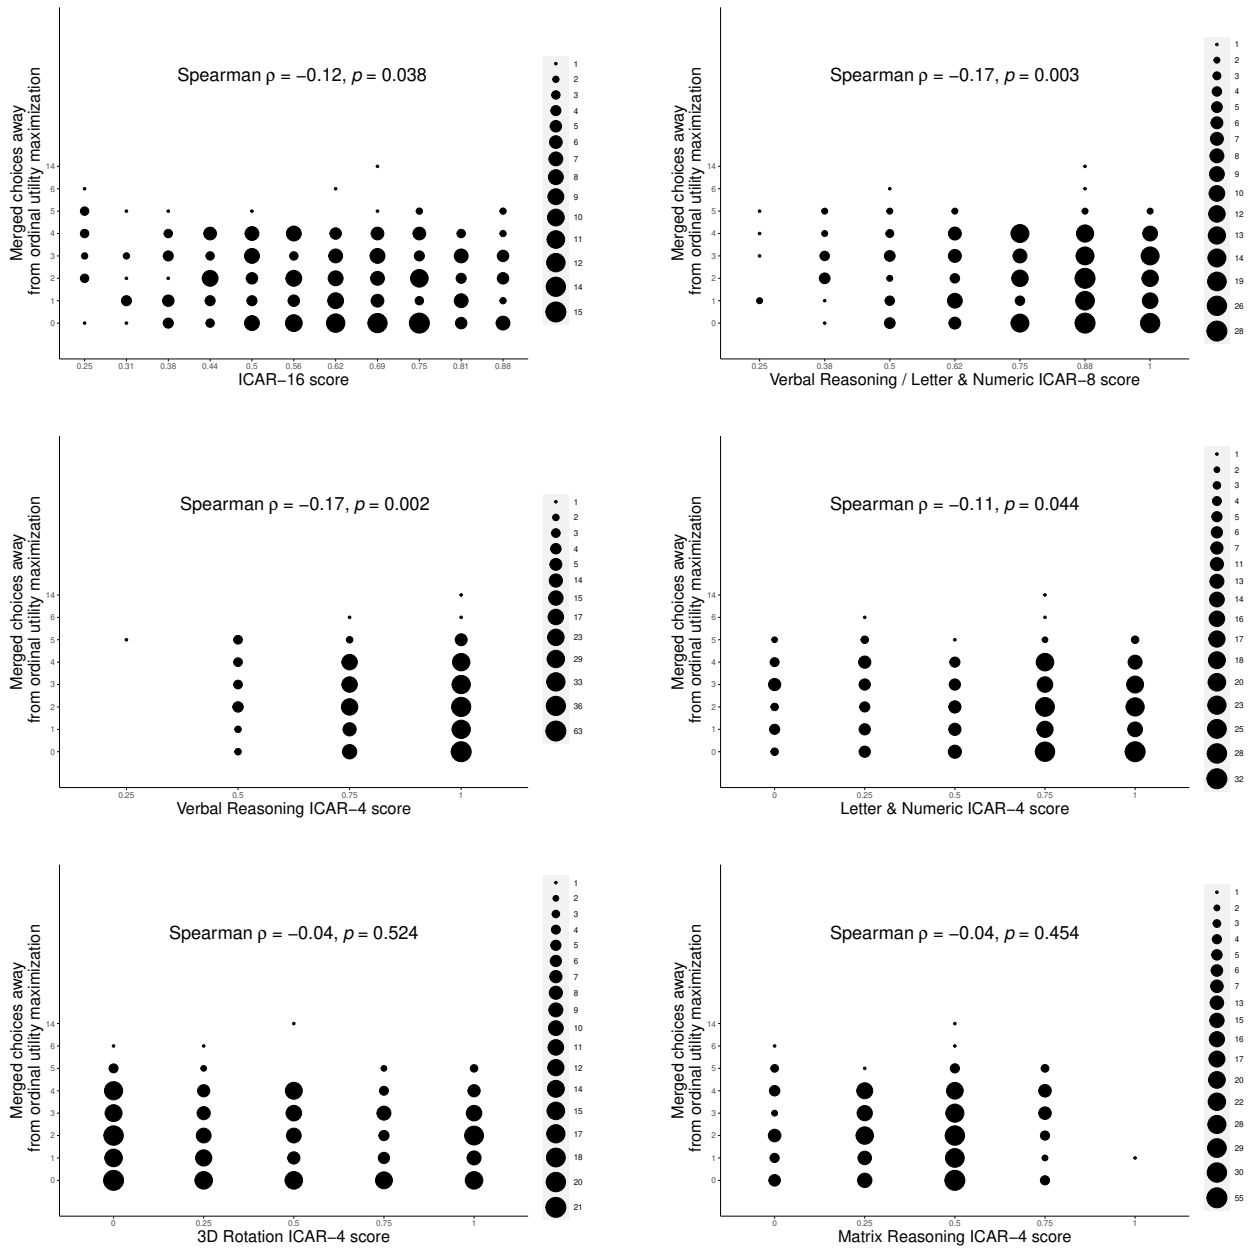


Figure 5 shows the density-inclusive correlograms between subjects' HM scores in their merged 15 decisions (where relevant, also penalizing deferrals) and their ICAR-16 scores, as well as a variety of more theme-focused subscores. Cognitive ability was significantly positively correlated with HM-consistency in subjects' overall choice behaviour under the ICAR-16 measure ( $\rho = -0.12; p = 0.038$ ). Furthermore, testing separately whether there is an association between consistency and each of the ICAR-4 measures that are formed by the LN, VR, MR and 3DR items in the ICAR-16 questionnaire, we find positive relations with all four, but only those with VR ( $\rho = -0.17; p = 0.002$ ) and LN ( $\rho = -0.11; p = 0.044$ ) are economically and statistically significant, as is the score formed by combining these two ( $\rho = -0.17; p = 0.003$ ).

Table 12: Relationships between cognitive ability and (not) learning to maximize utility.

Subject groups under comparison (group size in parenthesis)	Average ICAR Test Scores*							
	VR LN 3DR MR	VR LN 3DR	VR LN	LN 3DR	VR	LN	3DR	MR
UM in first 15 decisions or throughout (113) vs UM in neither (195)	0.626 vs 0.600 <i>p</i> = 0.271	0.697 vs 0.664 <i>p</i> = 0.230	0.833 vs 0.771 <i>p</i> = 0.009	0.580 vs 0.557 <i>p</i> = 0.558	0.931 vs 0.878 <i>p</i> = 0.004	0.735 vs 0.664 <i>p</i> = 0.103	0.425 vs 0.450 <i>p</i> = 0.510	0.412 vs 0.409 <i>p</i> = 0.776
EUM in first 15 decisions or throughout (87) v EUM in neither (221)	0.62 vs 0.605 <i>p</i> = 0.605	0.685 vs 0.673 <i>p</i> = 0.709	0.826 vs 0.781 <i>p</i> = 0.067	0.557 vs 0.568 <i>p</i> = 0.712	0.94 vs 0.881 <i>p</i> = 0.002	0.713 vs 0.681 <i>p</i> = 0.517	0.402 vs 0.456 <i>p</i> = 0.247	0.425 vs 0.404 <i>p</i> = 0.302
Non-UM in first 15 or throughout & UM in last 15 (78) v UM in neither first or last 15 or throughout (117)	0.632 vs 0.579 <i>p</i> = 0.036	0.706 vs 0.636 <i>p</i> = 0.026	0.804 vs 0.749 <i>p</i> = 0.129	0.614 vs 0.519 <i>p</i> = 0.021	0.891 vs 0.870 <i>p</i> = 0.503	0.718 vs 0.628 <i>p</i> = 0.178	0.510 vs 0.410 <i>p</i> = 0.098	0.410 vs 0.408 <i>p</i> = 0.810
Non-EUM in first 15 or throughout & EUM in last 15 (61) v EUM in neither first or last 15 or throughout (160)	0.632 vs 0.595 <i>p</i> = 0.183	0.714 vs 0.657 <i>p</i> = 0.103	0.816 vs 0.768 <i>p</i> = 0.319	0.631 vs 0.545 <i>p</i> = 0.045	0.881 vs 0.881 <i>p</i> = 0.643	0.75 vs 0.655 <i>p</i> = 0.153	0.512 vs 0.434 <i>p</i> = 0.209	0.385 vs 0.411 <i>p</i> = 0.462

\*Each of VR, LN, 3DR and MR refers to the respective subset of 4 items in the ICAR-16 test that includes Verbal Reasoning, Letter-Numeric, 3-Dimensional Rotation and Matrix Reasoning questions, respectively. The different columns report average scores and 2-sided Mann-Whitney *U* test *p*-values for the corresponding combinations of questions.

We also investigated potential associations between cognitive ability and learning. These findings are summarised in Table 12, where average scores are reported for various cognitive ability measures that the ICAR-16 gives rise to. The main insights from this analysis are as follows:

1. “*Early*” or “*stable*” utility maximizers, i.e., those who are rational in the ordinal sense at either their first 15 or all 75 decisions (the latter possibly with indifferences), have a higher VR (0.93 vs 0.88;  $p = 0.004$ ) and combined VR-LN (0.83 vs 0.77;  $p = 0.009$ ) cognitive score than those who did not behave as utility maximizers at either their first 15 or all 75 decisions (the latter necessarily with indifferences).
2. “*Late*” utility maximizers, or “*learners*”. That is, subjects who are not rational in this sense in either their first 15 decisions or throughout, but do become so in their last 15 decisions, are more cognitively able than those who do not have this status at any of these points, both according to the more holistic ICAR-16 measure (0.63 vs 0.58;  $p = 0.036$ ) and the more focused ones that combine VR, LN, 3DR (0.71 vs 0.64;  $p = 0.036$ ) and LN, 3DR (0.61 vs 0.52;  $p = 0.021$ ), as well as according to 3DR alone (0.51 vs 0.41;  $p = 0.098$ ) and, not significantly, LN too (0.72 vs 0.63;  $p = 0.178$ ).

Notably, as is detailed in Table 12, the direction in both these findings is the same when the rationality criterion is expected-utility rather than ordinal-utility maximization. In this case, however, the differences in average scores between the relevant (non-)learning groups are not always statistically significant.

We conclude this section by recapitulating the main findings from these investigations:

**Highlight 7.** *Subjects who are ordinal or expected-utility maximizers in the first round or throughout the experiment tend to have a higher cognitive ability than those who are not, particularly in the (individual or combined) verbal reasoning and letter-numeric scores. Moreover, subjects who learn to be so by the last round tend to have a higher cognitive ability than those who do not, particularly in the combined letter-numeric and 3-dimensional rotation score.*

## 7.1 Demographics

We finally turn to the potential role that subjects’ demographic characteristics may have on their economic rationality and learning. Starting with their field of study, we find no significant differences in the HM scores or proportions of ordinal and expected-utility maximizers between subjects studying Economics or Business (81) and the rest (227). This is true both for the round-per-round and merged choices. Moving to the gender demographic, there were 184 female, 120 male and 4 non-binary-gender subjects. Males had slightly lower HM scores than females at each round (1.24 vs 1.55; 1.02 vs 1.3; 0.89 vs 1.23; 0.73 vs 1.12; 0.63 vs 1) and overall (1.74 vs 2.13;  $p = 0.079$ ). All round-per-round differences except the first one

were significant at the 5% level. In addition, a higher proportion of males were ordinal or expected-utility maximizers (30% vs 24% and 20% vs 16%), but insignificantly so ( $p = 0.286$  and  $p = 0.356$ , respectively). The analysis in the next subsection studies the effects of these demographic variables in more detail.

## 7.2 Regression Analysis

Table 13 shows marginal effect estimates of probit models in which the dependent variable equals 1 if the individual made choices in a round, i.e., across 15 choices, that are consistent utility maximizing (columns (1) and (2)) or with expected utility maximizing (columns (3) and (4)). The estimates reveal that—consistent with Table 6—the number of utility maximizers and expected utility maximizers increases significantly from round to round. More cognitively able individuals are more likely to be utility- or expected-utility maximizers. In particular, the verbal-reasoning scores are significant. Furthermore, men tend to be more likely to make choices consistent with utility maximization. Relative to humanities students (base category in the probit models), those studying economics and business, computer science, mathematics, psychology, or law are more likely to be utility maximizers of any kind. These results are robust to restricting the sample to subjects who never deferred a choice at any menu (columns (2) and (4)).

An analysis of response times—summarized in Table A.1 in the Appendix—aligns with, and bolsters, the findings on learning. In particular, response times become shorter from round to round, and they are shorter for individuals who have a higher verbal reasoning score. Response times are also significantly shorter for utility maximizers—and particularly shorter for expected utility maximizers. Notably, response times are longer at more difficult and more complex menus. Deferring a choice also significantly reduces response times at that menu, indicating that individuals who defer spend less time and effort to think about the choice. Interestingly, while this is not the case in the first 15 decisions, during which the average response times were similar for deferrals and active choices (21.6 vs 20.6 seconds, respectively;  $p = 0.125$  from 2-sided Mann-Whitney test), it does become so in the last 15 decisions (5.57 vs 8.2 seconds;  $p < 0.001$ ). These facts suggest that subjects spend a similar amount of time at both kinds of decisions initially, but quickly learn which menus are hard for them and then spend little time before opting for the deferral option when they see those menus again. Also interestingly, this response-time reduction when deferring a choice is higher for more cognitively able individuals, suggesting that they use deferring more efficiently and decide to defer a choice more quickly to save cognitive effort.

These findings beg the question of how these individual characteristics are related to particular violations of utility maximization. Table 14 reports marginal effects of probit models, evaluated at the means of independent variables, in which the dependent variable is an indicator variable that equals 1 if a the choices in a round of 15 menus violate First-Order Stochastic Dominance, Independence, StAR, WARP, Contraction Consistency, and

Table 13: Utility Maximizers at all menus: Probit marginal effects

	Dependent variable:			
	1 if U-maximizer across 15 menus		1 if EU-maximizer across 15 menus	
	(1)	(2)	(3)	(4)
1 if round 2	0.082*** (0.031)	0.090** (0.036)	0.042 (0.032)	0.041 (0.041)
1 if round 3	0.132*** (0.033)	0.127*** (0.038)	0.093*** (0.033)	0.081** (0.041)
1 if round 4	0.194*** (0.033)	0.209*** (0.036)	0.135*** (0.035)	0.156*** (0.043)
1 if round 5	0.247*** (0.031)	0.270*** (0.032)	0.176*** (0.034)	0.194*** (0.041)
Verbal reasoning score	0.463*** (0.162)	0.521*** (0.182)	0.269* (0.149)	0.284 (0.180)
Letter-numeric sequence score	0.062 (0.081)	0.149 (0.096)	0.045 (0.074)	0.084 (0.091)
Matrix reasoning score	0.085 (0.123)	0.011 (0.158)	-0.014 (0.105)	-0.056 (0.139)
3-dimensional rotation score	-0.045 (0.066)	-0.095 (0.076)	-0.054 (0.059)	-0.100 (0.069)
1 if female	-0.102** (0.050)	-0.101* (0.055)	-0.029 (0.044)	0.007 (0.053)
1 if gender non-binary	-0.251* (0.139)	-0.310* (0.178)	-0.074 (0.132)	-0.077 (0.162)
1 if undergraduate student	0.073 (0.149)	-0.133 (0.209)	-0.089 (0.154)	-0.327 (0.218)
1 if Master student	0.045 (0.157)	-0.107 (0.237)	-0.122 (0.128)	-0.292* (0.153)
1 if PhD student	0.252 (0.156)	0.152 (0.217)	0.161 (0.186)	0.045 (0.253)
Economics and business	0.147** (0.074)	0.193*** (0.075)	0.143** (0.071)	0.202** (0.082)
Earth Sciences and agriculture	0.120 (0.107)	0.244*** (0.093)	0.058 (0.111)	0.145 (0.144)
Physics and chemistry	0.022 (0.089)	-0.060 (0.099)	0.048 (0.078)	-0.031 (0.089)
Life science	-0.044 (0.080)	-0.030 (0.104)	-0.073 (0.068)	-0.051 (0.095)
Computer science	0.217*** (0.084)	0.261*** (0.071)	0.151** (0.077)	0.208*** (0.080)
Mathematics	0.222*** (0.076)	0.177** (0.078)	0.157** (0.078)	0.141* (0.085)
Languages	0.070 (0.107)	0.079 (0.119)	-0.024 (0.083)	-0.019 (0.099)
Psychology	0.266*** (0.096)	0.286*** (0.080)	0.184* (0.101)	0.190* (0.110)
Law	0.245** (0.103)	0.330*** (0.069)	0.168 (0.107)	0.256** (0.124)
Other field of study	0.133 (0.102)	0.059 (0.108)	0.091 (0.095)	0.049 (0.105)
Observations	1,540	1,130	1,540	1,130

**Notes:** Marginal effect estimates of probit models evaluated at the means of independent variables. Robust standard errors clustered at the subject level in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ .

In the first two columns, the dependent variable is an indicator variable that equals one if a subject made choices in the 15 menus of a round that are consistent with ordinal utility maximization. In column (1) all subjects are included, in column (2) the sample is restricted to subjects who never deferred at any of the 75 menus. In columns (3) and (4), the dependent variable is an indicator variable that equals one if a subject made choices in the 15 menus of a round that are consistent with expected utility maximization. In column (3) all subjects are included, in column (4) the sample is restricted to subjects who never deferred at any of the 75 menus.

Decisiveness, in columns (1) – (6) respectively. The estimates reveal that the probability of any violation is lower in rounds 2 to 5 than in round 1. First-order Stochastic Dominance violations become particularly less likely from round 1 to round 2, while violations of Independence and StAR become particularly less likely in the last round. WARP and Contraction Consistency violations, on the other hand, become steadily less likely from round to round. Consistent with the findings in Table 6, violations of Decisiveness are not statistically significantly different across rounds. Higher cognitive ability, and specifically better verbal reasoning ability, reduces WARP and Contraction consistency violations in particular.

## 8 Related Literature

Although, to our knowledge, no previous study has raised a similar set of questions or reported an analogous set of results, we note that Hey (2001), van de Kuilen and Wakker (2006) and Birnbaum and Schmidt (2015) have also tested aspects of the learning question that constitutes the main focus of our paper. Hey’s (2001) experiment included 53 subjects who, over the course of 5 experimental sessions, were shown five times the same 100 binary menus of lotteries with two outcomes. The five sessions were conducted on different days and with no less than two days between them, thereby enabling subjects to acquire information and experience outside the lab environment. Leaving aside the significant differences in motivation, design and sample sizes between that study and ours, no evidence that subjects learned to behave rationally over time was provided in that paper. van de Kuilen and Wakker (2006) reported on a repeated-choice experiment under risk with two treatments and 52 student subjects of various levels and fields, who could either learn “by experience and by thought” (in this treatment subjects’ played their chosen lottery after each decision<sup>20</sup>) or “only by thought”. The 26 participants in each treatment made decisions in two trial and fifteen actual rounds from two binary menus of money lotteries with two outcomes that featured “common-ratio” types of tests of Independence. Importantly, and unlike our study, the lotteries in the two menus differed in each round. The authors found that the aggregate behaviour resulting from subjects’ two decisions tended to converge to expected utility maximization in the dual-learning treatment but not in the “only by thought” one.

Following an approach which, in their own words, is a synthesis of Hey (2001) and van de Kuilen and Wakker (2006), Birnbaum and Schmidt (2015) recruited 54 mainly economics and business undergraduate student subjects and presented them four times with the same 20 binary menus of money lotteries. These menus were designed to test Coalescing (splitting vs non-splitting an outcome’s probability should not alter choices between otherwise identical lotteries), Independence (in their case, the “common-ratio” and “common-consequence” implications thereof), and risk-attitude inconsistencies (manifested in that study when choices between the same risky and safe lotteries are reversed). The authors found evidence of “by thought” learning in all three dimensions, as evidenced by the significant decrease in the

---

<sup>20</sup>A recent survey of the literature on learning by experience in risky choice is Hertwig and Erev (2009).

Table 14: Violations of Utility Maximization at All Menus: Probit Marginal Effects

	Dependent variable: 1 if violating (across 15 menus)					
	FOSD	Independence	StAR	WARP	Contraction consistency	Decisiveness
	(1)	(2)	(3)	(4)	(5)	(6)
1 if round 2	-0.015** (0.007)	-0.043 (0.032)	-0.021 (0.034)	-0.120*** (0.030)	-0.127*** (0.030)	-0.013 (0.018)
1 if round 3	-0.014** (0.006)	-0.034 (0.032)	-0.052* (0.031)	-0.165*** (0.031)	-0.150*** (0.032)	-0.011 (0.018)
1 if round 4	-0.021*** (0.006)	-0.040 (0.034)	-0.037 (0.032)	-0.222*** (0.030)	-0.227*** (0.030)	-0.008 (0.018)
1 if round 5	-0.019*** (0.006)	-0.069** (0.031)	-0.103*** (0.030)	-0.291*** (0.027)	-0.293*** (0.027)	-0.014 (0.018)
Verbal reasoning score	-0.005 (0.025)	0.053 (0.110)	-0.232* (0.125)	-0.532*** (0.156)	-0.499*** (0.157)	-0.031 (0.112)
Letter-numeric sequence score	-0.027** (0.013)	-0.041 (0.057)	-0.013 (0.067)	-0.067 (0.081)	-0.057 (0.080)	0.038 (0.056)
Matrix reasoning score	0.002 (0.022)	0.029 (0.080)	0.133 (0.090)	0.001 (0.121)	-0.020 (0.118)	-0.102 (0.089)
3-dimensional rotation score	0.008 (0.014)	0.055 (0.046)	0.080 (0.052)	0.079 (0.065)	0.052 (0.065)	-0.058 (0.047)
1 if female	0.019* (0.010)	0.059* (0.035)	-0.009 (0.041)	0.078 (0.049)	0.079 (0.049)	0.018 (0.036)
1 if gender non-binary	0.278 (0.173)	0.010 (0.113)	-0.014 (0.164)	0.325** (0.142)	0.309** (0.143)	-0.079 (0.055)
1 if undergraduate student	0.219*** (0.044)	0.085 (0.124)	-0.058 (0.126)	-0.049 (0.134)	-0.033 (0.132)	-0.091 (0.121)
1 if Master student	0.976*** (0.016)	0.153 (0.156)	-0.055 (0.121)	-0.108 (0.135)	-0.089 (0.136)	-0.006 (0.100)
1 if PhD student	0.986*** (0.001)	0.019 (0.157)	-0.216*** (0.080)	-0.224* (0.127)	-0.216 (0.132)	-0.051 (0.083)
Economics and business	-0.014 (0.011)	0.033 (0.053)	-0.088 (0.055)	-0.101 (0.070)	-0.111 (0.070)	-0.050 (0.046)
Earth Sciences and agriculture	0.006 (0.022)	-0.035 (0.068)	-0.051 (0.077)	-0.132 (0.090)	-0.105 (0.096)	0.015 (0.072)
Physics and chemistry	-0.026*** (0.007)	0.008 (0.059)	-0.044 (0.064)	0.043 (0.091)	0.037 (0.090)	-0.045 (0.048)
Life science	-0.015* (0.009)	-0.007 (0.052)	0.049 (0.063)	0.054 (0.082)	0.044 (0.081)	-0.006 (0.051)
Computer science	-0.012 (0.011)	0.083 (0.066)	-0.128** (0.065)	-0.185** (0.075)	-0.183** (0.077)	-0.060 (0.045)
Mathematics		-0.017 (0.066)	-0.109* (0.058)	-0.145** (0.073)	-0.146** (0.074)	-0.096*** (0.034)
Languages	-0.027*** (0.007)	0.070 (0.072)	-0.003 (0.090)	0.033 (0.106)	0.005 (0.106)	-0.054 (0.050)
Psychology	-0.016 (0.011)	-0.018 (0.068)	-0.175*** (0.062)	-0.194** (0.089)	-0.185** (0.091)	-0.085** (0.040)
Law	-0.008 (0.015)	0.043 (0.080)	-0.135* (0.078)	-0.271*** (0.074)	-0.265*** (0.077)	0.001 (0.077)
Other field of study	-0.027*** (0.007)	0.020 (0.068)	-0.069 (0.071)	0.001 (0.104)	-0.019 (0.103)	-0.153*** (0.020)
Observations	1,375	1,540	1,540	1,540	1,540	1,540

**Notes:** Marginal effect estimates of probit models evaluated at the means of independent variables. Robust standard errors clustered at the subject level in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ .

The dependent variable in column (1) is an indicator variable that equals one if a subject makes at least one first-order stochastically dominated choice. Likewise the dependent variables in columns (2) to (6) are indicator variables that equal 1 if a subject's choices in a round of 15 menus violate Independence, StAR, WARP, Contraction consistency, and Decisiveness, respectively.

respective total violations.

Compared to these earlier studies, ours differs in several important ways: (i) it features *free choices*; (ii) presents both binary and non-binary menus; (iii) was designed to test several implications of deterministic ordinal- and expected-utility maximization that go beyond the Independence axiom; (iv) was also designed to test implications of *random* utility maximization; (v) uses rather involved computational methods to assess with precision each subject’s conformity with ordinal- and expected-utility maximization, as well as with each one of several behavioural axioms implied by these or by random utility theories; (vi) has 6 times as large a sample; (vii) analyses the data both at the individual and aggregate levels; (viii) accounts for potential indifferences in the former type of analysis and quantifies the importance of doing so; (ix) finds significant evidence of learning without feedback in what we believe were considerably more challenging decision environments; and (x) relates subjects’ overall consistency and (non-) learning to cognitive ability.

Our paper also differs significantly from the recent studies by Art and Haruvy (2017), Charness et al. (2023); Charness and Chemaya (2023), Nielsen and Rehbeck (2022), Benjamin et al. (2021), and Breig and Feldman (2024). A common feature in the latter three works is the finding that subjects show a higher conformity with principles of rational choice under risk when they are asked if they want to revise their previous choices, with or without receiving relevant feedback prior to the revision opportunity. Unlike our experiment, the ones in these studies were not designed to test for the existence and evolution of subjects’ neutral learning (i.e., only “by-thought”) to be (expected-)utility maximizers, either with or without also accounting for the possibility of them being occasionally indifferent. Art and Haruvy (2017) (60 subjects) and, concurrent to our own study, Charness et al. (2023) and Charness and Chemaya (2023) (99 subjects) on the other hand studied stability of risk preferences using several rounds—with feedback and trialling between rounds—of the multiple price list format due to Holt and Laury (2002) and a variation of the one-menu-with-six-lotteries method due to Eckel and Grossman (2002), respectively. All lotteries in these studies featured two monetary outcomes. The first study found generally unstable preferences between rounds, and subjects’ tendency to become more risk neutral with experience. Charness and Chemaya (2023), analysing the data from Charness et al. (2023), also found more than 50% of subjects changing their choices between the first and last decisions that preceded and succeeded the intermediating trial/feedback rounds. Consistent with the main message that comes from our own analysis that builds on substantially different experimental procedures and data-analytic methods, no feedback, free choices, and more than three times their sample size, Charness and Chemaya (2023) highlight the importance of providing participants with experience in risk-elicitation tasks towards reducing preference-measurement errors.

Although this is outwith the focus of our study, finally, we note that there is an extensive body of research that documents subjects’ learning to conform with various types of equilibrium predictions in strategic games as they become more experienced. In ultimatum bargaining, for example, Slonim and Roth (1998), followed by List and Cherry (2000)



and Grimm and Mengel (2011), among others, find evidence in favour of such learning by proposers and responders, respectively. In the case of the last paper, moreover, this occurred without play repetitions and outcome disclosures, which might be viewed as a form of feedback; instead, it occurred by imposing a short delay before responders got to respond.

## 9 Summary and Discussion

Our aim in this study has been to conduct a detailed and targeted investigation of behaviour in experimental choices under risk when all decision problems are presented to subjects multiple times (in our case, five) and in a carefully structured semi-random order. The lotteries and decision problems were designed to test subjects' conformity with a variety of ordinal- and expected-utility maximization principles, including ones that pertain to binary menus only as well as ones that require revealed-preference consistency across binary and/or non-binary menus. The data collected from 308 subjects who participated in our free-/non-forced choice experiment in the UK and in Germany allowed us to carry out rather comprehensive and, we hope, illuminating tests of some important questions that are relevant for our understanding of choice under risk, which underpins much of real-world economic decision-making, and for the design of experiments or questionnaires that aim to elicit risk preferences:

1. Do subjects learn to maximize ordinal or expected utility as they progress in the experiment, without receiving any feedback or invitations to revise their previous choices, and without being forced to make active choices?
2. Can proper accounting for the possibility of subjects' being indifferent between some lotteries partly explain the often observed choice reversals across different instances where the same menu was presented to them?
3. Which principles of utility maximization appear to be the hardest for subjects to (learn to) conform with?
4. What is the relation between cognitive ability and utility maximization or learning thereof?

In response to the first question, our analysis suggests that a substantial fraction of participants in choice experiments can learn, without any feedback or other interventions, to conform in a strict sense with the benchmark models of economic rationality when they are repeatedly exposed to the same decision problems, even when several of these problems involve relatively high degrees of decision difficulty. In our data, those fractions were 23% and 12% for ordinal- and expected-utility maximization between the first and last decision round, respectively, representing highly significant increases in theory-abiding subjects. This conclusion is in contrast to those in some—and in line to those in others—pre-existing studies that explored similar themes using more constrained analytical methods and experimental designs, smaller sample sizes and, typically, fewer repetitions of the same menus.

Our analysis pertaining to the second question uncovers that choice reversals between different presentations of the same decision problem do indeed often stem from the decision makers' rational indifference between the respective choice alternatives. More specifically, 15% of all subjects in our sample exhibited behaviour across their 75 decisions that is perfectly consistent with indifference-revealing ordinal utility maximization, with half of them also being perfectly consistent with expected-utility maximization. Without testing for the indifference hypothesis, this sizeable proportion of subjects would have been discarded as non-rational. While these possibilities had been acknowledged in the related literature as early as Davidson and Marschak (1959), our study appears to be the first to document and quantify them, and it does so using a sophisticated combinatorial-optimization method that is freely accessible. Furthermore, of potential interest to both empirical economists and behavioural decision theorists is our test for and detection of substantial overlaps between those subjects who are classified as deterministic utility maximizers of either type (including those with some indifference ties) in their overall decisions and subjects who conform with the five implications of random utility theory that we examined. This newly discovered association, in particular, suggests that the decision analyst may have more degrees of freedom than previously thought in interpreting how individuals behave in experiments that involve repeated presentations of the same problems.

As far as specific rationality principles are concerned, analysing subjects' deferring behaviour (equivalently, conformity with the Decisiveness axiom) in our free-/non-forced choice experiment reveals that such decisions are made by a sizeable minority of 15-16% participants, and they are persistent across rounds and relatively predictable in their occurrence. Indeed, in line with intuition and the decision-difficulty theoretical channel to choice avoidance/deferral that has been suggested and documented in the existing literature, deferrals tend to occur at menus that are more complex than others, either because they lack a (first-/second-order) stochastically dominant lottery or because, in addition, they present information in a complicated way. Considering that deferring in this experiment comes with a positive expected cost for subjects, the ensuing violations of the Decisiveness axiom cannot be attributed to indifference but would be better thought of as being caused by incomplete/imprecise preferences or complexity-aversion considerations.

Subjects' violations of all other axioms exhibited a steady decline over the course of the experiment. Notably, however, despite there being only 1 and 3 pairs of menus, respectively, where violations in the Independence and Stability of Attitudes to Risk (StAR) axioms could be observed in any one round (as opposed to 20 such pairs for the Contraction Consistency axiom, for example), 25% and 26% of all subjects still violated the strict-preference versions of these axioms in the fifth round. Unlike many tests of Independence in the literature that revolve around patterns inspired by the Allais-paradox, ours does not involve probabilities near/equal to one or zero in any of the lotteries involved. Instead, it features the novelty whereby each of the relevant two binary menus contains lotteries that are not ranked by second-order stochastic dominance, thereby representing a decision with difficult trade-offs.

Introducing and testing StAR, moreover, enabled us to uncover patterns of risk-attitude reversals that go beyond those in the classic “reflection effect” (Tversky and Kahneman, 1981) where the said reversals are mediated by the framing of decision in terms of gains or losses. More specifically, our test of StAR relies on checking whether a decision maker always opts for the second-order stochastically dominant (risk-averse) or dominated (risk-seeking) lottery at every collection of pairs that contain such a lottery and have an overlapping set of monetary outcomes. The vast majority of subjects respect this axiom, and do so in the direction of exhibiting consistent risk aversion. We note the possibility that those who violate it in any given round do so because they are risk-neutral, a hypothesis that cannot be definitively rejected in an experiment with single-valued choice data from individual rounds. Moreover, the relatively limited variation in the expected values of the lotteries in our specific experiment poses an obstacle toward testing for the possibility that some subjects start comparing lotteries by looking at their expected values before turning to other criteria. The risk-neutrality hypothesis, however, *can be tested* in the indifference-permitting analysis where the possibly distinct single choices per menu across different rounds are merged. There, risk neutrality would be detected by subjects’ revealed indifference in *every* binary menu that contained SOSD-ranked lotteries. Yet only one expected-utility maximizing subject was revealed to be risk-neutral thus. In light of this fact, StAR violations suggest that context-dependent risk attitudes exist even in without gain/loss framing. In particular, they also emerge when SOSD-ranked lotteries feature the same or similar three or more outcomes. This, in conjunction with our other findings, points to a potentially relevant descriptive role for the development of theories of choice under risk that allow for similarity-based reversals in attitudes to risk while simultaneously respecting first-order stochastic dominance and the basic axioms of ordinal utility theory.

Concerning the role of cognitive ability, finally, our results indicate that not only is it related to overall choice consistency, but it also plays a predictive role in individuals’ capacity to autonomously adapt towards rational decision-making over the course of the experiment. In particular, we find an important association between choice consistency and learning on the one hand and the verbal-reasoning and letter-numeric tasks of the administered ICAR-16 on the other. Assuming that the latter causes the former, which is something that we are clearly unable to test, this suggests that efforts to improve people’s decision-making quality in the real-world, for example their financial literacy, can benefit from the inclusion of not only numerical problems but also of verbal ones.

The clear presence of feedback-independent learning in our data carries important implications for experimental design, theory testing and preference elicitation. Indeed, it suggests that in conducting choice experiments or surveys where participants encounter the same scenarios repeatedly, focusing on subjects’ decisions in the final instance of these scenarios, and properly accounting for the possibility of indifferences across all their decisions, could yield significantly more accurate information about the subjects’ underlying decision process and preferences, whether these were “discovered” (Plott, 1996) or “constructed” (Kahneman,

1996). This postulated more accurate elicitation could in turn—as in our study—paint a relatively more favourable picture of the baseline models of economic rationality as descriptive theories of choice under risk than what is often inferred. Indeed, combining our answers to the first two questions above, as these are summarized in the relevant entries of Tables 3 and 6, leads to the conclusion that 51% and 36%, of all subjects, respectively, are revealed to be ordinal or expected-utility maximizers either across all 75 decisions or in their last 15 ones, with 6.5% and 3% of the indifference-revealing subjects in these groups breaking their indifference ties in the last round. While our experiment is not directly comparable to any pre-existing one that we are aware of, it is probably fair to say that such a degree of conformity with those two models is considerably higher than what one might have anticipated in the decision environment of our experiment.

This discussion suggests that learning without feedback or other exogenous interventions—such as choice-revision queries—raises the possibility that an additional, neutral and cost-effective way towards a meaningful reduction of measurement error in risk-preference elicitation in the lab or in the field (Schildberg-Hörisch, 2018; Gillen et al., 2019; Dohmen and Jagelka, 2024) could be the wider use of appropriately structured repeated-choice experiments or surveys. In particular, the set of findings presented earlier allows us to forcefully echo the claims made in Charness and Chemaya (2023) and Breig and Feldman (2024) concerning a better experimental-economics practice where the analyst focuses on subjects’ later choices in order to understand their “true” preferences. In determining the number of repetitions, of course, the researcher’s challenge is to strike the right balance between enabling subjects to learn and reveal this learning via their choices while preventing them from becoming fatigued and revealing inaccurate information instead.

The substantial evidence in favour of the learning hypothesis notwithstanding, non-trivial proportions of subjects in our study still deviated from expected or even ordinal utility maximization by the end of the experiment, even after accounting for the possibility of indifference. This fact reinforces the widely held belief that favours the development of deterministic and stochastic models of bounded-rational choice under risk. It is outside this paper’s scope to explore which of the numerous existing such models might explain those subjects’ behaviour better or to provide detailed outlines of potentially new models that might do so. Our approximate-rationality analysis for the ordinal model does suggest that the vast majority of subjects make up to one “mistaken” choice by the fifth time they are asked to decide from the same menus. This fact and our remarks above about the potential role of similarity-driven risk-attitude reversals point towards a potentially promising avenue in this respect. We hope that the rich new dataset that we are introducing with this study will facilitate further empirical and theoretical exploration of these important questions.

## References

- AGRANOV, M. AND P. ORTOLEVA (2017): “Stochastic Choice and Preferences for Randomization,” *Journal of Political Economy*, 125, 40–68.
- ALLAIS, M. (1953): “Le Comportement de l’Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l’École Américaine,” *Econometrica*, 21, 503–546.
- ALÓS-FERRER, C., E. FEHR, H. FEHR-DUDA, AND M. GARAGNANI (2024): “Common Ratio and Common Consequence Effects Arise from True Preferences,” *Working Paper*.
- ALÓS-FERRER, C., E. FEHR, AND N. NETZER (2021): “Time Will Tell: Recovering Preferences When Choices Are Noisy,” *Journal of Political Economy*, 129, 1828–1877.
- APESTEGUIA, J. AND M. A. BALLESTER (2015): “A Measure of Rationality and Welfare,” *Journal of Political Economy*, 123, 1278–1310.
- (2018): “Monotone Stochastic Choice Models: The Case of Risk and Time Preferences,” *Journal of Political Economy*, 126, 74–106.
- (2021): “Separating Predicted Randomness from Residual Behavior,” *Journal of the European Economic Association*, 19, 1041–1076.
- ART, E. AND E. HARUVY (2017): “Revisiting Risk Aversion: Can Risk Preferences Change with Experience?” *Economics Letters*, 151, 91–95.
- AUMANN, R. J. (1962): “Utility Theory without the Completeness Axiom,” *Econometrica*, 30, 445–462.
- BALAKRISHNAN, N., E. A. OK, AND P. ORTOLEVA (2021): “Inferential Choice Theory,” *Working Paper*.
- BARSEGHYAN, L. AND F. MOLINARI (2023): “Risk Preference Types, Limited Consideration, and Welfare,” *Journal of Business and Economic Statistics*, 41, 1011–1029.
- BARSEGHYAN, L., F. MOLINARI, AND M. THIRKETTLE (2021): “Discrete Choice under Risk with Limited Consideration,” *American Economic Review*, 111, 1972–2006.
- BECKER, A., T. DECKERS, T. DOHMEN, A. FALK, AND F. KOSSE (2012): “The Relationship Between Economic Preferences and Psychological Personality Measures,” *Annual Review of Economics*, 4, 453–478.
- BENJAMIN, D. J., M. A. FONTANA, AND M. S. KIMBALL (2021): “Reconsidering Risk Aversion,” *NBER Working Paper 28007*.
- BIRNBAUM, M. H. AND U. SCHMIDT (2015): “The Impact of Learning by Thought on Violations of Independence and Coalescing,” *Decision Analysis*, 12, 144–152.

- BLOCK, H. D. AND J. MARSCHAK (1960): “Random Orderings and Stochastic Theories of Response,” in *Contributions to Probability and Statistics: Essays in Honor of Harold Hotelling*, ed. by I. Olkin, Stanford, CA: Stanford University Press, 97–132.
- BOCK, O., I. BAETGE, AND A. NICKLISCH (2014): “hroot: Hamburg Registration and Organization Online Tool,” *European Economic Review*, 71, 117–120.
- BORDALO, P., N. GENNAIOLI, AND A. SHLEIFER (2012): “Salience Theory of Choice Under Risk,” *Quarterly Journal of Economics*, 127, 1243–1285.
- BREIG, Z. AND P. FELDMAN (2024): “Revealing Risky Mistakes through Revisions,” *Journal of Risk and Uncertainty*, 68, 227–254.
- CAPLIN, A. (2025): “Data Engineering for Cognitive Economics,” *Journal of Economic Literature*, 63, 164–196.
- CERREIA-VIOGLIO, S., D. DILLENBERGER, AND P. ORTOLEVA (2015): “Cautious Expected Utility and the Certainty Effect,” *Econometrica*, 83, 693–728.
- CERREIA-VIOGLIO, S., D. DILLENBERGER, P. ORTOLEVA, AND G. RIELLA (2019): “Deliberately Stochastic,” *American Economic Review*, 109, 2425–2445.
- CHAPMAN, J., M. DEAN, P. ORTOLEVA, E. SNOWBERG, AND C. CAMERER (2023): “Econographics,” *Journal of Political Economy: Microeconomics*, 1, 115–161.
- CHARNESS, G. AND N. CHEMAYA (2023): “Repeated Experience and Consistent Risk Preferences,” *Economics Letters*, 233, 111375.
- CHARNESS, G., N. CHEMAYA, AND D. TRUJANO-OCHOA (2023): “Learning Your Own Risk Preferences,” *Journal of Risk and Uncertainty*, 67, 1–19.
- CHERNEV, A., U. BÖCKENHOLT, AND J. GOODMAN (2015): “Choice Overload: A Conceptual Review and Meta-Analysis,” *Journal of Consumer Psychology*, 25, 333–358.
- CHERNOFF, H. (1954): “Rational Selection of Decision Functions,” *Econometrica*, 4, 422–443.
- CONDON, D. M. AND W. REVELLE (2014): “The International Cognitive Ability Resource: Development and Initial Validation of a Public-Domain Measure,” *Intelligence*, 43, 52–64.
- COSTA-GOMES, M., C. CUEVA, G. GERASIMOU, AND M. TEJIŠČÁK (2022): “Choice, Deferral and Consistency,” *Quantitative Economics*, 13, 1297–1318.
- CUBITT, R., D. NAVARRO-MARTINEZ, AND C. STARMER (2015): “On Preference Imprecision,” *Journal of Risk and Uncertainty*, 50, 1–34.
- DANAN, E. AND A. ZIEGELMEYER (2006): “Are Preferences Complete? An Experimental Measurement of Indecisiveness Under Risk,” *Working Paper*.

- DAVIDSON, D. AND J. MARSCHAK (1959): “Experimental Tests of a Stochastic Decision Theory,” in *Measurement: Definitions and Theories*, ed. by C. W. Churchman and P. Ratoosh, New York: Wiley, 233–269.
- DEAN, M. AND N. NELIGH (2023): “Experimental Tests of Rational Inattention,” *Journal of Political Economy*, 131, 3415–3461.
- DEAN, M., D. RAVINDRAN, AND J. STOYE (2022): “A Better Test of Choice Overload,” *arXiv:2212.03931*.
- DEMBO, A., S. KARIV, M. POLISSON, AND J. K.-H. QUAH (2021): “Ever Since Allais,” *Working Paper*.
- DHAR, R. (1997): “Consumer Preference for a No-Choice Option,” *Journal of Consumer Research*, 24, 215–231.
- DHAR, R. AND I. SIMONSON (2003): “The Effect of Forced Choice on Choice,” *Journal of Marketing Research*, 40, 146–160.
- DOHMEN, T., A. FALK, D. HUFFMAN, AND U. SUNDE (2010): “Are Risk Aversion and Impatience Related to Cognitive Ability?” *American Economic Review*, 100, 1238–1260.
- (2018): “On the Relationship between Cognitive Ability and Risk Preference,” *Journal of Economic Perspectives*, 32, 115–34.
- DOHMEN, T. AND T. JAGELKA (2024): “Accounting for Individual-Specific Reliability of Self-Assessed Measures of Economic Preferences and Personality Traits,” *Journal of Political Economy: Microeconomics*, forthcoming.
- ECHENIQUE, F., T. IMAI, AND K. SAITO (2023): “Approximate Expected Utility Rationalization,” *Journal of the European Economic Association*, 21, 1821–1864.
- ECKEL, C. C. AND P. J. GROSSMAN (2002): “Sex Differences and Statistical Stereotyping in Attitudes Toward Financial Risk,” *Evolution and Human Behavior*, 23, 281–295.
- FALMAGNE, J.-C. (2002): *Elements of Psychophysical Theory*, Oxford: Oxford University Press.
- FISHBURN, P. C. (1978): “Choice Probabilities and Choice Functions,” *Journal of Mathematical Psychology*, 18, 205–219.
- FRIEDMAN, M. AND L. J. SAVAGE (1948): “Utility Analysis of Choices Involving Risk,” *Journal of Political Economy*, 56, 279–304.
- GANSNER, E. R. AND S. C. NORTH (2000): “An Open Graph Visualization System and its Applications to Software Engineering,” *Software: Practice and Experience*, 30, 1203–1233.
- GERASIMOU, G. (2018): “Indecisiveness, Undesirability and Overload Revealed Through Rational Choice Deferral,” *Economic Journal*, 128, 2450–2479.

- (2021): “Model-Rich Approaches to Eliciting Possibly Weak or Incomplete Preferences: Evidence from a Multi-Valued Choice Experiment,” *arXiv:2111.14431*.
- GERASIMOU, G. AND M. TEJIŠČÁK (2018): “Prest: Open-Source Software for Computational Revealed Preference Analysis,” *Journal of Open Source Software*, 3, <https://doi.org/10.21105/joss.01015>, available at <https://prestsoftware.com>.
- GILL, D. AND V. PROWSE (2016): “Cognitive Ability, Character Skills, and Learning to Play Equilibrium: A Level-k Analysis,” *Journal of Political Economy*, 124, 1619–1676.
- GILL, D. AND Y. ROSOKHA (2024): “Beliefs, Learning, and Personality in the Indefinitely Repeated Prisoner’s Dilemma,” *American Economic Journal: Microeconomics*, 16, 259–283.
- GILLEN, B., E. SNOWBERG, AND L. YARIV (2019): “Experimenting with Measurement Error: Techniques with Applications to the Caltech Cohort Study,” *Journal of Political Economy*, 127, 1826–1863.
- GREINER, B. (2015): “Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE,” *Journal of the Economic Science Association*, 1, 114–125.
- GRIMM, V. AND F. MENGEL (2011): “Let Me Sleep on It: Delay Reduces Rejection Rates in Ultimatum Games,” *Economics Letters*, 111, 113–115.
- GUL, F. AND W. PESENDORFER (2006): “Random Expected Utility,” *Econometrica*, 74, 121–146.
- HADAR, J. AND W. R. RUSSELL (1969): “Rules for Ordering Uncertain Prospects,” *American Economic Review*, 59, 25–34.
- HANOCH, G. AND H. LEVY (1969): “The Efficiency Analysis of Choices Involving Risk,” *Review of Economic Studies*, 36, 335–346.
- HARLESS, D. W. AND C. F. CAMERER (1994): “The Predictive Utility of Generalized Expected Utility Theories,” *Econometrica*, 62, 1251–1289.
- HE, J. AND P. NATENZON (2024): “Moderate Utility,” *American Economic Review: Insights*, 6, 176–195.
- HERTWIG, R. AND I. EREV (2009): “The Description–Experience Gap in Risky Choice,” *Trends in Cognitive Sciences*, 13, 517–523.
- HEY, J. D. (2001): “Does Repetition Improve Consistency?” *Experimental Economics*, 4, 5–54.
- HEY, J. D. AND C. ORME (1994): “Investigating Generalizations of Expected Utility Theory Using Experimental Data,” *Econometrica*, 62, 1291–1326.
- HOLT, C. A. AND S. K. LAURY (2002): “Risk Aversion and Incentive Effects,” *American Economic Review*, 92, 1644–1655.
- HOUTMAN, M. AND J. A. MAK (1985): “Determining All Maximal Data Subsets Consistent with Revealed Preference,” *Kwantitatieve Methoden*, 19, 89–104.



- HURWICZ, L. (1986): “On the Implementation of Social Choice Rules in Irrational Societies,” in *Essays in Honor of Kenneth J. Arrow*, ed. by W. P. Heller, R. M. Starr, and D. A. Starett, Cambridge University Press, vol. I, chap. 4.
- IYENGAR, S. S. AND M. R. LEPPER (2000): “When Choice is Demotivating: Can One Desire Too Much of a Good Thing?” *Journal of Personality and Social Psychology*, 79, 995–1006.
- KAHNEMAN, D. (1996): “Comment on ‘Rational Individual Behavior in Markets and Social Choice Processes: the Discovered Preference Hypothesis’ by Charles R. Plott,” in *The Rational Foundations of Economic Behavior*, ed. by K. J. Arrow, E. Colombatto, M. Perlman, and C. Schmidt, Macmillan, chap. 10, 251–254.
- KAHNEMAN, D. AND A. TVERSKY (1979): “Prospect Theory: An Analysis of Decision under Risk,” *Econometrica*, 47, 263–291.
- KREPS, D. M. (1990): *A Course in Microeconomic Theory*, Harlow, UK: Prentice Hall.
- (2012): *Microeconomic Foundations I: Choice and Competitive Markets*, Princeton: Princeton University Press.
- LEVY, H. (2008): “First-Degree Stochastic Dominance Violations: Decision Weights and Bounded Rationality,” *Economic Journal*, 118, 759–774.
- LIST, J. A. AND T. L. CHERRY (2000): “Learning to Accept in Ultimatum Games: Evidence from an Experimental Design that Generates Low Offers,” *Experimental Economics*, 3, 11–29.
- LUCE, R. D. AND H. RAIFFA (1957): *Games and Decisions*, New York: Dover.
- MACHINA, M. J. (1985): “Stochastic Choice Functions Generated from Deterministic Preferences over Lotteries,” *Economic Journal*, 95, 575–594.
- MARSCHAK, J. (1960): “Binary Choice Constraints on Random Utility Indications,” in *Stanford Symposium on Mathematical Methods in the Social Sciences*, ed. by K. Arrow, Stanford: Stanford University Press, 312–329.
- MAS-COLELL, A., M. D. WHINSTON, AND J. R. GREEN (1995): *Microeconomic Theory*, New York: Oxford University Press.
- MCGRANAGHAN, C., K. NIELSEN, T. O’DONOGHUE, J. SOMERVILLE, AND C. D. SPRENGER (2024): “Distinguishing Common Ratio Preferences from Common Ratio Effects Using Paired Valuation Tasks,” *American Economic Review*, 114, 307–347.
- NIELSEN, K. AND J. REHBECK (2022): “When Choices Are Mistakes,” *American Economic Review*, 112, 2237–2268.
- NIELSEN, K. AND L. RIGOTTI (2022): “Revealed Incomplete Preferences,” *arXiv:2205.08584*.
- PLOTT, C. R. (1996): “Rational Individual Behavior in Markets and Social Choice Processes: the Discovered Preference Hypothesis,” in *The Rational Foundations of Economic Behavior*, ed. by K. J. Arrow, E. Colombatto, M. Perlman, and C. Schmidt, Macmillan, chap. 10, 225–250.

- PROTO, E., A. RUSTICHINI, AND A. SOFIANOS (2019): “Intelligence, Personality, and Gains from Cooperation in Repeated Interactions,” *Journal of Political Economy*, 127, 1351–1390.
- (2022): “Intelligence, Errors and Strategic Choices in the Repeated Prisoners’ Dilemma,” *Review of Economic Studies*, 89, 2723–2767.
- ROTHSCHILD, M. AND J. STIGLITZ (1970): “Increasing Risk. I. A Definition,” *Journal of Economic Theory*, 2, 225–243.
- SCHEIBEHENNE, B., R. GREIFENEDER, AND P. M. TODD (2010): “Can There Ever Be Too Many Options? A Meta-Analytic Review of Choice Overload,” *Journal of Consumer Research*, 37, 409–425.
- SCHILDBERG-HÖRISCH, H. (2018): “Are Risk Preferences Stable?” *Journal of Economic Perspectives*, 32, 135–154.
- SEN, A. (1997): “Maximization and the Act of Choice,” *Econometrica*, 65, 745–779.
- SEN, A. K. (1971): “Choice Functions and Revealed Preference,” *Review of Economic Studies*, 38, 307–317.
- SLONIM, R. AND A. E. ROTH (1998): “Learning in High Stakes Ultimatum Games: An Experiment in the Slovak Republic,” *Econometrica*, 66, 569–596.
- STRZALECKI, T. (2025): *Stochastic Choice Theory*, Econometric Society Monographs, Cambridge: Cambridge University Press.
- THURSTONE, L. L. (1927): “A Law of Comparative Judgement,” *Psychological Review*, 34, 278–286.
- TVERSKY, A. (1969): “Intransitivity of Preferences,” *Psychological Review*, 76, 31–48.
- TVERSKY, A. AND D. KAHNEMAN (1981): “The Framing of Decisions and the Psychology of Choice,” *Science*, 211, 453–458.
- TVERSKY, A. AND E. SHAFIR (1992): “Choice under Conflict: The Dynamics of Deferred Decision,” *Psychological Science*, 3, 358–361.
- VAN DE KUILEN, G. AND P. P. WAKKER (2006): “Learning in the Allais Paradox,” *Journal of Risk and Uncertainty*, 33, 155–164.
- VON NEUMANN, J. AND O. MORGENSTERN (1947): *Theory of Games and Economic Behavior*, Princeton: Princeton University Press, 2nd ed.

# Appendix

## A Analysis of Response Times

Table A.1: Response Times

	Dependent variable:		
	Response time at choice menu in seconds		
	(1)	(2)	(3)
1 if EU maximizer			-1.298** (0.643)
1 if utility maximizer			-1.774** (0.685)
1 if choice deferred		-2.976* (1.519)	-4.331*** (1.522)
1 if round 2	-7.392*** (0.339)	-7.394*** (0.339)	-7.215*** (0.343)
1 if round 3	-9.924*** (0.427)	-9.926*** (0.427)	-9.608*** (0.426)
1 if round 4	-11.227*** (0.466)	-11.226*** (0.467)	-10.752*** (0.468)
1 if round 5	-12.514*** (0.497)	-12.514*** (0.497)	-11.899*** (0.497)
1 if menu A1,A2	-4.671*** (0.516)	-5.017*** (0.534)	-5.175*** (0.533)
1 if menu A1,A2,C1	-3.703*** (0.508)	-4.047*** (0.534)	-4.203*** (0.533)
1 if menu A1,A2,C1,C2	-4.705*** (0.525)	-5.051*** (0.552)	-5.208*** (0.551)
1 if menu A1,A2,C2	-5.748*** (0.539)	-6.088*** (0.568)	-6.243*** (0.568)
1 if menu A1,B1	-6.464*** (0.500)	-6.804*** (0.530)	-6.959*** (0.530)
1 if menu A1,B1,B2	-2.412*** (0.491)	-2.740*** (0.536)	-2.890*** (0.539)
1 if menu A1,B1,B2,D	0.753 (0.557)	0.430 (0.590)	0.283 (0.594)
1 if menu A1,B2	-7.736*** (0.523)	-8.076*** (0.550)	-8.231*** (0.548)
1 if menu A1,D	-6.633*** (0.512)	-6.979*** (0.546)	-7.136*** (0.545)
1 if menu A2,D	-7.111*** (0.526)	-7.445*** (0.554)	-7.598*** (0.554)
1 if menu B1,B2	-1.580*** (0.436)	-1.837*** (0.458)	-1.954*** (0.458)
1 if menu B1,B2,D	-2.185*** (0.494)	-2.512*** (0.530)	-2.660*** (0.530)
1 if menu B1,D	-4.978*** (0.468)	-5.312*** (0.498)	-5.465*** (0.495)
1 if menu B2,D	-6.461*** (0.519)	-6.782*** (0.553)	-6.928*** (0.552)

Verbal reasoning score	-4.371** (1.967)	-4.270** (1.947)	-3.198* (1.901)
Letter-numeric sequence score	0.033 (1.000)	0.069 (0.995)	0.228 (0.959)
Matrix reasoning score	0.332 (1.706)	0.305 (1.706)	0.393 (1.653)
3-dimensional rotation score	0.867 (0.814)	0.868 (0.813)	0.727 (0.768)
1 if female	-1.376** (0.687)	-1.354** (0.683)	-1.555** (0.669)
1 if gender non-binary	-3.218*** (1.195)	-3.212*** (1.195)	-3.743*** (1.197)
1 if undergraduate student	-1.784 (1.901)	-1.798 (1.885)	-1.806 (1.791)
1 if Master student	-2.466 (2.040)	-2.372 (2.020)	-2.428 (1.938)
1 if PhD student	-1.601 (2.342)	-1.584 (2.329)	-0.978 (2.207)
Economics and business	1.246 (1.065)	1.270 (1.063)	1.677 (1.031)
Earth Sciences and agriculture	-1.075 (1.086)	-1.049 (1.084)	-0.781 (1.064)
Physics and chemistry	0.094 (1.218)	0.070 (1.216)	0.143 (1.198)
Life science	-1.248 (0.967)	-1.252 (0.961)	-1.408 (0.946)
Computer science	-1.262 (1.049)	-1.273 (1.045)	-0.760 (1.074)
Mathematics	-0.053 (1.148)	-0.085 (1.153)	0.448 (1.111)
Languages	1.335 (1.590)	1.304 (1.590)	1.383 (1.513)
Psychology	1.172 (1.293)	1.131 (1.295)	1.773 (1.224)
Law	-1.870* (1.061)	-1.770* (1.037)	-1.127 (1.029)
Other field of study	-0.822 (1.201)	-0.876 (1.199)	-0.605 (1.107)
Constant	31.059*** (3.146)	31.307*** (3.130)	31.136*** (3.053)
Observations	23,100	23,100	23,100
R-squared	0.191	0.193	0.204

**Notes:** OLS estimates. Robust standard errors clustered at the subject level in parentheses. Menu  $\{C1, C2\}$  is the reference category. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ . The dependent variable is the response time in seconds at a choice menu.

## B The 7 Lotteries and 15 Menus

Figure B.1: The 7 lotteries.




Figure B.2: The 15 distinct menus.



## C Instructions

Welcome, and thank you for your participation!

This is an experiment on decision making.

During its main phase you will be presented with 15 distinct **decision problems**.

A decision problem in this experiment is a menu that consists of 2 - 4 **lotteries** with different monetary prizes and/or different probabilities of winning the different prizes.

Each decision problem will be shown to you 5 times, in random order.

You will therefore see **75 decision problems in total**.

At every decision problem you will be asked to either choose one of the lotteries that are available at the relevant menu or to select "*I'm not choosing now*".

At the end of the experiment one of the 75 decision problems will be selected for you at random (each decision problem is equally likely to be selected).

You will then be reminded of the decision you made at that problem.

Your rewards will be determined as follows:

- If you had chosen a lottery at your randomly selected decision problem, this lottery will then be played out for you. **You will receive the lottery's prize accordingly, and an additional £5.**
- If you had selected "*I'm not choosing now*" at your randomly selected decision problem, you will be asked to choose a lottery at that problem, and this will be then played out for you. **You will receive the lottery's prize accordingly, and an additional £4.50.**

Please note that, regardless of how early you complete all tasks, your rewards won't be determined until at least 60 minutes have passed since the beginning of the experiment.

## D Experimental Interface: Screenshots

### D.1 Understanding Quiz

English (United Kingdom) ▾

Suppose that you chose lottery A from the set of lotteries A, B and C once, and that you chose lottery B from this set at a different point. Which one of the following is true?

- If one of these two menus is randomly selected for me at the end, I will get to choose again which lottery I want from these three.
- I will win lottery A if the former menu is randomly selected for me and lottery B if the latter menu is randomly selected for me. In both cases I will also receive 4.50 euro.
- I will win lottery A if the former menu is randomly selected for me and lottery B if the latter menu is randomly selected for me. In both cases I will also receive 5 euro.

English (United Kingdom) ▾

Suppose that during the main part of the experiment you had selected "I'm not choosing now" at what later turned out to be your randomly selected menu. Which one of the following is true?

- I will be asked to choose a lottery at that menu, which will then be played out for me. In addition to the lottery prize, I will also receive 4.50 euro.
- I will be asked to choose a lottery at that menu, which will then be played out for me. In addition to the lottery prize, I will also receive 5 euro.
- I will not be able to choose a lottery. My only reward will be 5 euro.

English (United Kingdom) ▾


Each one of the 75 menus that I will see in the main part of the experiment is equally likely to be selected for me at the end. The decisions I made in all other menus during the main phase will not affect my final rewards.

- True
- False




## D.2 Main Part


Q1-1. Choose one of the following two lotteries, or select "I'm not choosing now":



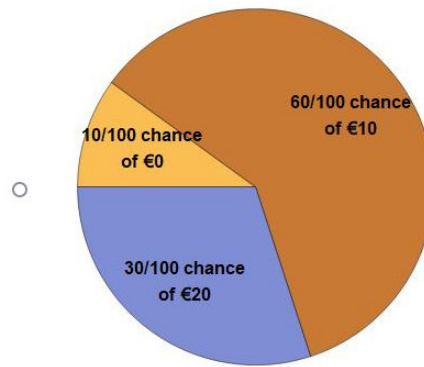
Q10-1. Choose one of the following three lotteries, or select "I'm not choosing now":




Q15-4. Choose one of the following four lotteries, or select "I'm not choosing now":




A2



A1



C1



C2

I'm not choosing now

### D.3 Randomly Selected Menu


English (United Kingdom) ▾

One of the 75 decision problems that you were presented with during the main phase of the experiment will now be selected at random for you.

You will first be reminded of the choice you made at that decision problem.


Then, the experiment will be concluded as described in the instructions, and your ultimately chosen lottery will be played out for you.

Your randomly selected menu is Q9-1 and shown below:



You chose I'm not choosing now at this menu.

Please choose one of the following two lotteries. As per the experiment's instructions, this lottery will be played out for you, and you will also receive 6.50 euro.



## E Stochastic Dominance (Non-)Relations

Table E.1: The cumulative density functions and expected values of lotteries involved in a FOSD claim.

Lottery	$x \in$	[0, 9)	[9, 10)	[10, 20)	[20, 24)	[24, $\infty$ )	$\mathbb{E}(x)$
<b>A1</b>		0.1	0.1	0.7	1	1	12
<b>A2</b>		0.2	0.2	0.7	1	1	11
<b>C1</b>		0.625	0.625	0.775	1	1	6
<b>C2</b>		0.625	0.825	0.825	0.825	1	6
<b>D</b>		0.15	0.15	0.65	1	1	12

1. A1 FOSD A2.
2. A1 FOSD C1.
3. A1 “nearly” FOSD C2: A1 dominates C2 in  $[0, 20)$  with a weighted-average difference in probability mass of 0.335; C2 dominates A1 in  $[20, 24]$  with a probability-mass difference of 0.175; for  $x \in [0, 24]$ , the net difference in weighted-average probability mass is 0.25 in favour of A1; finally  $\mathbb{E}_{A1}(x) = 12 = 2 \times \mathbb{E}_{C2}(x)$ .
4. A2 “nearly” FOSD C2: A2 dominates C2 in  $[0, 20)$  with a weighted-average difference in probability mass of 0.285; C2 dominates A2 in  $[20, 24]$  with a probability-mass difference of 0.175; for  $x \in [0, 24]$ , the net difference in weighted-average probability mass is  $\approx 0.21$  in favour of A2; finally  $\mathbb{E}_{A2}(x) = 11 > 6 = \mathbb{E}_{C2}(x)$ .
5. A1 and D are FOSD-unranked: compare their cdf values at  $[0, 10)$  and  $[10, 20)$ .

Table E.2: The area under the cumulative density function of each lottery involved in a SOSD claim, evaluated in the range  $[0, x]$  for each integer  $x \in \{1, 2, \dots, 24\}$ .

$x =$	Lottery	A1	B1	B2	D
1		0.1	0.25	0.25	0.15
2		0.2	0.50	0.50	0.30
3		0.3	0.75	0.75	0.45
4		0.4	1.00	1.00	0.60
5		0.5	1.25	1.25	0.75
6		0.6	1.50	1.50	0.90
7		0.7	1.75	1.75	1.05
8		0.8	2.00	2.00	1.20
9		0.9	2.25	2.25	1.35
10		1.0	2.50	2.25	1.50
11		1.7	3.05	2.90	2.15
12		2.4	3.60	3.55	2.80
13		3.1	4.15	4.20	3.45
14		3.8	4.70	4.85	4.10
15		4.5	5.25	5.50	4.75
16		5.2	5.80	6.15	5.40
17		5.9	6.35	6.80	6.05
18		6.6	6.90	7.45	6.70
19		7.3	7.45	8.10	7.35
20		8.0	8.0	8.75	8.0
21		9.0	9.0	9.40	9.0
22		10.0	10.0	10.05	10.0
23		11.0	11.0	10.70	11.0
24		12.0	12.0	11.35	12.0

1. A1 SOSD D: The claim can be established with a line-by-line inspection of the second and fifth columns of Table E.2.
2. D SOSD B1: -//- third and fifth columns.
3. A1 SOSD B1: -//- second and third columns.
4. A1 “nearly” SOSD B2: A1 dominates B2 in this sense for  $x \in [0, 22.143]$ .
5. D “nearly” SOSD B2: D dominates B2 in this sense for  $x \in [0, 22.143]$ .
6. B1 & B2 are SOSD-unranked: Compare values at  $x = 12$  vs  $x = 13$  and at  $x = 23$  vs  $x = 24$ .
7. C1 & C2 are SOSD-unranked: The areas under the cumulative density functions of C1, C2 are half those of B1, B2. Hence, the claim follows from the test of B1 v B2 above.