# Axiom Preferences and Choice Mistakes under Risk*

Fabian Herweg† Svenja Hippel‡ Daniel Müller§ Fabio Römeis¶

February 18, 2026

*We investigate whether violations of canonical axioms of choice under risk are mistakes or a manifestation of true preferences. First, we elicit axiom and gamble preferences and then allow subjects to revise their potentially conflicting preference statements. Among the behavioral patterns that allow for a clear-cut interpretation on the axiom level, we find that roughly 70% of axiom violations are intentional whereas only 30% are mistakes. On the subject level we can clearly categorize almost half of our subjects. Among those, roughly 24% behave like expected utility maximizers, 24% make occasional mistakes, and 52% refute the normative value of these axioms.*

*JEL classification: C91; D01; D81; D91.*

*Keywords: Axiomatic rationality; Choice Mistakes; Choice under risk; Intentional axiom violation; Subjective rationality*

*We would like to thank Jean-Michel Benkert, Eugen Dimant, John Duffy, Georg Nöldeke, Asri Özgümüs, Catherine Roux, Navid Sabet, Marco A. Schwarz, Stefan Trautmann, Marie Claire Villeval, as well as the participants at the MPI Lab Meeting (2022, Bonn), IMEBESS (2023, Lisbon), ESA World Meeting (2023, Lyon), SAET (2023, Paris), SABE (2023, San Diego), Nordic Conference on Behavioral and Experimental Economics (2023, Lund), U Freiburg for helpful comments and suggestions. Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation), grants HE 7106/4-1 and MU 4439/2-1 (GEPRIS project number 440656074) as well as under Germany's Excellence Strategy - EXC 2126/1-390838866. The experiment was preregistered with AsPredicted The anonymized pre-registration for peer-review can be found under https://aspredicted.org/DNP_VZJ.

†University of Bayreuth and CESifo, Faculty RW, Universitätsstr. 30, D-95440 Bayreuth, Germany, E-mail address: fabian.herweg@uni-bayreuth.de.

‡University of Bonn, Center for the Advanced Studies of Law and Economics, Adenauerallee 24-42, D-53113, Bonn, E-mail address: shippel@uni-bonn.de.

§University of Würzburg, Chair for Contract Theory and Information Economics, Sanderring 2, D-97070, Würzburg, E-mail address: daniel.mueller2@uni-wuerzburg.de.

¶University of Würzburg, Chair for Contract Theory and Information Economics, Sanderring 2, D-97070, Würzburg, E-mail address: fabio.roemeis@uni-wuerzburg.de.

## 1. Introduction

> *"[...] people are generally quite rational; that is to say, they usually have reasons for what they do."*
>
> —Herbert A. Simon (1985)

Expected utility theory (EUT) has long served as the canonical model of decision making under risk. Its normative appeal stems from its foundation on a small set of axioms that appear both intuitively compelling and behaviorally reasonable. EUT also enforces desirable prescriptions, such as respect for first-order stochastic dominance. For decades, these features have justified the use of EUT as the benchmark of rational choice.

Yet, a vast body of experimental evidence documents systematic and replicable violations of EUT. Subjects regularly deviate from its axioms in predictable ways, raising doubts about whether EUT reflects actual human rationality. Gilboa (2010) proposed a constructive response to this discrepancy: rather than assuming that axiom violations signal mistakes, researchers should directly ask whether decision makers themselves consider their behavior 'subjectively' rational. If, after being made aware of the axioms and the implications of their own choices, subjects revise their axiom-violating decisions, then they consider their own original behavior as irrational and education may indeed 'rescue' the canonical theory. If, however, subjects persist in their original choices, the implication is that they regard their own axiom violations as rational and maintain these violations intentionally – suggesting that EUT may not be an adequate normative standard.[1]

Since Gilboa's (2010) call, only a handful of experimental studies have attempted to empirically assess whether axiom violations are mistakes or intentional choices. Two broad approaches have emerged, both of which involve some sort of 'education' or 'preaching'. One approach, proposed by Nielsen and Rehbeck (2022), introduces subjects to the logic of canonical axioms in abstract terms, elicits their endorsement of those axioms, and then confronts them with their own inconsistent choices in subsequent lottery tasks. Subjects are then given the opportunity to revise their earlier axiom endorsement and their lottery choices. A second approach advocated by Humphrey and Kruse (2024, 2025) instead confronts participants directly with situation-specific normative statements that challenge their own previous behavior, thereby making motives for potential violations salient, before allowing lottery choices to be revised.

The results of these approaches diverge. Humphrey and Kruse (2024, 2025) find that participants often maintain their earlier axiom-violating behavior, interpreting their

[1] Notably, the definition of rationality advocated by Gilboa (2010) is subjective, as a mode of behavior that is considered rational by one person may not be deemed rational by another.

choices as reasonable given the decision context.[2] In contrast, Nielsen and Rehbeck (2022) report that many subjects endorse canonical axioms upfront and later revise their behavior to align with them, suggesting that violations are largely mistakes. Critics, however, argue that the experimental paradigm proposed by Nielsen and Rehbeck (2022) may overstate the normative force of axioms. For example, presenting canonical choice axioms alongside counter-intuitive or nonsensical 'control axioms' may have made endorsement of canonical axioms artificially attractive while failing to highlight reasons why an agent might want to reject them, e.g. anticipation of regret (Benjamin et al., 2021). Moreover, the design by Nielsen and Rehbeck (2022) only offers revision opportunities to participants who both endorsed an axiom and subsequently violated it. This truncation may bias conclusions by excluding the behavioral patterns of subjects who do not endorse axioms outright (Humphrey and Kruse, 2025).

This paper builds directly on this debate by modifying the Nielsen and Rehbeck (2022) paradigm and assessing the normativity of canonical choice axioms under risk (first-order stochastic dominance, transitivity, consistency and brach independence). Our results indicate that concerns about bias in the framework by Nielsen and Rehbeck (2022) are well-founded. In our experiment, initial selection rates are quite dispersed over canonical choice axioms, ranging from about 40% to above 80%. At the axiom level, while roughly 60% of participants behave in ways consistent with EUT, a sizeable 25% exhibit axiom-violating lottery choices that allow for a clear-cut interpretation – and of these nearly 70% are intentional rather than mistakes. At the subject level, about half of participants can be unambiguously categorized: 24% never violate axioms, 24% commit violations that can be regarded as mistakes, but a majority of 52% intentionally violate at least one axiom. Overall, significantly more subjects violate axioms intentionally rather than by mistake. This also holds true for each axiom separately, except first-order stochastic dominance.

To address the aforementioned concerns that were voiced with regard to the experimental paradigm by Nielsen and Rehbeck (2022), we introduce three key changes to their design. First, to avoid artificially inflating endorsement rates, we present only canonical – that is, intuitively plausible – axioms, without including implausible control axioms. Second, we present the lottery choices in a format that clearly reveals the correlations between them. The lotteries we use are ones for which existing evidence suggests that subjects may respond to the juxtaposition of outcomes in ways that lead to violations of the relevant axiom. Importantly, for these lotteries, a decision maker might rationally prefer to violate an axiom – for instance, to avoid ex post regret. Third, we alter the revision stage by implementing both batch revisions (subjects revise all choices related to

[2] Adhering to contemporary experimental standards, Humphrey and Kruse (2024, 2025) thereby validate the findings in Slovic and Tversky (1974), who – with a sample size of only 29 subjects and non-incentivized, hypothetical choices – already analyzed the normativity of the sure-thing principle.

an axiom simultaneously) and placebo revisions (all choices are revisited, not only those conflicting with stated preferences). This ensures that participants' own past choices, rather than the experimenter's framing, serve as arguments for or against an axiom. We also adopt more neutral language, e.g., by referring to axioms as 'choice assistants' rather than 'decision rules.'

Furthermore, regarding the empirical approach, the elicitation of axiom preferences in our experiment is not used to exclude observations but to be understood as an opportunity to familiarize subjects with the axioms on neutral grounds upfront. That is, we classify behavioral patterns that emerge from lottery choices and revision behavior irrespective of initially stated axiom preferences. If a participant violates an axiom but later changes her lottery choices to align with it – and ultimately states a preference in favor of the axiom – we classify the violation as a mistake. If, however, the participant continues to violate the axiom and ultimately does not endorse it, we classify the violation as intentional. This design avoids truncating the dataset and allows us to observe behavior across all relevant groups of subjects.

Our findings within this modified framework cast doubt on the normativity of canonical choice axioms as normative standards. While many individuals behave in ways compatible with EUT, a substantial share deliberately and persistently reject at least one axiom. In line with the above quote by Herbert A. Simon (1985), our evidence suggests that even when subjects violate axioms, they do so for reasons they themselves regard as rational.

The consistency and persistence of lottery choices that violate canonical axioms have been investigated well before Gilboa's (2010) call to distinguish between standard economic rationality and subjective rationality. The approaches used by Nielsen and Rehbeck (2022) and Humphrey and Kruse (2024, 2025) were inspired by the pioneering work of MacCrimmon (1968), who engaged participants in free-form discussions of several decision-theoretic postulates. A complementary experimental approach analyzes deliberation over earlier choices without explicitly educating subjects about, or emphasizing, the axioms themselves. In the spirit of this alternative line of research, several studies allow respondents to revisit earlier lottery choices, either by presenting the same problems across multiple rounds (van de Kuilen and Wacker, 2006; van de Kuilen, 2009; Nicholls et al., 2015; Birnbaum and Schmidt, 2015) or by offering the opportunity to revise initial decisions (Breig and Feldman, 2024; Yu et al., 2021; Crosetto and Gaudeul, 2023). Overall, these studies find that reconsidering earlier choices tends to increase consistency with expected utility. Other work explores the presentation of identical choice problems in different frames (McNeil et al., 1988; Miller and Fagley, 1991; Sieck and Yates, 1997; Druckman, 2001). More recently, Benjamin et al. (2021) examined the normative appeal of the independence axiom's implications for choices between simple and compound

lotteries. When confronted with both frames simultaneously – after having previously made separate choices – respondents, on average, revised their decisions to align with the axiom's implications.

The rest of the paper is structured as follows: Section 2 introduces the theoretical framework, followed by the experimental design in Section 3. After outlining our empirical strategy to differentiate mistakes from intentional axiom violations in Section 4, we present the results of our experiment in Section 5. Section 6 concludes.

## 2. Theoretical Framework

In our experiment, we consider four canonical axioms that have a strong normative appeal and violations of which can be observed in simple pairwise choice problems that do not involve mixtures of lotteries (i.e., compound lotteries):

**First-order Stochastic Dominance (FOSD):** If under one lottery the probability to obtain at least a particular monetary amount is weakly higher for all amounts and strictly higher for at least one amount than under the alternative lottery, then the decision maker must prefer the former lottery over the latter.

**Transitivity (TRANS):** If a first lottery is preferred over a second and the second lottery is preferred over a third, then the first lottery must also be preferred over the third.

**Consistency (CONS):** If a first and a second lottery both represent one and the same probability distribution over the set of feasible payoffs (i.e., if the two lotteries are stochastically equivalent) and the first lottery is preferred over a third lottery, then the second lottery must also be preferred over the third lottery.

**Branch Independence (BRANCH):** If a first and a second lottery yield a given payoff $x$ with equal probability and the first lottery is preferred over the second lottery, then the preference must not change if the payoff $x$ is interchanged with some other payoff $y$.

Notably, motives like the avoidance of ex-post regret may make a violation of these axioms attractive.[3] To see this, consider the choice between the two lotteries $L_1$ and $L_2$ depicted in Table 1, which is based on Loomes et al. (1992).[4] With the probability of obtaining at least a given payoff being at least as high under $L_1$ as under $L_2$ for all payoffs, $L_1$ dominates $L_2$ in the sense of FOSD. The choice of $L_1$, however, exposes the decision

[3]The view that the preference for avoiding post-decisional regret is rational was already advocated by Loomes and Sugden (1982). While supported by the neurodata reported in Bourgeois-Gironde (2010), this view is not uncontroversial. For opposing arguments, see the discussion in Bleichrodt and Wakker (2015).

[4]The choice problem depicted in Table 1 is an example from our actual study. The difference to Loomes et al. (1992) is that we triple all payoffs and choose the currency to be €. For a complete overview of the choice problems in our study see Appendix B.3.

maker to the risk of experiencing severe ex-post regret if state 4 turns out to be the true state of the world – in which case the decision maker obtains €0 but would have gotten €30 if she had chosen the dominated lottery $L_2$. Thus, a regret-averse decision maker may well end up choosing the dominated lottery $L_2$.[5]

| | $s = 1$ | $s = 2$ | $s = 3$ | $s = 4$ |
|---|---|---|---|---|
| $L_1$ | €0 | €30 | €9 | €0 |
| $L_2$ | €9 | €9 | €0 | €30 |
| Prob. | 18% | 20% | 42% | 20% |

Table 1: Lotteries $L_1$ and $L_2$ depicted in the action-state matrix.

Precise formal definitions of all four relevant axioms as well as specific examples for how violations of TRANS, CONS and BRANCH may be rooted in the avoidance of ex-post regret are provided in Appendix A.

## 3. Experimental Design

To identify whether axiom violations are intentional or not, we build upon the three-stage experimental design proposed by Nielsen and Rehbeck (2022): first, elicitation of preferences over canonical axioms; second, lottery choices; third, providing the opportunity to revise axiom preferences and lottery choices.[6] Next, we describe the specifics of the experiment's three stages in more detail, outline the general procedures (including the incentivization of the decisions in each of the three stages) and discuss our design choices. The complete instructions are available in Appendix B.1.

### 3.1. Stage 1: Axiom Preferences

To elicit whether a subject considered a given axiom as a desirable guiding principle, the subject was asked whether she wants the axiom – which was labeled as 'choice assistant' – to make choices on her behalf or to make the respective choices herself. Subjects had to make this decision for each of the four axioms that were introduced in Section 2: FOSD, TRANS, CONS and BRANCH. If a subject opted for an axiom to make choices on her behalf and if this decision later became payoff relevant, then the axiom would be applied to a choice situation (the specifics of which were not known to the subject at the point of this decision) and the subject would receive the outcome of that choice made by the axiom.

[5] In fact, Loomes et al. (1992) find that 42% of their 90 subjects choose the dominated lottery $L_2$ in the choice problem in Table 1 (with one-third of the payoffs in £).

[6] To avoid any order effects, decisions in each stage were randomized on the subject level.

With axioms being a highly abstract concept, we used a pictorial logic statement to explain each axiom as tangibly as possible, with colored circles representing otherwise unspecified lotteries. Whenever specific additional assumptions on the properties of the involved lotteries were imposed, these were clearly spelled out to our subjects. Furthermore, we complemented each such pictorial logic statement with a verbal explanation of the workings of the respective axiom as well as an easy-to-understand example that related directly to the pictorial presentation.

For example, consider the TRANS axiom as depicted in Figure 1. Here, it was explained to subjects that if they prefer a lottery represented by a blue circle over another lottery represented by an orange circle and, at the same time, prefer the lottery represented by an orange circle over a third lottery represented by a green circle, then, if they should select the axiom to make choices on her behalf, the axiom would choose the lottery represented by the blue circle instead of the lottery represented by the green circle. The working of the remaining three axioms was explained in an analogous fashion. The full description of all axioms can be found in Appendix B.2.

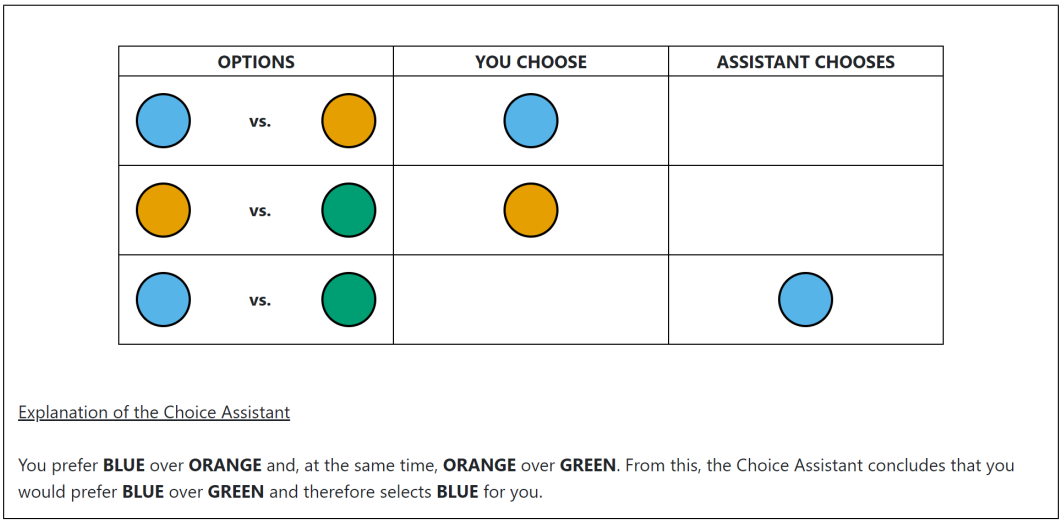


Figure 1: Depiction of transitivity (TRANS).

### 3.2. Stage 2: Lottery Choices

Each subject made 20 subsequent pairwise lottery choices. Specifically, subjects were informed that they would have to make all these choices on their own even if they had selected one or more of the axioms in Stage 1. In our experiment, a lottery assigned a monetary amount between €0 and €60 to each integer number from 1 to 100. At the end of the experiment, if a given lottery was selected as payoff relevant, the computer would randomly generate an integer number between 1 and 100 and the lottery's actual

payoff corresponded to the monetary amount that the lottery assigned to this randomly generated number.

The 20 pairwise lottery choices that subjects had to make in Stage 2 are depicted in Appendix B.3, from which it can be seen that Problems 1 to 16 allowed for each axiom to be violated in exactly two instances and that each of these sixteen pairwise choice problems was associated with exactly one axiom. To construct these sixteen choice problems, we relied on guidance obtained from previous experimental studies as far as possible.[7] Problems 17 to 20 did not allow for any axiom to be violated. These choices were used as preference input for axiom application in case that a subject's axiom choice from Stage 1 or her revised axiom choice from Stage 3 turned out to be payoff relevant.[8]

Inspired by earlier studies on juxtaposition-driven regret effects, we presented the 20 pairwise lottery choice problems in a matrix display, under which the correlation of outcomes is fully specified and transparently displayed.[9] As depicted in Figure 2, the two feasible lotteries correspond to the two rows labeled with A and B, respectively, and the numbers within these rows' cells denote amounts of money in Euro. For each of a lottery's potential monetary outcome, the numbers along the top of the matrix indicate the range of numbers for which the respective monetary amount would be paid out. The numbers along the bottom of the matrix show the probability with which the randomly generated number falls into the corresponding range indicated at the top of the matrix. The width of each column is proportional to the respective probability.

| | 1 18 | 19 38 | 39 80 | 81 100 |
|---|---|---|---|---|
| **A** | **€9** | **€9** | **€0** | **€30** |
| **B** | **€0** | **€30** | **€9** | **€0** |
| | 18% | 20% | 42% | 20% |

Figure 2: Matrix display.

Battalio et al. (1990) and Harless (1992), however, found the occurrence of systematic juxtaposition effects in line with the avoidance of ex-post regret to be highly susceptible

[7]Problems 1 and 2, which allow for a violation of FOSD, are based on Questions 3 and 5 in Loomes et al. (1992). The pair of Problems 3 and 4 as well as the pair of Problems 5 and 6 allow for a violation of CONS and are inspired by Questions 8 and 8D and Questions 2 and 2D in Starmer and Sugden (1993), respectively. The triple of Problems 7, 8 and 9 as well as the triple of Problems 10, 11 and 12, which both allow for a violation of TRANS, are based on Triples 8 and 3 in Loomes et al. (1991), respectively. The pair of Problems 13 and 14 as well as the pair of Problems 15 and 16 allow for a violation of BRANCH and are partly based on Starmer (1992).

[8]While Problems 17 and 18 were used as input for TRANS, Problem 19 and Problem 20 were fed into CONS and BRANCH, respectively.

[9]This display format was used, for example, by Loomes (1988a,b, 1989), Loomes and Sugden (1987), Loomes et al. (1989, 1991, 1992), Starmer (1992), Starmer and Sugden (1989, 1993, 1998), and Castillo (2020).

to the format of problem representation. Specifically, they provide evidence that the prevalence of such effects is greatly diminished if the choice problem is not represented in the matrix display but in a verbal display.[10] Therefore, to assess the robustness of our results, we presented the 20 pairwise lottery choice problems to a different group of subjects in the verbal display format depicted in Figure 3. For a given lottery, here denoted as either Option A or Option B, the left entries specify the monetary amounts that are potentially paid out, the entries in the middle denote the ranges of numbers for which the respective monetary amount is paid out, and the right entries list the probabilities of the corresponding events.

**OPTION A**

| | | | |
|---|---|---|---|
| **€9** | if | 1-18 | (18%) |
| **€9** | if | 19-38 | (20%) |
| **€0** | if | 39-80 | (42%) |
| **€30** | if | 81-100 | (20%) |

**OPTION B**

| | | | |
|---|---|---|---|
| **€0** | if | 1-18 | (18%) |
| **€30** | if | 19-38 | (20%) |
| **€9** | if | 39-80 | (42%) |
| **€0** | if | 81-100 | (20%) |

Figure 3: Verbal display.

Importantly, for each pairwise choice problem, both the matrix display and the verbal display contain the exact same information as they fully specify the correlation of the monetary outcomes of the two feasible lotteries.

### 3.3. Stage 3: Revision Behavior

After stating their axiom preferences in Stage 1 and making their lottery choices in Stage 2, subjects had the opportunity to revise any of their previously made decisions. These revision opportunities were grouped according to the four axioms. First, for a given axiom, a subject was shown whether she had endorsed this axiom to make choices on her behalf. At the same time, the subject was shown all her lottery choices from Stage 2 that had the potential to violate the axiom in question.[11] Thereafter, the subject could freely decide whether to change her axiom preference, to change any number of the lottery choices that related to the axiom in question, to change both axiom preference and lottery choices, or

[10] That is, to put it in the words of Starmer and Sugden (1998), "it seems that [regret effects come] into play only when decision problems are framed in ways that make within-event, cross-act comparisons particularly salient."

[11] In reference to Appendix B.3, a subject was shown her decisions in Problems 1 and 2 with regard to FOSD, in the pair of Problems 3 and 4 as well as in the pair of Problems 5 and 6 with regard to CONS, in the triple of Problems 7, 8 and 9 as well as in the triple of Problems 10, 11 and 12 with regard to TRANS, and in the pair of Problems 13 and 14 as well as in the pair of Problems 15 and 16 with regard to BRANCH.

leave each of these decisions unchanged. Importantly, every axiom choice from Stage 1 and each of the 16 core lottery choices from Stage 2 could be revised exactly once.

Each axiom from Stage 1 could be violated in exactly two instances in Stage 2. In consequence, in Stage 3, subjects had to make their revision decisions with regard to a given axiom after having been informed that either they had adhered to the axiom on all occasions or that they had violated the axiom on all occasions or that they had adhered to the axiom on one occasion and violated the axiom on the other occasion. The associated revision process comprised a series of information screens followed by a decision screen. On the first information screen, an example of which is shown in Figure 4, a subject was informed about her axiom decision in Stage 1, where the pictorial logic statement of the axiom in question was presented once again, and she was shown her decisions for the first constellation of lotteries from Stage 2 that potentially allowed for a violation of that axiom. Notably, lottery decisions were presented in the exact same display format as in Stage 2. On the next information screen the subject was shown the second constellation of lotteries from Stage 2 that potentially allowed for a violation of the axiom. The last information screen contained the subject's axiom decision as well as her decisions in both constellations of lotteries. While the subject's previous decisions from Stages 1 and 2 were locked-in throughout all these information screens, they could be freely changed in the final decision screen, which the subject could only access by actively clicking on the "Make Changes" button positioned in the lower part of the last information screen.

As can be seen from Figure 4, two further features were implemented to make the revision process as transparent as possible. First, to highlight how a subject's lottery choices in Stage 2 relate to her axiom preferences in Stage 1, the pairwise choice problems in each constellation of lotteries that allowed for a violation of a given axiom were matched to the colored circles of the pictorial logic statement used to describe that axiom. Furthermore, for each such constellation of lotteries, a short written explanation in neutral language reminded the subject of what option she had chosen and which option the axiom would have prescribed to be chosen.

### 3.4. Procedures

At the beginning of the experiment, the subjects were informed that the main part of the experiment consisted of three stages and that they had to make a series of decisions in each of these stages. The instructions proceeded with an illustration of what a lottery is. The concept of a lottery was explained as a dice roll with each number between 1 and 100 being equally likely to come up and with a monetary amount between €0 and €60 being assigned to each number.

Furthermore, subjects were informed about the incentivization procedure which strictly

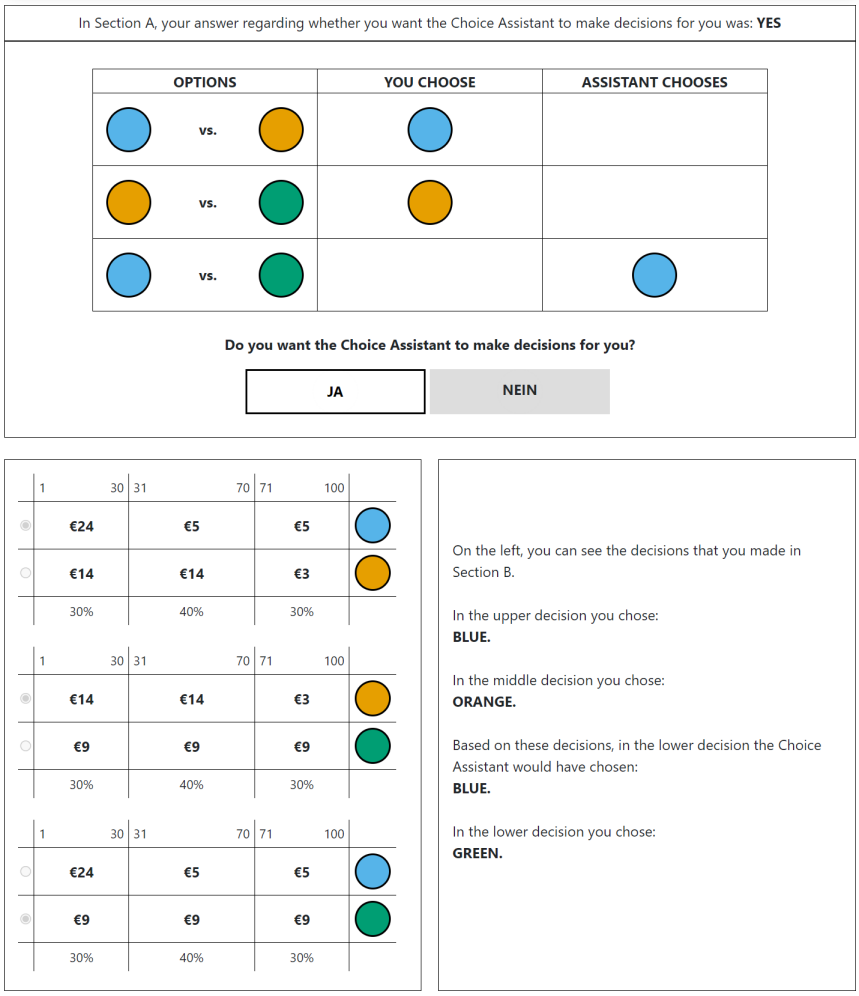


Figure 4: Information page about decisions made in Stage 1 and 2.

followed Nielsen and Rehbeck (2022) and is described in detail in Appendix B.4. Exactly one decision from any of the stages described above became payoff-relevant, i.e. a subject could be paid for one of her original axiom choices from Stage 1, one of her original lottery choices from Stage 2, or one of her (potentially revised) decisions from Stage 3. Which decision became payoff-relevant was determined by a random draw of the computer at the end of the experiment. The probability of being payoff-relevant was identical for each and every decision made.

The experiment was conducted at the *BonnEconLab* of the University of Bonn, Germany, in November 2022. It was computerized using the software *o-Tree* (Chen et al., 2016). Subjects sat in visually isolated cubicles. Before each stage in the main experiment, subjects had to answer check questions to ensure their understanding. From a database of more than 5,000 people, we recruited 349 subjects using *hroot* (Bock et al., 2014). Each subject faced only one display format, with 122 subjects making lottery choices under the matrix display and 115 under the verbal display. 112 additional subjects were recruited for a robustness check where the 20 lottery choice problems from our

experiment were presented in the exact same display format as in Nielsen and Rehbeck (2022).[12] Subjects were mainly undergraduate students from a variety of disciplines. On average a session took 90 minutes and subjects were paid €25.95, including a show-up fee of €5. After the main experiment, we elicited subject's risk attitudes, regret aversion and cognitive reflection ability.[13] Additionally, we had a questionnaire asking for each subject's sociodemographics (gender, age, field of study, aspired university degree, and final math grade in high school) as well as the strategy pursued when making their decisions in each of the three stages of the main experiment.

### 3.5. Discussion of Design Choices

As noted by Benjamin et al. (2021) and Humphrey and Kruse (2025), Nielsen and Rehbeck (2022) report surprisingly high endorsement rates (between 82% and 90%) for canonical choice axioms in the first stage. This finding may partly reflect that, in addition to the canonical axioms, Nielsen and Rehbeck (2022) also presented 'control' axioms (which prescribe the exact opposite) and meaningless 'distractor' axioms. The inclusion of these may have led subjects to perceive the canonical axioms as more intuitively plausible.

In the revision stage, where Nielsen and Rehbeck (2022) found that nearly four times as many participants adjusted their lottery choices to match an axiom rather than retract their endorsements, several factors seem relevant. Axioms were presented as abstract propositions in the form of a 'decision rule' the simplicity of which may have implicitly framed them as 'correct' without highlighting reasons why one might want to deviate (e.g., anticipated regret). This may have resulted in an "experimenter demand effect pushing participants to endorse the axioms while eliminating an effect in the opposite direction" (Benjamin et al., 2021, p.36). Likewise, the display format used by Nielsen and Rehbeck (2022) did not embody any rationale for why an axiom violation might be deemed reasonable (e.g., regret avoidance) and participants in some cases might even have felt nudged into violations. Finally, participants corrected multiple violations of the same axiom independently rather than collectively, which may have reinforced the impression that the axiom was compelling instead of highlighting the repeated violations. Overall, the conditions made it comparatively easy for subjects to stand by their initial endorsement of the axiom.

In our experiment, we retained pictorial-logic statements similar to those used by Nielsen and Rehbeck (2022), but presented only canonical axioms.[14] Furthermore, we

[12] The results of this additional variation in the display format are discussed in Section 5.3.

[13] See Appendix B.5 for more details on the post-experimental tests.

[14] Nielsen and Rehbeck (2022) used a very simple and therefore intuitively plausible pictorial-logic statement to convey each axiom's logic. For example, FOSD was illustrated with compound lotteries. We used more precise pictorial-logic statements, partly accompanied by short verbal qualifications regarding specific characteristics of the considered lotteries, as in the case of FOSD. See Appendix C

used a state-space-based display format for pairwise lottery choices that highlights correlations between lotteries, thereby allowing for discerning why it might actually make sense to violate an axiom (e.g., to avoid ex-post regret). In the revision stage, we introduced placebo and batch revisions. That is, respondents revised all their decisions (rather than only those in which axiom preferences and lottery choices are inconsistent) and all lottery choices concerning the same axiom were revised simultaneously (rather than independently of each other in subsequent revisions). This way, individuals' own behavior across choices could serve as an argument for or against an axiom. Finally, instead of labeling the axioms as 'decision rules,' we adopted the more neutral term 'choice assistant.' In the same spirit, whereas Nielsen and Rehbeck (2022) explicitly informed their subjects that their "choices were inconsistent with [the] rule" (Nielsen and Rehbeck (2022), p. 45, Online Appendix), we merely reminded subjects of their initial lottery choices and indicated which option the corresponding choice assistant would have selected.

## 4. Empirical Strategy

Next to the aforementioned design aspects, also the empirical approach used by Nielsen and Rehbeck (2022) has been criticized as biased in favor of canonical axioms being portrayed as normative because they collected data on revision behavior in the third stage only for participants who endorsed a given axiom in the first stage *and* violated this axiom in the second stage. As argued by Humphrey and Kruse (2025, p.19), "truncating the dataset to the preclusion of observations relevant to the normative status of choice axioms seems unwarranted." For example, consider a subject who did not endorse an axiom initially in Stage 1, violated that axiom through her lottery choices in Stage 2, and then refrained from revising any of her decisions in Stage 3. Such a case would represent the literal embodiment of an intentional refusal to accept the axiom – and a deliberate violation of it – reflecting behavior the subject perceives as subjectively rational. Similarly, a subject who initially endorsed an axiom in Stage 1 and did not violate it in Stage 2 but later deselected the axiom and made choices inconsistent with it during the revision phase in Stage 3 could likewise be viewed as intentionally violating the axiom – that is, as behaving in a subjectively rational manner. Thus, one cannot dismiss the possibility that truncating the dataset may affect conclusions.

In light of these concerns, we follow an empirical strategy aimed at capturing the full extent not only of axiom violations by mistake, but also of intentional axiom violations. To this end, we divide the observed behavioral patterns from Stages 2 and 3 into different categories, which we then use to classify our subjects. This categorization of intentional or

for a detailed comparison.

| *Category* | *S2: Violation* | *S3: Axiom Preference* | *S3: Violation* | *Label* |
|---|---|---|---|---|
| **[1]** | No | Yes/No | No | **EUT** |
| **[2]** | Yes | Yes | No | **Mistake** |
| **[3]** | Yes | No | Yes | **Intention** |
| **[4]** | Yes | No | No | |
| **[5]** | No | No | Yes | **Other** |
| **[6]** | Yes/No | Yes | Yes | |

Table 2: Categorization of behavioral patterns on the axiom level.

non-intentional axiom violations is not based on whether subjects did or did not endorse an axiom in Stage 1; in particular, the elicitation of axiom preferences in Stage 1 is not used to exclude observations. This is not to say that we consider the first stage of the design in Nielsen and Rehbeck (2022) as irrelevant per se. To the contrary, we believe that having subjects assess an axiom's appeal as a desirable guiding principle before contrasting it with their own specific lottery choices should allow for a more neutral evaluation of the axiom.

**Categorization on the Axiom Level.** To categorize a subject's behavior with regard to a particular axiom, we ask three questions: First, did at least one of the subject's lottery choices in Stage 2 violate the axiom? Second, did the subject ultimately endorse the axiom in Stage 3? Third, did at least one of the subject's (potentially revised) lottery choices in Stage 3 violate the axiom? Based on the answers to these questions, a subject's behavior in Stages 2 and 3 with regard to a particular axiom can take one of eight possible behavioral patterns, which we group into the six categories displayed in Table 2.

In category [1], a subject's lottery choices never violate the axiom in question, neither in Stage 2 nor in Stage 3. With this behavioral pattern being compatible with expected utility theory (von Neumann and Morgenstern, 1947), we take a conservative stand here and label category [1] as *EUT*. Moreover, we stack the deck in favor of this *EUT* category as we do not require the axiom in question to be explicitly endorsed in Stage 3.

In category [2], a subject violates the axiom in question in Stage 2 but then, in Stage 3, endorses this axiom as a desirable guiding principle and revises her lottery choices to be consistent with the axiom. This behavioral pattern, where a subject corrects her initially conflicting lottery choices to be consistent with an axiom that she ultimately deems worthwhile following, reflects an unintentional violation of the axiom and is labeled as *Mistake*.

In category [3], a subject violates the axiom in question in Stage 2 and then, in Stage 3, renounces this axiom and keeps violating the axiom with her lottery choices. We classify

| *Axiom-level behavioral patterns* | *Label* |
|---|---|
| Category [1] only | **S-EUT** |
| Category [2] (at least once) or category [1] | **S-Mistake** |
| Category [3] (at least once) or category [1] | **S-Intention** |
| Any other combination | **S-Other** |

Table 3: Categorization of behavioral patterns on the subject level.

this behavioral pattern, where a subject ultimately rejects and knowingly violates the axiom in question, as an intentional violation and label it as *Intention*.[15]

The remaining behavioral patterns do not allow for such a straightforward interpretation. In category [4], a subject revises her initially conflicting lottery choices to align with the axiom in the end, which could be interpreted as the initial axiom violation having been a mistake. In contrast to category [2], however, in category [4] there is no ultimate endorsement of the axiom as desirable guiding principle, which makes category [4] less indicative of a mistake than category [2]. In category [5] a subject's behavior in Stage 3 coincides with the behavior in category [3]: ultimately the axiom in question is violated and not endorsed. This behavior in Stage 3 suggests that the final violation of the axiom is intentional. This intentional violation is less clear though than for category [3], where, in contrast to category [5], a subject violated the axiom already at Stage 2. Finally, in category [6], a subject's behavior is inherently inconsistent because ultimately, in Stage 3, the subject states a strict preference for adhering to the axiom in question but locks in (possibly revised) lottery choices that violate the axiom.

To provide an as clear-cut interpretation as possible of whether lottery choices that violate a given axiom are intentional or not, in the following, we pool categories [4], [5] and [6] under the label *Other*.

**Categorization on the Subject Level.** To categorize behavior on the subject level, we individually 'aggregate' each subject's behavior on the axiom level. The resulting subject-level categories are summarized in Table 3, the left column of which shows the axiom-level categories into which a subject's behavioral patterns with regard to the four axioms fall.

Following our conservative approach from the axiom-level categorization, we categorize subjects who never violate any axiom – i.e., whose behavior on the axiom level falls into category [1] for all axioms – under the label *S-EUT*, where S refers to the subject level.

[15] One might be worried that behavioral patterns in category [3] will not allow for a clear-cut interpretation as they might comprise behavior under which a given axiom is violated in Stage 2 in only one of the two possible lottery constellations and then, in Stage 3, this violation is revised to align with the axiom whereas the originally non-conflicting lottery choice is revised to violate the axiom. Foreshadowing our results, however, such behavior is virtually non-existent with at most one (out of 237) patterns for each axiom exhibiting this behavior.

To reduce the scope for any ambiguity in the interpretation of our results as far as possible, we extend the axiom-level categories labeled *Mistake* and *Intention* to the subject level in the strictest possible sense. Requiring that a subject's axiom-level behavior for each axiom falls exclusively into either category [2] or category [3], however, seems too hard a requirement, because even behavioral theories (e.g., regret theory) do not predict axioms to be violated in each and every circumstance (but rather allow for axioms to be violated occasionally). Therefore, to establish the subject-level category *S-Mistake*, we require a subject to violate *at least one* of the four axioms in Stage 2 and then, in Stage 3, to address *any* axiom violation in the spirit of axiom-level category [2] while any non-violated axiom remains non-violated. Likewise, to fall into the subject-level category *S-Intention*, a subject has to violate at least one axiom in Stage 2 and then, in Stage 3, has to address any axiom violation according to axiom-level category [3] while leaving any non-violated axiom non-violated.[16]

Any other combination of axiom-level behavioral patterns bears a less clear-cut interpretation and therefore is subsumed under the category *S-Other*. In particular, this includes profiles where the axiom-level behavioral pattern for at least one axiom falls into the categories [4], [5] or [6] as well as profiles under which one axiom violation is addressed according to category [2] and another axiom violation is addressed according to category [3].

## 5. Results

In this section, we categorize the experimentally observed behavioral patterns both at the axiom level and at the subject level, following the empirical strategy outlined in Section 4. As we do not find systematic differences between the matrix and verbal display formats in decision behavior at either Stage 2 or Stage 3 (see Appendix D for details), we pool the observations from both display formats for conciseness. This results in 237 behavioral patterns for each of the four canonical choice axioms.[17]

### 5.1. Axiom Endorsement and Axiom Violations

Figure 5a depicts the initial (i.e., Stage 1) endorsement rates for the four canonical axioms in our experiment. The four axioms are selected at an average rate of 61.29%, with relatively high endorsement rates of 82.70% for FOSD and 68.35% for TRANS and lower

[16] Thus, while behavioral patterns in the subject-level categories *S-Mistake* and *S-Intention* may contain up to three axiom-level behavioral patterns from category [1], they otherwise comprise only axiom-level behavioral patterns from category [2] and category [3], respectively.

[17] Unless stated otherwise, in the following all p-values are from probit regressions with robust standard errors clustered at the subject level. Controls include risk attitudes, regret aversion and cognitive reflection ability, as well as sociodemographics (gender, age, field of study, aspired university degree and final math grade in high school). Qualitative results are the same for Fisher exact tests.

endorsement rates of 54.43% for CONS and 39.66% for BRANCH. More than half of the subjects (52.32%) found at least three of the four axioms appealing, whereas only a minority of 18.57% chose only one or none axiom.

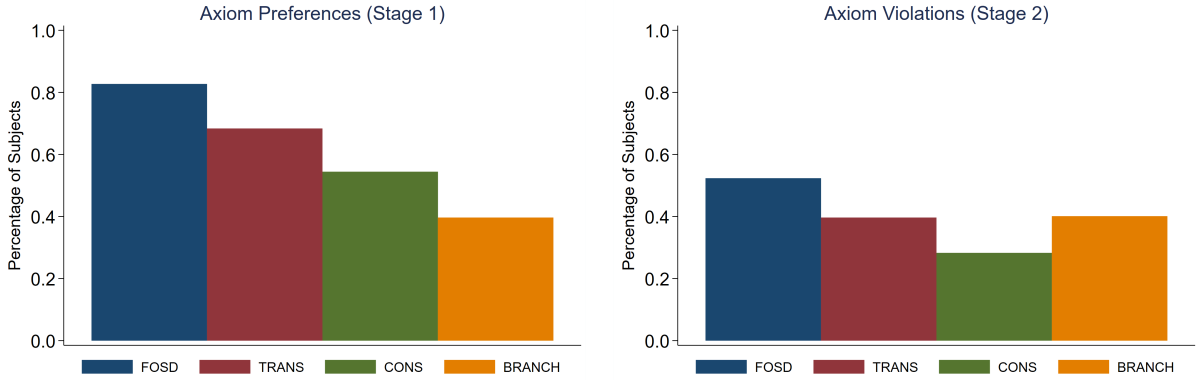


(a) Stage 1 axiom endorsement rates.

(b) Stage 2 axiom violation rates.

Figure 5: Initial axiom endorsement and violation shares for matrix and verbal display.

**Result 1.** *Initial endorsement rates are quite dispersed over canonical choice axioms, ranging from about 40% to above 80%.*

Initial (i.e., Stage 2) violation rates split up by axioms are depicted in Figure 5b. 52.32% of our subjects violated FOSD, 28.27% violated TRANS, 39.66% violated CONS and 40.08% violated BRANCH at least in one of the two constellations that targeted the axiom in question. Conditioning the probability to violate a given axiom at least once on the axiom's prior endorsement, in aggregate, we find that those subjects who endorsed an axiom violated it in 40.79% of all cases whereas subjects who did not endorse the axiom violated it in 38.86% of all cases. That is, prior endorsement of an axiom does not make it any less likely to violate the axiom in question.[18]

**Result 2.** *Initial violation rates vary across canonical choice axioms, ranging from about 30% to above 50%. Aggregate initial violation rates are nearly identical for those subjects who endorsed a given axiom in Stage 1 and those who did not endorse the axiom.*

### 5.2. Disentangling Mistakes from Intention

Based on the pooled data from the two display formats, our analysis of whether axiom violations are by mistake or intentional comprises 948 behavioral patterns that result from 237 subjects making choices with regard to four axioms. As outlined in Section 4,

[18]For FOSD, the violation rates are 52.05% vs. 53.66% (Chi-squared test, p=0.85); for TRANS, 27.16% vs. 30.67% (Chi-squared test, p=0.58); for CONS, 38.76% vs. 40.74% (Chi-squared test, p=0.76); and for BRANCH, 43.62% vs. 37.76% (Chi-squared test, p=0.37). That is, the differences in violation rates between subjects who endorsed the axiom and those who did not are not statistically significant.

a behavioral pattern here refers to a subject's revised preference statement regarding a given axiom together with her original and revised lottery choices in the two constellations that allow for a violation of the axiom in question.

On the axiom level, we categorize the 948 behavioral patterns according to Table 2. Viewed across all four axioms, a majority (59.18%) of behavioral patterns involves no violation of the axiom in question and therefore falls into category [1], and only a minority (15.51%) of behavioral patterns does not allow for an unambiguous interpretation and thus falls into categories [4] to [6]. The remaining quarter (25.32%) of behavioral patterns falls into the categories [2] and [3], such that a sizeable share of behavioral patterns represent an axiom violation that is either intentional or by mistake. Of these behavioral patterns, a clear majority of 69.58% can be categorized as an intentional violation of the axiom in question (category [2]), whereas only 30.42% represent mistakes (category [3]). Thus, viewed across all axioms, more than twice as many axiom violations are by intention rather than by mistake.

Figure 6 shows the axiom-level categorization of behavioral patterns. Overall, the picture looks quite similar for each axiom, with the majority of behavioral patterns (ranging from 47.68% to 70.46%) falling into category [1] and only a minority (ranging from 12.24% to 20.68%) falling into categories [4] to [6]. The share of behavioral patterns representing an intentional violation of the axiom in question (category [2]) ranges between 12.66% and 24.47%, whereas the share of behavioral patterns that embody a mistake (category [3]) is only between 2.53% and 17.30%. Notably, except for FOSD, where categories [2] and [3] are of about equal size (Binomial test, p=0.82), for TRANS, CONS and BRANCH there are significantly more behavioral patterns categorized as having violated the respective axiom intentionally rather than by mistake (Binomial test, $p < 0.01$ for TRANS, CONS and BRANCH).

On the subject level we aggregate the four behavioral patterns that a subject exhibited for the four different axioms into a subject-level categorization according to Table 3. Our findings are depicted in Figure 7. Of the 237 subjects that participated under the matrix display and the verbal display, only 10.97% never violated any of the four axioms and therefore are classified as *S-EUT*, and 54.58% at least once exhibited behavior that is not to be interpreted unambiguously and therefore are classified as *S-Other*.[19] The remaining 34.18% of subjects fall in the two categories of primary interest, *S-Intention* and *S-Mistake*. Of these subjects, sizable 67.90% addressed any axiom violation according to decision-level category [3] and therefore are categorized as *S-Intention*, whereas only 32.10% addressed any axiom violation according to decision-level category [2] and there-

[19] Hardly surprising, by construction, the share of subjects in category *S-EUT* is lower and the share of subjects in *S-Other* is higher than in the corresponding decision-level category *EUT* (category [1]) and *Other* (categories [4], [5] and [6]), respectively.

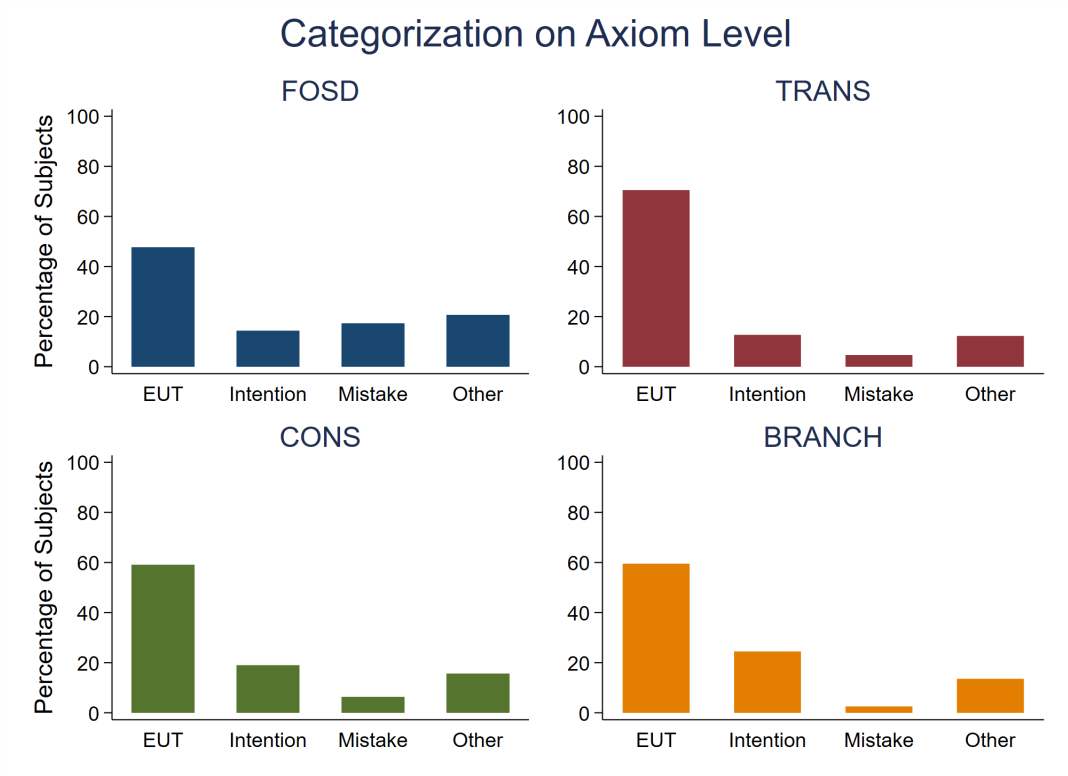


Figure 6: Categorization of behavioral patterns on the axiom level.

fore fall into the category *S-Mistake*. Thus, the ratio of these two subject-level categories is very similar to the ratio of the corresponding decision-level categories, with more than twice as many subjects violating axioms intentionally rather than by mistake – a difference that is highly significant (Binomial test, $p < 0.01$).

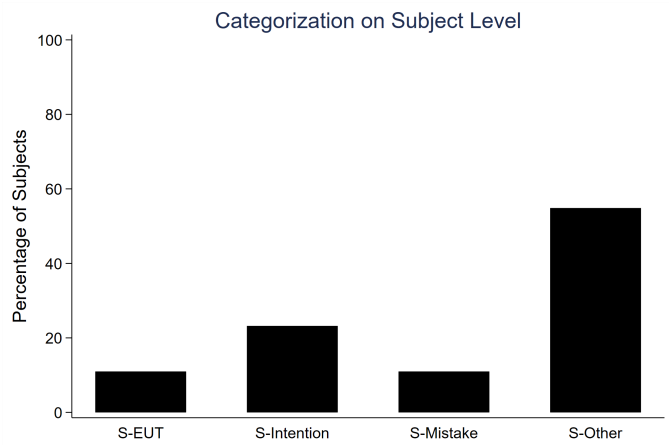


Figure 7: Categorization of behavioral patterns on the subject level.

**Result 3.** *Overall, significantly more subjects violate axioms intentionally rather than by mistake. This also holds true for all axioms separately, except first-order stochastic dominance.*

To complete the picture, we now carve out to what extent the behavior of subjects in our two subject-level categories of main interest, *S-Intention* and *S-Mistake*, differs. First of all, the average number of violations is higher in category *S-Intention* (1.61) than in category *S-Mistake* (1.31). Notably, the share of subjects violating exactly (a) two, (b) three and (c) four axioms is – for all three cases – higher in *S-Intention* than in *S-Mistake*: (a) 34.55% vs. 23.08%, (b) 10.91% vs. 3.85%, (c) 1.81% vs. 0.00%.

Subjects in the two categories also tend to violate different axioms. While the share of intentional violations of FOSD is significantly lower than the share of violations by mistake (43.64% vs. 69.23%, Fisher exact test, $p = 0.036$), the contrary is true for BRANCH (49.09% vs. 15.38%, Fisher exact test, $p = 0.004$) and CONS (40.00% vs. 19.23%, Fisher exact test, $p = 0.08$). The difference for TRANS is not significant (29.09% vs. 26.92%, Fisher exact test, $p = 1.00$).

Finally, initial axiom selection in Stage 1 is highly predictive of the axiom-level category into which subjects fall. Overall, in 89.66% of all cases, subjects do not change their axiom preferences between Stage 1 and Stage 3. This raises the concern that our finding regarding the relative size of the *S-Intention* and *S-Mistake* groups may be driven primarily by choice fatigue. However, when restricting attention to behavioral patterns from categories [2] and [3], we observe substantially more changes in axiom preferences for intentional violations (29.34%) than for violations by mistake (2.74%).

### 5.3. Robustness

We modified the experimental paradigm by Nielsen and Rehbeck (2022) to account for criticism voiced in the literature (Benjamin et al., 2021; Humphrey and Kruse, 2025) that the normativity of canonical choice axioms was portrayed too favorably. The results that we obtained within this modified framework are quite the opposite of what Nielsen and Rehbeck (2022) find: while in their experiment axioms were violated significantly more often by mistake than intentionally, in our data significantly more axiom violations were deliberate rather than by mistake. In the following, we assess the robustness of our results regarding two of the key modifications that we introduced.

**Empirical Approach.** For a given axiom, Nielsen and Rehbeck (2022) restrict their analysis of revision behavior to subjects who endorsed the axiom in question in Stage 1. This truncation of the data set has been criticized for potentially biasing conclusions as it ignores subjects who did not endorse the axiom at the outset and whose behavior therefore might be particularly indicative of an intentional violation (Humphrey and Kruse, 2025). Therefore, to attach equal weight to the questions of when axiom-violating behavior is intentional and when it is a mistake, we applied a categorization of behavioral patterns

that is not based on whether a subject initially endorsed a given axiom in Stage 1 or not. To assess whether this difference in the empirical approach indeed affects the obtained results, we next truncate the data set resulting from our matrix and verbal display on the axiom level by excluding all those behavioral patterns under which a subject did not endorse a given axiom in Stage 1. Of the original 237 patterns associated with each axiom, this truncation leaves us with 196 for FOSD, 129 for TRANS, 162 for CONS and 94 for BRANCH.

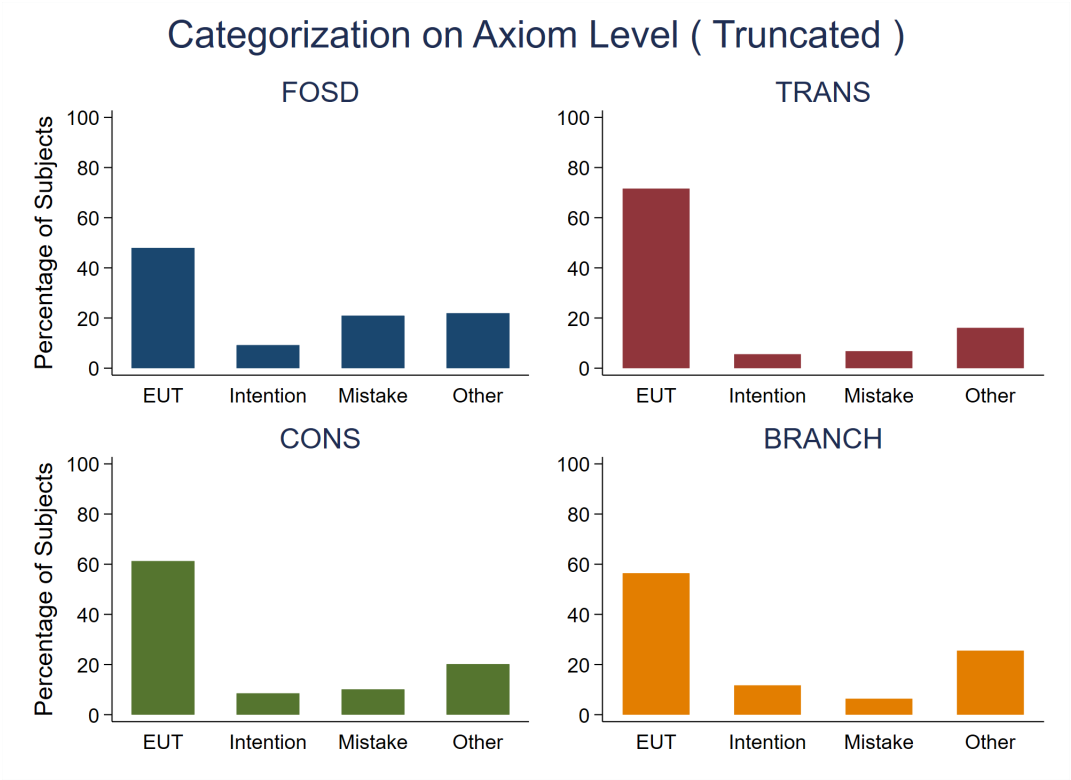


Figure 8: Categorization of behavioral patterns on the axiom level for truncated data.

Figure 8 depicts the axiom-level categorization for these remaining behavioral patterns. For those patterns that fall in our axiom-level categories [2] (i.e., *Intention*) and [3] (i.e., *Mistake*), the share of intentional axiom violations amounts to 12.8% for FOSD, 43.4% for TRANS, 45.0% for CONS and 62.5% for BRANCH. Remarkably, these shares are about 30 percentage points lower for each axiom than under the matrix and verbal display (where the share of intentional violations amounted to for 45.3% for FOSD, 73.2% for TRANS, 75.0% for CONS and 90.6% for BRANCH).

**Result 4.** *Excluding those behavioral patterns under which a given axiom was not endorsed in Stage 1 significantly shifts the scales in favor of axiom violations being portrayed as mistakes rather than intentional.*

Truncating the data set by excluding those behavioral patterns under which an axion was initially not endorsed brings our results closer to those in Nielsen and Rehbeck (2022). However, also under this alternative empirical approach, only FOSD is violated significantly more often by mistake than intentionally (Binomial test, $p<0.01$). For TRANS and CONS the slight imbalance in favor of violations by mistake is not significant (Binomial test, $p=0.82$ for TRANS and $p=0.84$ for CONS), and BRANCH still is violated more often intentionally than by mistake, but this difference is no longer significant (Binomial test, $p=0.33$).

**Display Format.** As outlined before, we used display formats to present lottery choice problems under which the correlation of the two lotteries' payoffs was conveyed transparently. The rationale for this choice of problem presentation was that earlier studies found evidence for regret avoidance to shape lottery choices under these display formats, which might foster the deliberate violation of canonical choice axioms. Nielsen and Rehbeck (2022), in contrast, used a stripped-down display format without any mentioning of how lottery payoffs are correlated; each lottery's payoffs are listed from top to bottom in increasing order together with the associated occurrence probabilities.

**OPTION A**

| | | |
|---|---|---|
| **€0** | with | 42% |
| **€9** | with | 38% |
| **€30** | with | 20% |

**OPTION B**

| | | |
|---|---|---|
| **€0** | with | 38% |
| **€9** | with | 42% |
| **€30** | with | 20% |

Figure 9: Lottery question to target FOSD with stripped-down display.

To assess the extent to which this difference in the display format affects results, we ran an additional treatment with 112 participants in which we presented our lottery choice problems in the exact same display format used by Nielsen and Rehbeck (2022). As an example, consider Figure 9, where we depict one of the lottery constellations from our experiment that target FOSD. Here, for each lottery, the left entries list the monetary amounts (in increasing order) that are potentially paid out and the right entries correspond to the occurrence probabilities associated with these outcomes.

As one would expect, with this variation of the display format being administered *after* Stage 1, original axiom endorsement rates closely resemble those from our matrix and verbal displays – cf. Figure 10a.[20] In contrast to Stage 1 behavior, it is reasonable to

[20]The four axioms are endorsed at an average rate of 64%, with relatively high selection rates of 77% for FOSD and 74% for TRANS and lower selection rates of 59% for CONS and 46% for BRANCH.

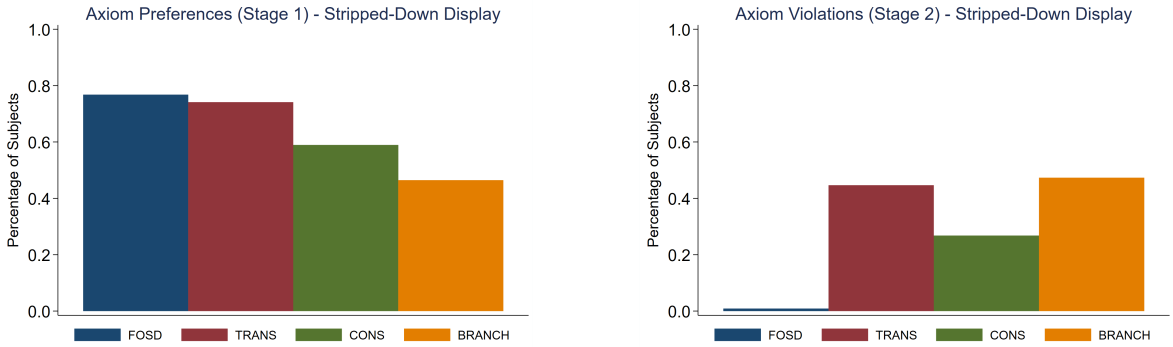


(a) Stage 1 axiom endorsement rates.

(b) Stage 2 axiom violation rates.

Figure 10: Initial axiom endorsement and violation shares for stripped-down display.

expect effects of the display format on axiom-violating lottery choices at Stage 2. As can be seen from Figure 10b, with payoffs being listed in increasing order, the dominance relationship in our pairwise lottery choice problems that targeted FOSD was very easy to detect. In consequence, it is not surprising that we observed no violations of FOSD.[21] Regarding CONS we observe fewer violations under the stripped-down display format than under the matrix and the verbal display (26.79% vs. 39.66%). This is intuitively plausible as under the stripped-down display format each lottery constellation targeting CONS involves two identical pairwise lottery choices, whereas these lottery choices differed regarding the juxtaposition of their outcomes under the matrix and the verbal display. Regarding TRANS and BRANCH, we did not have any prior expectation of how the change in the display format might affect axiom violation rates. Here, in comparison to the matrix and verbal display, we find increases in the corresponding violation rates (44.64% vs 28.27% for TRANS, 47.32% vs 40.08% for BRANCH).

Figure 11 depicts the axiom-level categorization of overall behavioral patterns (i.e., including revision decisions). Clearly, with no violations of FOSD occurring under the stripped-down display format, a meaningful comparison across display formats regarding FOSD is not feasible. Regarding the remaining axioms, for those subjects who fall in our axiom-level categories [2] (i.e., *Intention*) and [3] (i.e., *Mistake*), the share of intentional axiom violations amounts to 83.3% for TRANS, 71.4% for CONS and 86.5% for BRANCH. These shares are very similar compared to our results under the matrix and verbal display, where the share of intentional violations amounted to 73.2% for TRANS, 75.0% for CONS and 90.6% for BRANCH. Thus, our key finding that significantly more subjects violate canonical choice axioms intentionally rather than by mistake turns out remarkably robust with regard to the administered change in the display format of problem presentation.

[21] In contrast, Nielsen and Rehbeck (2022) used lotteries designed to test FOSD, where the stripped-down display format encouraged subjects to make line-by-line comparisons, likely leading them to choose the dominated option.

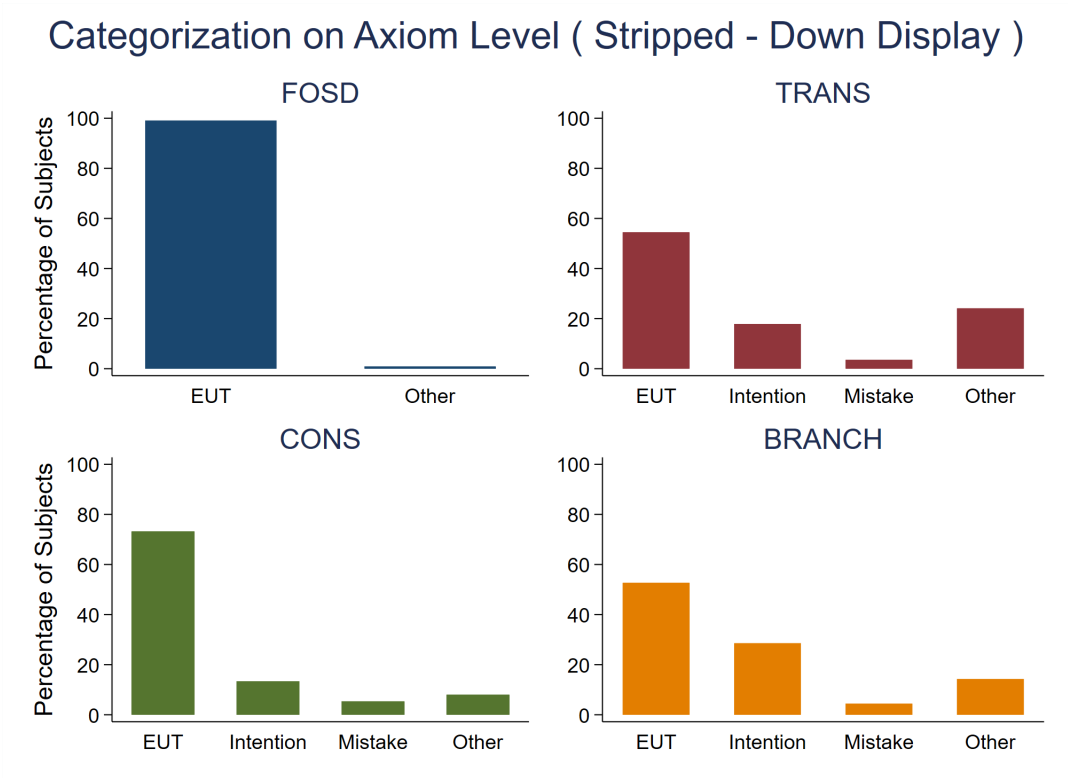


Figure 11: Categorization of behavioral patterns on the axiom level for stripped-down display.

**Result 5.** *Whether axioms are violated intentionally or by mistake does not depend on whether the correlation of lottery outcomes is conveyed by the display format of lottery choice problems or not.*

## 6. Conclusion

When facing choices under risk, human beings often make decisions that violate canonical choice axioms. This raises the question whether these violations are a manifestation of true preferences, in which case the respective decision maker rejects the normative value that economists typically ascribe to these axioms, or mistakes. Addressing this research question in a laboratory experiment, we categorize subjects whose behavior allows for a clear-cut interpretation into three groups: first, expected utility maximizers (24%); second, subjects who would like to adhere to the canonical axioms but occasionally make decisions in violation of these axioms (24%); third, subjects that violate the tested canonical axioms intentionally and thus reject their normative value (52%). Thus, roughly half of our unambiguously classifiable subjects consider the rationality axioms in question as desirable general guiding principles, whereas the other half do not. At the end of the day, we still cannot conclusively answer the question regarding the normative content of canonical axioms in choice under risk. In particular, some axioms seem to be normatively

more appealing than others (particularly FOSD). All in all, we see our results as complementary to recent empirical findings, echoing their insight that "it will take much more work to show where and when [...] violations are mistakes" (Nielsen and Rehbeck, 2022, p.2239).

## References


**Battalio, R. C., J. H. Kagel, and K. Jiranyakul**, "Testing between alternative models of choice under uncertainty: Some initial results," *Journal of Risk and Uncertainty*, 1990, *3*, 25–59.

**Benjamin, D. J., M. A. Fontana, and M. S. Kimball**, "Reconsidering risk aversion," 2021. Mimeo, NBER Working Paper 28007.

**Birnbaum, M. H. and T. Martin**, "Generalization across people, procedures, and predictions: Violations of stochastic dominance and coalescing," *Emerging Perspectives on Decision Research*, 2003, pp. 84–107.

— **and U. Schmidt**, "The impact of learning by thought on violations of independence and coalescing," *Decision Analysis*, 2015, *12*, 144–152.

**Bleichrodt, H. and P. P. Wakker**, "Regret theory: A bold alternative to the alternatives," *Economic Journal*, 2015, *125* (583), 493–532.

**Bock, O., I. Baetge, and A. Nicklisch**, "hroot: Hamburg registration and organization online tool," *European Economic Review*, 2014, *71*, 117–120.

**Bourgeois-Gironde, S.**, "Regret and the rationality of choices," *Philosophical Transactions of the Royal Society B: Biological Sciences*, 2010, *365*, 249–257.

**Breig, Z. and P. Feldman**, "Revealing risky mistakes through revisions," *Journal of Risk and Uncertainty*, 2024, *68*, 227–254.

**Castillo, G.**, "The attraction effect and its explanations," *Games and Economic Behavior*, 2020, *119*, 123–147.

**Chen, D. L., M. Schonger, and C. Wickens**, "oTree — An open-source platform for laboratory, online, and field experiments," *Journal of Behavioral and Experimental Finance*, 2016, *9*, 88–97.

**Crosetto, P. and A. Filippin**, "The "bomb" risk elicitation task," *Journal of Risk and Uncertainty*, 2013, *47*, 31–65.

_ **and A. Gaudeul**, "Fast then slow: Choice revisions drive a decline in the attraction effect," *Management Science*, 2023.

**Druckman, J. N.**, "Evaluating framing effects," *Journal of Economic Psychology*, 2001, *22*, 91–101.

**Frederick, S.**, "Cognitive reflection and decision making," *Journal of Economic Perspectives*, 2005, *19*, 25–42.

**Gilboa, I.**, "Questions in decision theory," *Annual Review of Economics*, 2010, *2*, 2–19.

**Harless, D.W.**, "Actions versus prospects: The effect of problem representation on regret," *American Economic Review*, 1992, *82*, 634–649.

**Holzmeister, F. and A. Pfurtscheller**, "oTree: The "bomb" risk elicitation task," *Journal of Behavioral and Experimental Finance*, 2016, *10*, 105–108.

**Humphrey, S. and N.-Y. Kruse**, "Who accepts Savage's axiom now?," *Theory and Decision*, 2024, *96*, 1–17.

_ **and** _ , "The normative force of Savage's axiom and the Allais paradoxes: More evidence," *Theory and Decision*, 2025, *forthcoming.*

**Loomes, G.**, "Further evidence of the impact of regret and disappointment in choice under uncertainty," *Economica*, 1988, *55*, 47–62.

_ , "When actions speak louder than prospects," *American Economic Review*, 1988, *78*, 463–470.

_ , "Predicted violations of the invariance principle in choice under uncertainty," *Annals of Operations Research*, 1989, *19*, 103–113.

_ **and R. Sugden**, "Regret theory: An alternative theory of rational choice under uncertainty," *Economic Journal*, 1982, *92*, 805–824.

_ **and** _ , "Testing for regret and disappointment in choice under uncertainty," *Economic Journal*, 1987, *97*, 118–129.

_ **, C. Starmer, and R. Sugden**, "Preference reversal: Information processing effect or rational non-transitive choice," *Economic Journal*, 1989, *99*, 140–151.

_ **,** _ **, and** _ , "Observing violations of transitivity by experimental methods," *Econometrica*, 1991, *59*, 425–439.

**_ , _ , and _** , "Are preferences monotonic? Testing some predictions of regret theory," *Economica*, 1992, *59*, 17–33.

**MacCrimmon, K. R.**, "Descriptive and normative implications of the decision-theory postulates," in K. Borch and J. Mossin, eds., *Risk and uncertainty: Proceedings of a conference held by the International Economic Association*, Palgrave Macmillan UK London 1968, pp. 3–32.

**McNeil, B. J., S. G. Pauker, and A. Tversky**, "On the framing of medical decisions," in D. E. Bell, H. Raiffa, and A. Tversky, eds., *Decision making: Descriptive, normative, and prescriptive interactions*, Cambridge University Press Cambridge 1988, pp. 562–568.

**Miller, P. M. and N. S. Fagley**, "The effects of framing, problem variations, and providing rationale on choice," *Personality and Social Psychology Bulletin*, 1991, *17*, 517–522.

**Nicholls, N., A. T. Romm, and A. Zimper**, "The impact of statistical learning on violations of the sure-thing principle," *Journal of Risk and Uncertainty*, 2015, *50*, 97–115.

**Nielsen, K. and J. Rehbeck**, "When choices are mistakes," *American Economic Review*, 2022, *112*, 2237–2268.

**Savage, L. J.**, *The foundations of statistics*, New York: Wiley, 1954.

**Sieck, W. and J. F. Yates**, "Exposition effects on decision making: Choice and confidence in choice," *Organizational Behavior and Human Decision Processes*, 1997, *70*, 207–219.

**Simon, H. A.**, "Human nature in politics: The dialogue of psychology with political science," *American Political Science Review*, 1985, *79* (2), 293–304.

**Slovic, P. and A. Tversky**, "Who accepts Savage's axiom?," *Behavioral Science*, 1974, *19*, 368–373.

**Starmer, C.**, "Testing new theories of choice under uncertainty using the common consequence effect," *Review of Economic Studies*, 1992, *59*, 813–830.

**_ and R. Sugden**, "Probability and juxtaposition effects: An experimental investigation of the common ratio effect," *Journal of Risk and Uncertainty*, 1989, *2*, 159–178.

**_ and _** , "Testing for juxtaposition and event-splitting effects," *Journal of Risk and Uncertainty*, 1993, *6*, 235–254.

**_ and _** , "Testing alternative explanations of cyclical choices," *Economica*, 1998, *65*, 347–361.

**Thomson, K. S. and D. M Oppenheimer**, "Investigating an alternate form of the cognitive reflection test," *Judgment and Decision making*, 2016, *11* (1), 99–113.

**van de Kuilen, G.**, "Subjective probability weighting and the discovered preference hypothesis," *Theory and Decision*, 2009, *67*, 1–22.

— **and P. P. Wacker**, "Learning in the Allais Paradox," *Journal of Risk and Uncertainty*, 2006, *33*, 155–164.

**von Neumann, J. and O. Morgenstern**, *The theory of games and economic behavior*, Princeton University Press, 1947. 2nd ed.

**Yu, C. W., Y. J. Zhang, and S. X. Zuo**, "Multiple switching and data quality in the multiple price list," *Review of Economics and Statistics*, 2021, *103*, 136–150.

## Online Appendix

## A. Supplementary Material to the Theoretical Framework

In the following, we formally define each of the axioms relevant for our study and explain, by means of the example of a regret-averse decision maker, how context-sensitive theories can lead to a violation of the respective axiom.[22]

We denote the set of feasible monetary payoffs by $\mathcal{X}$, the set of potential states of the world by $\mathcal{S}$, and the occurrence probability of state $s \in \mathcal{S}$ by $p(s)$. A lottery $L_i$ assigns a monetary payoff to each state of the world. Let $x_i(s)$ denote the monetary payoff assigned to state $s$, i.e., $x_i : \mathcal{S} \to \mathcal{X}$.[23] The set of all lotteries is denoted by $\mathcal{L}(\mathcal{X}, \mathcal{S})$. As usual, $\succ$ denotes the strict preference relation. Furthermore, for $L_i \in \mathcal{L}(\mathcal{X}, \mathcal{S})$ and $z \in \mathbb{R}$, define $F_i(z)$ as the probability of obtaining a payoff not larger than the amount $z$ under lottery $L_i$; i.e., $F_i(z) = \sum_{\{s \in \mathcal{S} | x_i(s) \leq z\}} p(s)$.

### A.1. First-Order Stochastic Dominance

First-order stochastic dominance requires that, if under one lottery the probability to obtain at least a particular monetary amount is weakly higher for all amounts and strictly higher for at least one amount than under the alternative lottery, then the decision maker must prefer the former lottery over the latter.

**Definition 1** (First-order Stochastic Dominance)**.** *For any two lotteries $L_i, L_j \in \mathcal{L}(\mathcal{X}, \mathcal{S})$ the following holds: If $1 - F_i(z) \geq 1 - F_j(z)$ for all $z \in \mathcal{X}$ and $1 - F_i(z') > 1 - F_j(z')$ for some $z' \in \mathcal{X}$, then $L_i \succ L_j$.*

An example for why first-order stochastic dominance may well be violated to avoid regret was given in Section 2.

### A.2. Transitivity

Transitivity requires a decision maker's pairwise choices to be consistently linked with each other in the sense that if a first lottery is preferred over a second and the second lottery is preferred over a third, then the first lottery must also be preferred over the third.

**Definition 2** (Transitivity)**.** *For any three lotteries $L_i, L_j, L_k \in \mathcal{L}(\mathcal{X}, \mathcal{S})$ the following holds: $[L_i \succ L_j \wedge L_j \succ L_k] \Longrightarrow L_i \succ L_k$.*

[22] The choice problems used to illustrate these potential axiom violations are all taken from our actual study. For a complete overview of these choice problems see Appendix B.3.

[23] Strictly speaking, the lottery $L_i$ is an *act* in the sense of Savage (1954).

Consider the three lotteries depicted in Table 4, which are based on Loomes et al. (1991).[24] Here, in the pairwise choice between $L_3$ and $L_4$, it may well be that a regret-averse decision maker chooses $L_4$ in order to avoid the post-decisional regret of receiving only €5 and missing out on €14 that she would experience if she had chosen $L_3$ and state 2 had been realized. Likewise, in the pairwise choice between $L_4$ and $L_5$, the same decision maker might choose $L_5$ in order to avoid the regret associated with receiving only €3 and missing out on €9 in case that she had chosen $L_4$ and state 3 had been realized. In the pairwise choice between $L_3$ and $L_5$, however, the decision maker might opt for $L_3$ in order to avoid the regret of receiving only €9 and missing out on €24 that she would experience if she had chosen $L_5$ and state 1 had been realized. In this example, we would have $L_5 \succ L_4$, $L_4 \succ L_3$ and $L_3 \succ L_5$ – a violation of the transitivity axiom.[25]

| | $s = 1$ | $s = 2$ | $s = 3$ |
|---|---|---|---|
| $L_3$ | €24 | €5 | €5 |
| $L_4$ | €14 | €14 | €3 |
| $L_5$ | €9 | €9 | €9 |
| Prob. | 40% | 30% | 30% |

Table 4: Lotteries $L_3$, $L_4$ and $L_5$ depicted in the action-state matrix.

### A.3. Consistency

Consistency requires that, if a first and a second lottery both represent one and the same probability distribution over the set of feasible payoffs (i.e., if the two lotteries are stochastically equivalent) and the first lottery is preferred over a third lottery, then the second lottery must also be preferred over the third lottery.

**Definition 3** (Consistency)**.** *For any three lotteries $L_i, L_j, L_k \in \mathcal{L}(\mathcal{X}, \mathcal{S})$ the following holds: If $\sum_{\{s \in \mathcal{S} | x_i(s) = x\}} p(s) = \sum_{\{s \in \mathcal{S} | x_j(s) = x\}} p(s)$ for all $x \in \mathcal{X}$, then $L_i \succ L_k \iff L_j \succ L_k$ .*

Consider the lotteries $L_6$, $L_7$ and $L_8$ depicted in Table 5, which are taken from Starmer and Sugden (1993).[26] In terms of probability distributions, $L_6$ is equivalent to $L_7$. Hence, if $L_6$ is chosen in the pairwise choice between $L_6$ and $L_8$, then consistency requires $L_7$

[24] The difference to Loomes et al. (1991) is that we double the amount of all payoffs and choose the payoffs' currency to be €.

[25] Presenting the three pairwise choice problems $L_3$ vs. $L_4$, $L_4$ vs. $L_5$, and $L_3$ vs. $L_5$ (with one-half of the payoffs in £) to 100 subjects, Loomes et al. (1991) find that 21% of their subjects exhibit the cyclical choice pattern $L_5 \succ L_4$, $L_4 \succ L_3$ and $L_3 \succ L_5$.

[26] Compared to Starmer and Sugden (1993), we once again triple the amount of all payoffs and choose the payoffs' currency to be €.

to be chosen in the pairwise choice between $L_7$ and $L_8$. A regret-averse decision maker, however, may perceive the pairwise choice between $L_6$ and $L_8$ as starkly different from the pairwise choice between $L_7$ and $L_8$. Once again, the reason for this is rooted in the juxtaposition of outcomes and the associated scope for experiencing post-decisional regret: While the scope for experiencing ex post regret when choosing $L_8$ rather than $L_6$ is rather small (in state 1 the decision maker obtains only €21 under $L_8$ and would have gotten €33 under $L_6$), the scope for experiencing ex post regret when choosing $L_8$ instead of $L_7$ might be perceived as rather large (in state 2 the decision maker obtains only €0 under $L_8$ and would have gotten €33 under $L_7$). Hence, for a regret-averse decision maker, we might well observe $L_8 \succ L_6$ and $L_7 \succ L_8$.[27]

| | $s = 1$ | $s = 2$ | $s = 3$ |
|---|---|---|---|
| $L_6$ | €33 | €0 | €0 |
| $L_7$ | €0 | €33 | €0 |
| $L_8$ | €21 | €0 | €21 |
| Prob. | 45% | 45% | 10% |

Table 5: Lotteries $L_6$, $L_7$ and $L_8$ depicted in the action-state matrix.

### A.4. Branch Independence

Branch independence requires that, if a first and a second lottery yield a given payoff $x \in \mathcal{X}$ with equal probability and the first lottery is preferred over the second lottery, then the preference must not change if the payoff $x \in \mathcal{X}$ is interchanged with some other payoff $y \in \mathcal{X}$.

**Definition 4** (Branch Independence)**.** *For any four lotteries $L_i, L_j, L_k, L_m \in \mathcal{L}(\mathcal{X}, \mathcal{S})$ and arbitrary $x, y \in \mathcal{X}$ the following holds: If there exist $\tilde{\mathcal{S}}, \hat{\mathcal{S}} \subsetneq \mathcal{S}$ such that $\sum_{s \in \tilde{\mathcal{S}}} p(s) = \sum_{s' \in \hat{\mathcal{S}}} p(s')$, $x_i(s) = x$ and $x_k(s) = y$ for all $s \in \tilde{\mathcal{S}}$, $x_i(s) = x_k(s)$ for all $s \in \mathcal{S} \backslash \tilde{\mathcal{S}}$, $x_j(s') = x$ and $x_m(s') = y$ for all $s' \in \hat{\mathcal{S}}$ and $x_j(s') = x_m(s')$ for all $s' \in \mathcal{S} \backslash \hat{\mathcal{S}}$, then $L_i \succ L_j \iff L_k \succ L_m$.*

Consider the lotteries $L_9$, $L_{10}$, $L_{11}$ and $L_{12}$ depicted in Table 6. Here, lottery $L_9$ pays out €15 in states 1 and 2, which have a combined occurrence probability of 0.5, and $L_{10}$ pays out €15 in state 3, which occurs with probability 0.5. Hence, if lottery $L_9$ is chosen in the pairwise choice between $L_9$ and $L_{10}$, then branch independence requires $L_{11}$ to be chosen in the pairwise choice between $L_{11}$ and $L_{12}$, where the payoff of €15 in states 1

[27] Presenting the two pairwise choice problems $L_6$ vs. $L_8$ and $L_7$ vs. $L_8$ (with one-third of the payoffs in £) to 90 subjects, Starmer and Sugden (1993) find that 32% of their subjects exhibit the choice pattern $L_8 \succ L_6$ and $L_7 \succ L_8$.

and 2 under $L_9$ and in state 3 under $L_{10}$ has been replaced by a payoff of €0. While the scope for experiencing ex post regret when choosing $L_9$ rather than $L_{10}$ is rather small (in state 2 the decision maker obtains only €15 under $L_9$ and would have gotten €24 under $L_{10}$), the scope for experiencing ex post regret when choosing $L_{11}$ instead of $L_{12}$ might be perceived as rather large (in state 2 the decision maker obtains only €0 under $L_{11}$ and would have gotten €24 under $L_{12}$). Hence, for a regret-averse decision maker, we might well observe $L_9 \succ L_{10}$ and $L_{12} \succ L_{11}$.

| | $s = 1$ | $s = 2$ | $s = 3$ |
|---|---|---|---|
| $L_9$ | €15 | €15 | €15 |
| $L_{10}$ | €0 | €24 | €15 |
| $L_{11}$ | €0 | €0 | €15 |
| $L_{12}$ | €0 | €24 | €0 |
| Prob. | 20% | 30% | 50% |

Table 6: Lotteries $L_9$, $L_{10}$, $L_{11}$ and $L_{12}$ depicted in the action-state matrix.

## B. Supplementary Material to the Experiment

### B.1. Instructions

*Note: The text below shows the instructions of Treatment M. Instructions displayed here are a translation into English. Original instructions were in German and are available from the authors upon request.*

**Welcome to this experiment!**

Today you will take part in an economic experiment. You can earn a substantial amount of money in this experiment. The amount of this money depends on your own decisions and on chance. Therefore, it is very important that you read the following explanations carefully.

Absolutely no communication with the other participants is allowed during the experiment. Please stop any conversations with your neighbors now and remove everything from your table that you do not need for the experiment. Anyone disobeying this rule will be excluded from the experiment and all payments.

All your income in this experiment will be calculated directly in Euros. At the end, we will pay you the total amount of Euros earned during the experiment in cash.

Today's experiment consists of **4 parts in total (Part I, II, III and IV)**. In each of these parts you can earn money. At the end of the experiment, all the money amounts from each part will be added up. In addition, you will receive **5 Euros for showing up on time** for this experiment.

At the beginning of each part you will receive detailed instructions for that part. The decisions you will make in one part will have no effect on the other parts.

On the next screen, the exact procedure of Part I of this experiment will be described. During the instructions, we will ask you to repeatedly answer some comprehension questions on the computer. These questions should make it easier for you to familiarize yourself with the decision situation.

The following applies to the entire experiment: Your decisions are anonymous. None of the other participants in the experiment will receive information about your identity during or after the experiment.

Should you have any questions at any point in the experiment, please raise your hand. We will then come to you and answer your questions.

## PART I: Instructions

Part I of the experiment consists of **a total of 3 sections (Sections A, B, and C)**. In each of these sections, you will make a series of choices.

At the end of the experiment, exactly **ONE of your decisions from ONE of the three sections** will be randomly selected by the computer. This decision will then determine the amount of money you earn in Part I of the experiment. **Each of your decisions from the entire Part I** can be selected with the **same probability**. You will learn which decision the computer randomly drew at the end of the entire experiment.

In the following, the **procedure of Section A** is described. The procedure of **Sections B and C** is explained in each case **directly before they begin**. Then you will also learn how exactly the decisions from the respective section are paid if they are drawn at random and thus become relevant for payment at the end of the experiment.

**Important:** For all sections, there are no right or wrong choices. We are only interested in what you prefer.

All three sections in Part I of today's experiment are about lotteries. To illustrate what a lottery is, imagine the following gamble: A number between 1 and 100 is rolled. Each number is equally likely. There is a monetary amount assigned to each of these numbers between 1 and 100. These amounts are between 0 Euro and 60 Euro. So there are no negative amounts of money and you cannot lose money in any case.

**Example of a lottery:** If one of the numbers from 1 to 30 is rolled, the lottery pays 8 Euros. If one of the numbers from 31 to 70 is rolled, then the lottery pays 37 Euros. If one of the numbers from 71 to 100 is rolled, then the lottery pays 0 Euros.

**Example of another lottery:** If one of the numbers from 1 to 100 is rolled, then the lottery pays 11 Euros.

## PART I - Section A: Choice Assistants

In **Section A**, you will be presented with **four** different, so-called **"Choice Assistants", one after the other**. A choice assistant is an **abstract DECISION SCHEME** that sometimes observes one or more initial choices between lotteries by you, and uses these to determine which lottery to choose in a subsequent decision. How a Choice Assistant works is represented by **colored balls** that represent **different lotteries**. The colors of the balls have no meaning in themselves, they only serve to distinguish different lotteries.

Your **task in this section** of the experiment is to determine for each Choice Assistant whether you want to **let the Choice Assistant make a decision for you**, or whether you want to **make the decision yourself**.

**Important:** Choose exactly the Choice Assistant whose decision scheme you feel **makes sense**. In this way, you can ensure that at the end of the experiment, the Choice Assistant makes a **decision in your favor without you having to make the decision yourself**. You are completely free to decide which and how many of the Choice Assistants should be used.

---

**Example 1** for the display of a possible Choice Assistant:

| OPTIONS | YOU CHOOSE | ASSISTANT CHOOSES |
|---|---|---|
| vs. | | |
| vs. | | |

Explanation of the Choice Assistant

You prefer BLUE over ORANGE. From this, the Choice Assistant concludes that you would prefer BLUE over ORANGE and therefore selects BLUE for you.

Example for Illustration:

For this example, do not think about amounts of money, but about types of ice cream.

- Imagine that there are the following possible ice cream sundaes:
  - BLUE sundae: Contains 1 scoop of latte macchiato ice cream.
  - ORANGE sundae: Contains 1 scoop of walnut ice cream.
- Imagine you have a choice between the BLUE sundae and the ORANGE sundae. Suppose you like latte macchiato ice cream better than walnut ice cream, so you choose the BLUE sundae.
- From this, the Choice Assistant concludes that if you were to choose again between the BLUE sundae and the ORANGE sundae, you would prefer the BLUE sundae and therefore selects BLUE for you.

You may want to use this Choice Assistant. The Choice Assistant uses the decision you made and would choose exactly the same as you if you made the same decision.

**Example** for Clarification: If you like latte macchiato ice cream better than walnut ice cream, it can be concluded that you will also like latte macchiato ice cream better than walnut ice cream in the future.

---

**Example 2** for the display of a possible Choice Assistant:

| OPTIONS | YOU CHOOSE | ASSISTANT CHOOSES |
|---|---|---|
| vs. | | |
| vs. | | |

Explanation of the Choice Assistant

You prefer BLUE to ORANGE. From this, the Choice Assistant concludes that you would prefer BLUE over GREEN and therefore selects BLUE for you.

Example for Illustration:

For this example, think not about amounts of money, but about types of ice cream.

- Imagine that there are the following possible ice cream sundaes:
  - BLUE sundae: Contains 1 scoop of stracciatella ice cream.
  - ORANGE sundae: Contains 1 scoop of caramel ice cream.
  - GREEN sundae: Contains 1 scoop of wild berry ice cream.
- Imagine you have a choice between the BLUE sundae and the ORANGE sundae. Suppose you like stracciatella ice cream better than caramel ice cream, so you choose the BLUE sundae.
- From this, the Choice Assistant concludes that if you had to choose between the BLUE sundae and the GREEN sundae, you would prefer the BLUE sundae and therefore selects BLUE for you.

You may not want to use this Choice Assistant. Both the blue lottery and the orange and green lotteries may contain different probabilities for different amounts of money. Just because you prefer the blue lottery over the orange lottery, you would not necessarily prefer the blue lottery over the green lottery.

**Example** for Clarification: Just because you like stracciatella ice cream better than caramel ice cream does not necessarily mean that you also like stracciatella ice cream better than wild berry ice cream.

---

While you make your choice regarding the Choice Assistants, you don't know what specific lotteries the colored balls stand for. The lotteries can be very different, some are risky, others involve safe payoffs, some tend to be lower amounts, others tend to be higher amounts.

**In summary:** A choice assistant is an abstract decision scheme that observes one or more initial choices between lotteries by you and uses them to determine which lottery to choose in a subsequent decision. So you want to select exactly the Choice Assistants that will **IN EVERY CASE** make the decision that is best for you. On the other hand, it may be worth making decisions yourself if the Choice Assistant does not necessarily choose the best lottery for you, as shown in the example above.

**Payment**

Each of your decisions for or against a Choice Assistant in Section A will have the same probability of becoming payoff relevant as any other decision you make in Part I of this experiment. Underlying each decision for or against a Choice Assistant are specific lotteries represented by the colored balls. If one of the decisions from this section is selected for payoff, there are two possibilities:

- If you decided not to use the Choice Assistant, then the associated lottery decision will be displayed at the end of the experiment and you will have to make the lottery decision yourself.
- If you have decided to apply the Choice Assistant, then you do not have to make an additional decision at the end of the experiment. Instead, you will be shown which lottery the Choice Assistant has chosen for you.

At the end of the experiment, the computer will randomly roll a number between 1 and 100. This will then determine the outcome of the lottery chosen by you or by the Choice Assistant.

## PART I - Section B: Lottery Decisions

In **Section B**, you will make a total of **20 decisions between two lotteries each**. In each of these lotteries you can receive certain amounts of money with certain probabilities. The amounts of money in the following lotteries range from 0 Euros to 60 Euros. The probabilities for the individual amounts are between 0% and 100%.

Your **task in this section** is to choose the **lottery you prefer** in each case. Each of these choices will be shown to you on a separate screen. Each time you can choose either the **upper lottery** or the **lower lottery**.

**Example** of the display of a possible lottery decision:

| | 1 30 | 31 70 | 71 100 |
|---|---|---|---|
| ○ | **€8** | **€17** | **€0** |
| ○ | **€0** | **€9** | **€9** |
| | 30% | 40% | 30% |

In order to play out a lottery, the computer randomly rolls a number between 1 and 100. For this, the lotteries assign a monetary amount to each number between 1 and 100. The **top line of the table** shown summarizes **number ranges**. The **probability** that the randomly drawn number falls into a certain number range is shown in the **bottom row of the table**.

In this example, the upper lottery pays 8 Euros for a number between 1 and 30 (30% probability), 17 Euros for a number between 31 and 70 (40% probability) and 0 Euros for a number between 71 and 100 (30% probability). The lower lottery, on the other hand, pays 0 Euros for a number between 1 and 30 (30% probability), 9 Euros for a number between 31 and 70 (40% probability) and 9 Euros for a number between 71 and 100 (30% probability).

**Payment**

Each of your lottery choices in Section B will have the same probability of becoming payoff as any other choice you make in Part I of this experiment. If any of the decisions in this section is selected for payoff, at the end of the experiment the computer will roll a **random number between 1 and 100**, and you will receive the amount of money that the lottery you selected assigns to the rolled number.

Assume that you have chosen the upper lottery in the example. If the random number is between 1 and 30, then you will receive a cash amount of 8 Euros. If the random number is between 31 and 70, then you will receive 17 Euros, and if the random number is between 71 and 100, then you will receive 0 Euros.

## PART I - Section C: Revision Opportunity

In Section A, you determined whether or not you wanted a Choice Assistant to make a decision for you. In Section B, you made lottery decisions. In doing so, it is possible that you selected a Choice Assistant in Section A that could have been applied to decisions in Section B as well.

You now have the opportunity in **Section C** to **re-evaluate your choices from the first two Sections A and B**. We will show you again your decisions on the **Choice Assistants from Section A** and at the same time the **matching lotteries from Section B** for which

the Choice Assistant could have made a decision. This way you can view your decisions from Sections A and B together. You have the possibility to compare if you have chosen a Choice Assistant and at the same time made a different choice in the lottery decisions than the Choice Assistant would have made. You do not have to make this comparison yourself. Instead, an information page will show you both graphically and verbally which choices you made in Section A and B, and which choice the Choice Assistant would have made for you.

**Example:** Here you see again a Choice Assistant that served as an example in the instructions for Section A. Assume that you have decided that this Choice Assistant **should not make a decision for you**:

| OPTIONS | YOU CHOOSE | ASSISTANT CHOOSES |
|---|---|---|
| vs. | | |
| vs. | | |

Also, here you see the lottery decision again that served as an example in the instructions to Section B. Assume that you have selected the **upper lottery** in this decision:

| | 1 – 30 | 31 – 70 | 71 – 100 |
|---|---|---|---|
| | **€8** | **€17** | **€0** |
| | **€0** | **€9** | **€9** |
| | 30% | 40% | 30% |

Last, imagine you were also facing the following lottery decision. Suppose you chose the **lower lottery** in this decision:

| | 1 – 30 | 31 – 70 | 71 – 100 |
|---|---|---|---|
| ○ | **€8** | **€17** | **€0** |
| ◉ | **€11** | **€25** | **€4** |
| | 30% | 40% | 30% |

The **information page** on these decisions would then look like this:

In the upper part of the screen you will see the choice you made in Section A regarding a Choice Assistant. In the lower part you see the respective matching lottery choices from Section B. Next to each of them is a verbal list of your choices and the choice that the Choice Assistant would have made. There are **two such information pages** for each Choice Assistant.

After these information pages you can then decide on a next screen whether you want to change your choice concerning the Choice Assistant or not. At the same time, you can

decide whether you want to change your choice regarding the lotteries or not. **You are completely free to change either none of your choices, only individual choices, or even all of the choices displayed.**

In total, you **will be given four opportunities to adjust your choices**. Each of these possibilities refers to one of the four Choice Assistants from Section A. Once again it should be emphasized here: You are completely free to decide whether to make adjustments or not, and if so, which decisions to adjust.

**Payment**

Each of your decisions for or against changing your previous choice is as likely to become payoff relevant as any other decision you make in Part I of this experiment. If any of the choices in Section C are selected for payoff, that choice will be paid off in the same way as described earlier in Section A and Section B, respectively.

**Important:** The decisions from Section C are independent. Thus, they do not overwrite any decision from Sections A and B. Therefore, both the original decision from Section A or B can be selected for payment and the possibly re-evaluated decision from Section C can be selected. So if the possibly re-evaluated choice regarding a Choice Assistant becomes relevant for payoff, and you have decided to (now) make the decision yourself, then the lottery decision will be displayed to you at the end of the experiment and you will make a decision yourself. If you have (now) decided to use the Choice Assistant, then at the end of the experiment you will be shown which lottery the Choice Assistant has chosen for you.

If the possibly re-evaluated choice regarding a lottery becomes relevant for the payoff, then the lottery (now) chosen by you will be played out in each case.

## B.2. Axiom Descriptions

*Note: The text below shows the description of axioms on the decisions screens in the first stage of the experiment, when preferences over axioms were elicited. Descriptions displayed here are a translation into English. Original descriptions where in German and are available from the authors upon request.*

### Choice Assistant for FOSD

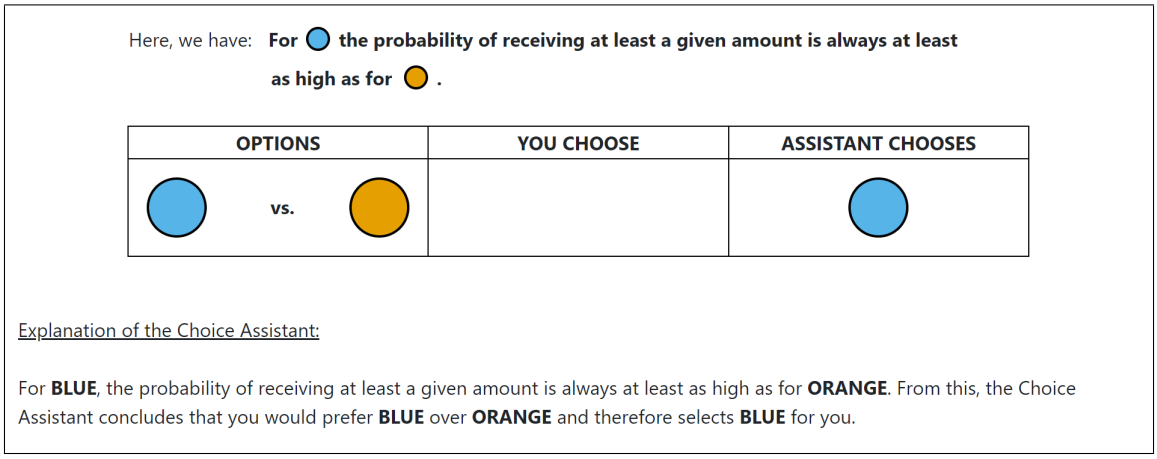


Explanation of the Choice Assistant

With **BLUE**, the probability of receiving at least a certain amount is always at least as high as with **ORANGE**. From this, the Choice Assistant concludes that you would prefer **BLUE** over **ORANGE** and therefore selects **BLUE** for you.

Example for Illustration:

For this example, do not think about amounts of money, but about types of ice cream.

- Imagine that there are the following possible ice cream sundaes:
  - **BLUE sundae:** A fair coin is tossed. Heads it contains 1 scoop of chocolate ice cream, tails it contains 2 scoops of chocolate ice cream.
  - **ORANGE sundae:** Contains 1 scoop of chocolate ice cream.
- From this, the Choice Assistant concludes that given the choice between the **BLUE** and **ORANGE** sundae, you would prefer the **BLUE** sundae and therefore selects **BLUE** for you.

### Choice Assistant for TRANS

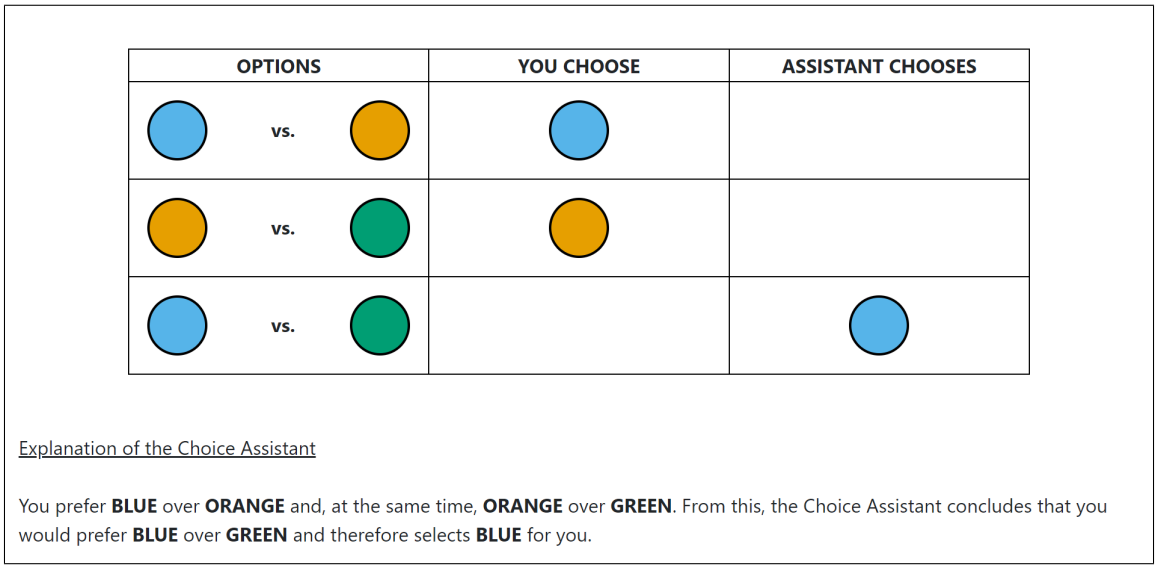


Explanation of the Choice Assistant

You prefer **BLUE** to **ORANGE** and at the same time **ORANGE** to **GREEN**. From this, the Choice Assistant concludes that you would prefer **BLUE** over **GREEN** and therefore selects **BLUE** for you.

Example for Illustration:

For this example, do not think about amounts of money, but about types of ice cream.

- Imagine that there are the following possible ice cream sundaes:
  - **BLUE sundae:** Includes 1 scoop of yogurt ice cream.
  - **ORANGE sundae:** Includes 1 scoop of pistachio ice cream.
  - **GREEN sundae:** Includes 1 scoop of cherry ice cream.
- Imagine you have a choice between the **BLUE** sundae and the **ORANGE** sundae. Suppose you like yogurt ice cream better than pistachio ice cream, so you choose the **BLUE** sundae.
- Also, imagine that you have a choice between the **ORANGE** sundae and the **GREEN** sundae. Suppose you like pistachio ice cream better than cherry ice cream and therefore choose the **ORANGE** sundae.
- From this, the Choice Assistant concludes that if you had to choose between the **BLUE** sundae and the **GREEN** sundae, you would prefer the **BLUE** sundae and therefore selects **BLUE** for you.

## Choice Assistant for CONS

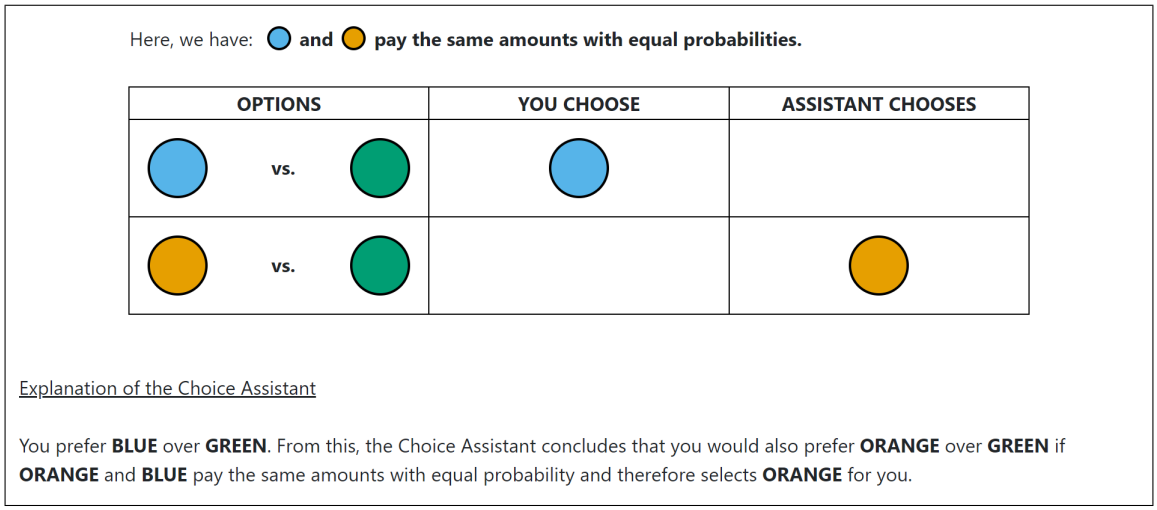


Explanation of the Choice Assistant

You prefer **BLUE** to **GREEN**. From this, the Choice Assistant concludes that you would also prefer **ORANGE** over **GREEN** if **ORANGE** paid out the same amounts with the same probabilities as **BLUE**, and therefore selects **ORANGE** for you.

Example for Illustration:

For this example, do not think about amounts of money, but about types of ice cream.

- Imagine that there are the following possible ice cream sundaes:
  - **BLUE sundae:** A fair coin is tossed. If heads, it contains 1 scoop of strawberry ice cream; if tails, it contains 1 scoop of raspberry ice cream.
  - **GREEN sundae:** Includes 1 scoop of hazelnut ice cream.
  - **ORANGE sundae:** A fair coin is tossed. Heads it includes 1 scoop of raspberry ice cream, tails it includes 1 scoop of strawberry ice cream.
- Imagine you have a choice between the **BLUE** sundae and the **GREEN** sundae. Suppose you like both strawberry ice cream and raspberry ice cream better than hazelnut ice cream, so you choose the **BLUE** sundae.
- From this, the Choice Assistant concludes that given the choice between the **ORANGE** sundae and the **GREEN** sundae, you would prefer the **ORANGE** sundae and therefore selects **ORANGE** for you.

## Choice Assistant for BRANCH

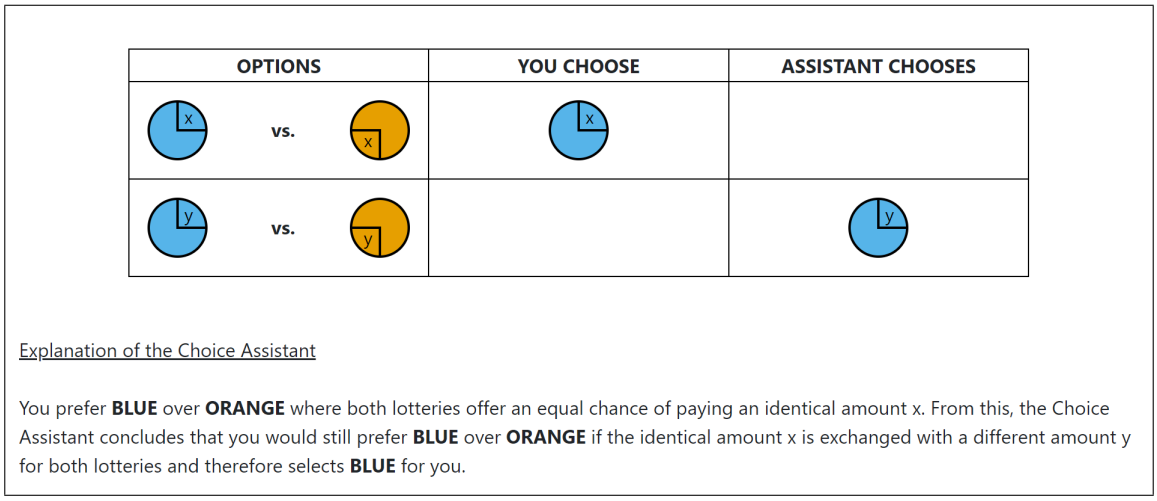


Explanation of the Choice Assistant

You prefer **BLUE** over **ORANGE**, with both lotteries paying an identical payout x with an equal chance. From this, the Choice Assistant concludes that you would continue to prefer **BLUE** over **ORANGE**, even if the identical amount x were exchanged for a different amount y in both lotteries, and therefore selects **BLUE** for you.

Example for Illustration:

For this example, do not think about amounts of money, but about types of ice cream.

- Imagine that there are the following possible ice cream sundaes:
  - **BLUE sundae:** A fair coin is tossed. On tails it contains 1 scoop of vanilla ice cream, on heads it contains 1 scoop of lemon ice cream.
  - **ORANGE sundae:** A fair coin is tossed. On tails it includes 1 scoop of lemon ice cream, on heads it includes 1 scoop of banana ice cream.
- Imagine you have a choice between the **BLUE** sundae and the **ORANGE** sundae. Suppose you like vanilla ice cream better than banana ice cream, so you choose the **BLUE** sundae.
- From this, the Choice Assistant concludes that you would still prefer the **BLUE** sundae over the **ORANGE** sundae, even if you replace the 1 scoop of lemon ice cream (x) with 1 scoop of mango ice cream (y) in both sundaes, and therefore selects **BLUE** for you.

### B.3. Overview of Choice Problems

| | $s=1$ | $s=2$ | $s=3$ | $s=4$ |
|---|---|---|---|---|
| $L_1$ | €0 | €30 | €9 | €0 |
| $L_2$ | €9 | €9 | €0 | €30 |
| Prob. | 18% | 20% | 42% | 20% |

Problem 1 (FOSD I).

| | $s=1$ | $s=2$ | $s=3$ | $s=4$ |
|---|---|---|---|---|
| $\widehat{L}_1$ | €0 | €39 | €9 | €0 |
| $\widehat{L}_2$ | €9 | €9 | €0 | €36 |
| Prob. | 20% | 20% | 40% | 20% |

Problem 2 (FOSD II).

| | $s=1$ | $s=2$ | $s=3$ |
|---|---|---|---|
| $L_6$ | €33 | €0 | €0 |
| $L_8$ | €21 | €0 | €21 |
| Prob. | 45% | 45% | 10% |

Problem 3 (CONS I.1).

| | $s=1$ | $s=2$ | $s=3$ |
|---|---|---|---|
| $L_7$ | €0 | €33 | €0 |
| $L_8$ | €21 | €0 | €21 |
| Prob. | 45% | 45% | 10% |

Problem 4 (CONS I.2).

| | $s=1$ | $s=2$ | $s=3$ | $s=4$ |
|---|---|---|---|---|
| $\widehat{L}_6$ | €0 | €0 | €0 | €42 |
| $\widehat{L}_8$ | €0 | €24 | €0 | €24 |
| Prob. | 50% | 10% | 20% | 20% |

Problem 5 (CONS II.1).

| | $s=1$ | $s=2$ | $s=3$ | $s=4$ |
|---|---|---|---|---|
| $\widehat{L}_7$ | €0 | €0 | €42 | €0 |
| $\widehat{L}_8$ | €0 | €24 | €0 | €24 |
| Prob. | 50% | 10% | 20% | 20% |

Problem 6 (CONS II.2).

| | $s=1$ | $s=2$ | $s=3$ |
|---|---|---|---|
| $L_3$ | €24 | €5 | €5 |
| $L_4$ | €14 | €14 | €3 |
| Prob. | 30% | 40% | 30% |

Problem 7 (TRANS I.1).

| | $s=1$ | $s=2$ | $s=3$ |
|---|---|---|---|
| $L_4$ | €14 | €14 | €3 |
| $L_5$ | €9 | €9 | €9 |
| Prob. | 30% | 40% | 30% |

Problem 8 (TRANS I.2).

| | $s=1$ | $s=2$ | $s=3$ |
|---|---|---|---|
| $L_3$ | €24 | €5 | €5 |
| $L_5$ | €9 | €9 | €9 |
| Prob. | 30% | 40% | 30% |

Problem 9 (TRANS I.3).

| | $s=1$ | $s=2$ | $s=3$ | $s=4$ |
|---|---|---|---|---|
| $\widehat{L}_3$ | €20 | €6 | €6 | €0 |
| $\widehat{L}_4$ | €15 | €15 | €2 | €0 |
| Prob. | 32% | 24% | 24% | 20% |

Problem 10 (TRANS II.1).

| | $s=1$ | $s=2$ | $s=3$ | $s=4$ |
|---|---|---|---|---|
| $\widehat{L}_4$ | €15 | €15 | €2 | €0 |
| $\widehat{L}_5$ | €10 | €10 | €10 | €0 |
| Prob. | 32% | 24% | 24% | 20% |

Problem 11 (TRANS II.2).

| | $s=1$ | $s=2$ | $s=3$ | $s=4$ |
|---|---|---|---|---|
| $\widehat{L}_3$ | €20 | €6 | €6 | €0 |
| $\widehat{L}_5$ | €10 | €10 | €10 | €0 |
| Prob. | 32% | 24% | 24% | 20% |

Problem 12 (TRANS II.3).

| | $s=1$ | $s=2$ | $s=3$ |
|---|---|---|---|
| $L_9$ | €15 | €15 | €15 |
| $L_{10}$ | €0 | €24 | €15 |
| Prob. | 20% | 30% | 50% |

Problem 13 (BRANCH I.1).

| | $s=1$ | $s=2$ | $s=3$ |
|---|---|---|---|
| $L_{11}$ | €0 | €0 | €15 |
| $L_{12}$ | €0 | €24 | €0 |
| Prob. | 20% | 30% | 50% |

Problem 14 (BRANCH I.2).

| | $s=1$ | $s=2$ | $s=3$ | $s=4$ |
|---|---|---|---|---|
| $\widehat{L}_9$ | €18 | €18 | €18 | €0 |
| $\widehat{L}_{10}$ | €0 | €22 | €18 | €0 |
| Prob. | 10% | 30% | 40% | 20% |

Problem 15 (BRANCH II.1).

| | $s=1$ | $s=2$ | $s=3$ | $s=4$ |
|---|---|---|---|---|
| $\widehat{L}_{11}$ | €0 | €0 | €18 | €0 |
| $\widehat{L}_{12}$ | €0 | €22 | €0 | €0 |
| Prob. | 10% | 30% | 40% | 20% |

Problem 16 (BRANCH II.2).

| | $s=1$ | $s=2$ | $s=3$ |
|---|---|---|---|
| $L_A$ | €20 | €6 | €6 |
| $L_B$ | €15 | €15 | €2 |
| Prob. | 40% | 30% | 30% |

Problem 17 (TRANS Inc.1).

| | $s=1$ | $s=2$ | $s=3$ |
|---|---|---|---|
| $L_B$ | €15 | €15 | €2 |
| $L_C$ | €10 | €10 | €10 |
| Prob. | 40% | 30% | 30% |

Problem 18 (TRANS Inc.2).

| | $s=1$ | $s=2$ | $s=3$ | $s=4$ |
|---|---|---|---|---|
| $L_D$ | €30 | €0 | €30 | €0 |
| $L_E$ | €60 | €0 | €0 | €0 |
| Prob. | 50% | 10% | 20% | 20% |

Problem 19 (CONS Inc.).

| | $s=1$ | $s=2$ | $s=3$ | $s=4$ |
|---|---|---|---|---|
| $L_F$ | €6 | €6 | €0 | €24 |
| $L_G$ | €0 | €28 | €6 | €0 |
| Prob. | 10% | 20% | 40% | 30% |

Problem 20 (BRANCH Inc.).

### B.4. Incentivization of Axiom Choices and Preference Input

The incentivization procedures were the same for original and revised axiom choices as well as for original and revised lottery choices. Lottery choices were incentivized in the standard manner; i.e., if a given lottery choice turned out to be payoff relevant, a subject was paid a realization from the lottery she selected. This realization was determined by another random draw from the computer.

With regard to the incentivization of (both original and revised) axiom choice, if a subject stated that she wants a given axiom to make decisions on her behalf and if this axiom choice was selected as the payoff-relevant decision at the end of the experiment, then the axiom was applied to a pairwise choice situation where it had implications. Any information necessary to apply the axiom – i.e., preference statements from the individual with regard to other pairwise choice situations – was elicited throughout Stage 2, where the original lottery choices were made. The subject then was paid a realization of the lottery 'chosen' by the axiom. If the subject did not select the axiom to make choices on her behalf and this decision was selected as payoff relevant, then the subject made the relevant lottery choice on her own.

Regarding the incentivization of revision decisions, subjects were explicitly told that they can change any of their previously made choices or leave them as they were. While the revision opportunity for each axiom occurred on a single screen, any decision associated with that screen (irrespective of whether the decision was to revise a previously made choice or to leave a previously made choice unchanged) had an independent chance of being selected for payment. The subject's original choices were not overturned by the revision opportunity but still could be chosen for payment.

To clarify how the information necessary to apply the axiom was elicited, consider the example of TRANS. Suppose that in Stage 1 a subject selects the TRANS axiom to make decisions on her behalf. Then, in reference to Appendix B.3, in Stage 2 this subject faces the pairwise choice between $L_A$ and $L_B$ in Problem 17 as well as the pairwise choice between $L_B$ and $L_C$ in Problem 18. Importantly, in Stage 2 the subject does not face the pairwise choice between $L_A$ and $L_C$. Suppose further that this individual chooses $L_A$ in Problem 17 (i.e., $L_A \succ L_B$) and $L_B$ in Problem 18 (i.e., $L_B \succ L_C$). Also, suppose that the subject's selection of the TRANS axiom in the Stage 1 turns out to be the payoff relevant decision. The subject then would be shown the pairwise choice between $L_A$ and $L_C$ and she would be informed that, based on her lottery decisions in Stage 2, the TRANS axiom prescribed the choice of lottery $L_A$. The monetary payoff of lottery $L_A$ would randomly be determined by the computer and paid out to the subject. If, in contrast, the subject did not select the TRANS axiom in the first stage and this decision is determined as payoff

relevant, then the subject would make the choice between $L_A$ and $L_C$ herself.[28]

[28]For the record, it could also happen that a subject selected the TRANS axiom and still had to make the choice between $L_A$ and $L_C$ herself. This happened if, in the pairwise choice between $L_A$ and $L_B$ and in the pairwise choice between $L_B$ and $L_C$ in Stage 2, she preferred both $L_A$ and $L_C$ over $L_B$ or if she preferred $L_B$ over both $L_A$ and $L_C$. In these cases the TRANS axiom "has no bite" and makes no prescription which lottery to choose.

### B.5. Post-Experimental Tests

**Risk Attitudes.** Risk attitudes were elicited with the "bomb" risk elicitation task (BRET) by Crosetto and Filippin (2013) using the ready-to-use oTree software module from Holzmeister and Pfurtscheller (2016). Subjects play one round of the BRET. They decide how many boxes of a $8 \times 8$ grid to collect, one of which contains a bomb. Between pressing the "Start" and "Stop" button, the computer would collect one box each second, starting in the upper left corner and proceeding row by row. Earnings increase by €0.30 with each box collected, but are zero if the bomb is also collected. We use the number of collected boxes as our measure of an individual's risk attitude.

**Regret Aversion.** Regret aversion was assessed in two different ways. For the first regret measure, subjects answer two questions about the amount of regret they feel in a specific situation described to them. There are 10 answer options ranging from no regret at all to extreme regret. We use the arithmetic mean over the two questions as our main measure for regret aversion. The questions read:

- Imagine you usually play the lottery every month and always choose your date of birth as the numbers. This month, however, you have forgotten to hand in your lottery ticket. You have just found out that your exact numbers have been drawn and that you would have won €50 if you had submitted your ticket. How much do you regret not having submitted your lottery ticket?
- Imagine it's the last soccer match in the championship. Your team is awarded a penalty shortly before the final whistle. If the penalty is converted, your team wins the game and the championship. If the penalty is missed, the game ends in a draw and your team does not win the championship. You step up to take the penalty and decide to shoot into the left corner. Unfortunately, the goalkeeper has also chosen the left corner and is able to save your shot. How much do you regret choosing the left corner?

For the second regret measure, subjects choose between two lotteries, were they either win €5 for sure (option A) or €80 with a probability of 4% and €0 with a probability of 96% (option B). Then they are asked, whether they want to prevent that the result of the lottery they have not chosen is displayed to them. In order to do so, the subjects indicate with which probability they want to have the result not shown, by specifying a number between 0% and 100%. We use the probability given by individuals as an additional measure for regret aversion. However, this measure is only valid for subjects that choose the safe option A and is therefore not used in regression analysis.

**Cognitiv Reflection Ability.** The subject's propensity to override an intuitive, but incorrect, response with a more analytical correct response was measured using the cognitive reflection test, using seven questions in total. For each correct answer, subjects earn €0.30. We use the average number of correct answers as our measure of an individual's cognitive reflection ability.

The three original CRT questions taken from Frederick (2005) read:

- A bat and a ball together cost 110 cents. The bat costs 100 cents more than the ball. How much does the ball cost? (intuitive answer: 10 cents; correct answer: 5 cents).
- If it takes 5 machines 5 minutes to make 5 widgets, how long would it take 100 machines to make 100 widgets? (intuitive answer: 100 min; correct answer: 5 min).
- In a lake, there is a patch of lily pads. Every day, the patch doubles in size. If it takes 48 days for the patch to cover the entire lake, how long would it take for the patch to cover half of the lake? (intuitive answer: 24 days; correct answer: 47 days)

The additional four CRT-2 questions taken from Thomson and Oppenheimer (2016) read:

- If you're running a race and you pass the person in second place, what place are you in? (intuitive answer: first; correct answer: second)
- A farmer had 15 sheep and all but 8 died. How many are left? (intuitive answer: 7; correct answer: 8)
- Emily's father has three daughters. The first two are named April and May. What is the third daughter's name? (intuitive answer: June; correct answer: Emily)
- How many cubic feet of dirt are there in a hole that is 3' deep x 3' wide x 3' long? (intuitive answer: 27; correct answer: none)

## C. Differences in the Experimental Design compared to Nielsen and Rehbeck (2022)

### Different Axioms

Next to FOSD, TRANS, CONS and BRANCH, Nielsen and Rehbeck (2022) (henceforth NR) also include the canonical choice axioms *mixture independence* and *independence of irrelevant alternatives*. We include only those axioms from NR that apply to pairwise choice and for which avoidance of regret predicts systematic violations once the correlation structure of the lotteries in question is appropriately specified.

Furthermore, to control for potential experimenter demand effects or other underlying drivers of axiom selection, NR included six so-called control axioms, which prescribe the exact reverse implication of the six canonical choice axioms of main interest, as well as six rather meaningless distractor axioms. To streamline our experiment and to avoid choice fatigue, we refrained from including control and distractor axioms. Instead, we rather take as given the insight of NR that axiom selection is not driven by experimenter demand effects, avoidance of responsibility or algorithm aversion but reflects a manifestation of genuine preferences over axioms.

### Different Depictions of Axioms

To make the axioms under consideration easily understandable, we followed NR and used pictorial logic statements to illustrate each axiom, with colored circles representing otherwise unspecified lotteries. However, our depiction of FOSD, CONS and BRANCH differs from NR. Only the pictorial logic statement to depict TRANS is identical in both studies.

Our depiction of FOSD is logically different from NR. Specifically, in accordance with the standard definition of first-order stochastic dominance, our depiction of FOSD is based exclusively on a property of the (cumulative) probability distributions of the lotteries in question and does not require any previous decision based on which the axiom then is applied. In contrast, as depicted in Figure 12, the application of FOSD in NR requires an initial choice between two lotteries which then serves as an input for the choice between two mixture lotteries that randomize over the two initial lotteries with different mixing probabilities. While our depiction of FOSD requires spelling out the dominance relationship between two lotteries, it might be considered as simpler than the depiction in NR as it does not require understanding mixture lotteries.

The pictorial logic statement for CONS in NR shows two identical decisions, see Figure 13. We, in contrast, represent in our pictorial logic statement of CONS the two lotteries with identical probability distributions over payoffs with different colors, see Appendix

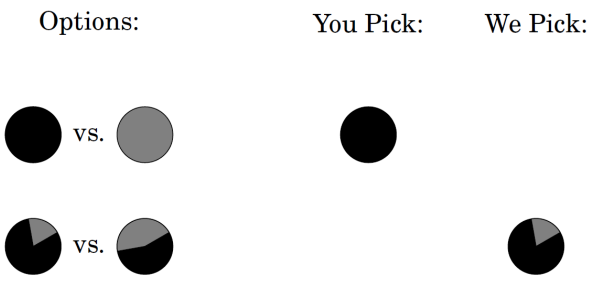


Figure 12: FOSD in NR

B.2. The reason is that the two lotteries with identical probability distributions over outcomes are not identical but different Savage acts; i.e., the juxtaposition of outcomes differs in comparison to an identical alternative lottery (the green circle in our pictorial logic statement). Note that the difference in juxtaposition of outcomes is observable only if information regarding the lotteries' correlation is provided.

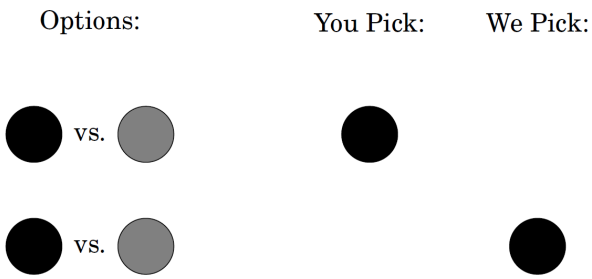


Figure 13: CONS in NR

As shown in Figure 14, not only the depiction of FOSD, but also the pictorial logic statement to explain BRANCH in NR involves mixture lotteries. Specifically, the choice between two mixture lotteries that both play out a given common lottery with the same probability is used as an input for the choice between two other mixture lotteries which are identical to the first two mixture lotteries except for the initial common lottery being replaced by a different common lottery. Our pictorial logic statement to explain BRANCH refrains from invoking the notion of mixture lotteries by indicating (within the colored circles representing lotteries) that a given outcome (denoted by $x$) which is paid out with the same probability under the initial two lotteries is replaced by a different outcome (denoted by $y$).

### Different Lottery Choices

The decision problems used by us as well as by NR are inspired by existing evidence reporting relatively high choices in violation of the targeted axiom. We relied on evidence that reported high violation rates of canonical axioms for lottery choices displayed in the

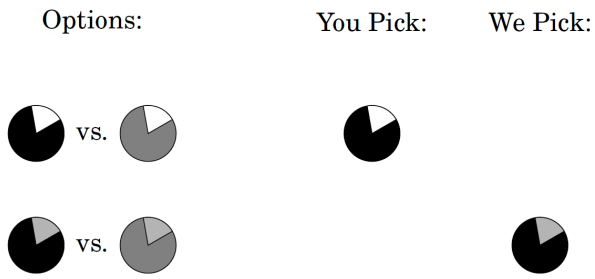


Figure 14: BRANCH in NR

matrix format, where choices that violate the axiom in question are potentially triggered by juxtaposition effects. NR, on the other hand, relied partly on studies reporting high violation rates for display formats not containing information on how the lotteries' outcomes are juxtaposed/correlated. This point can be best illustrated for the case of FOSD. All four questions asked by NR to target FOSD followed the structure in Birnbaum and Martin (2003). For example, one choice problem asked by them is the following:

| | | | | | |
|---|---|---|---|---|---|
| 10% | chance of | \$2 | 5% | chance of | \$2 |
| 5% | chance of | \$16 | 5% | chance of | \$3 |
| 85% | chance of | \$19 | 90% | chance of | \$19 |
| | **Lottery A** | | | **Lottery B** | |

Lottery B dominates Lottery A in the sense of first-order stochastic dominance. A line-wise comparison of the two lotteries may trick a subject into choosing Lottery A.

The choice problems used in our experiment rely on a particular juxtaposition of outcomes in order to trigger a violation of FOSD. Consider our Problem 1 (FOSD I) described in Appendix B.3. The fear of the high potential post-decisional regret in state $s = 4$ may lead a subject to choose the dominated option $L_2$. When this choice problem is displayed in the stripped-down display format used by NR, as we do in our robustness treatment, it takes the following form:

| | | | | | |
|---|---|---|---|---|---|
| 38% | chance of | €0 | 42% | chance of | €0 |
| 42% | chance of | €9 | 38% | chance of | €9 |
| 20% | chance of | €30 | 20% | chance of | €30 |
| | **Lottery** $L_1$ | | | **Lottery** $L_2$ | |

With this display format it is rather obvious that Lottery $L_1$ dominates Lottery $L_2$. Not surprisingly, in Treatment R, none of our subjects chose the dominated lottery and thus none violated FOSD.

To conclude, one might argue that the choice problems we use are designed to trigger juxtaposition effects that lead to choices in violation of the targeted canonical axiom. The choice problems used by NR, on the other hand, are designed to trick subjects into making lottery choices that violate the targeted canonical axiom; i.e.,they work akin to

optical illusions. These violations, by design, are rather mistakes than a manifestation of true preferences.

### Different Structures for the Revision Opportunity

We modified the design of the third stage in the experiment by NR, where subjects have an opportunity to revise their axiom preferences and lottery choices, in two central aspects: the scope and the nature of the revision opportunity.

First, regarding the scope of the revision opportunity, we do not restrict the opportunity to revise one's choices to the case where a subject's stated axiom preferences and her subsequent lottery choices turned out to be inconsistent. Instead, we allow subjects to revise *all* their axiom preferences and lottery choices. Of particular interest here to us is the case where a subject's lottery choices are necessarily consistent with her previously stated axiom preferences because the subject did not express a preference for the given axiom to make choices on her behalf in the first place. We believe that a subject violating an axiom and then sticking to her lottery choices in this case is highly indicative of an intentional violation.

Second, with regard to the nature of the revision opportunity, we implement batch revisions for each canonical axiom rather than independent revisions to underline the axiomatic character of our choice assistants. In NR, each revision decision was made independent of any other revision decisions. That is, if a subject violated a given axiom on multiple occasions, this subject had as many revision opportunities *with regard to the same axiom*. In particular, a subject that violated an axiom more than once could stick with the axiom in one instance and renounce the same axiom in another instance. In order to highlight each axiom's nature as a universally valid guiding principle, we impose batch revisions in which a subject has to decide whether to revise her decisions in the light of *all* the lottery choices that are related to a given axiom.[29]

While primarily meant to widen the scope of our analysis and to highlight the axioms' inherent character, these two modifications have the additional benefit of making the experiment more symmetric across subjects. In NR, with revisions regarding one and the same axiom being independent and with the revision opportunity being feasible only in case of lottery choices that are inconsistent with a selected axiom, the number of revision decisions ranged from 2 to 22. In our experiment, each subject faces the same number of four revision decisions.

[29]The reason for NR to opt for independent revision decisions rather than for badge revisions was to reduce the cognitive load and to avoid choice fatigue in their experiment. We do not believe that batch revisions will necessarily increase cognitive load and choice fatigue significantly in our experiment as we have significantly fewer axiom choices (4 instead of 18) and lottery choices (20 instead of 33).

### Different Wordings

First of all, our study was conducted in the German language, whereas the study by NR was conducted in the English language. This already implies modifications and differences in wordings.

More specifically, while axioms were presented in the spirit of algorithms that are capable of making choices on a subject's behalf in both studies, the wording used to refer to axioms throughout the experiment was different. Specifically, while NR labeled axioms as "rules", we referred to axioms as "choice assistants" (*Auswahlassistenten*). With our subject pool consisting mostly of German students and with Germans being widely regarded as sticklers for rules and order, we wanted to avoid that our subjects feel that they were expected to follow a rule in the sense of following a regulation or a law.

Furthermore, we opted for a somewhat more neutral wording for the revision opportunity in Stage 3. We did not directly tell subjects that their "choices were inconsistent with [the] rule" (NR, p.45 in the Online Appendix). Instead, we only reminded subjects of the lottery choices they had made in Stage 2 and told them what option the corresponding choice assistant would have chosen. Hence, in our experiment, subjects had to draw the conclusion themselves whether their lottery choices were consistent with an axiom's prescription or not.

## D. Variation of the Display Format

To account for potential effects of the format of problem presentation, our main experiment comprised two different display formats for lottery choice problems, a matrix display and a verbal display. Notably, this variation in display format was administered in Stage 2, i.e, after subjects had already stated their axiom preferences in Stage 1. Thus, we should not expect any differences between the matrix display and the verbal display with regard to subjects' axiom preferences elicited in Stage 1 of the experiment. This conjecture is confirmed by the data. Overall, subjects chose an axiom in 61.29% of cases, with no significant difference between the two display formats (p=0.69). This also holds true for each of the four axioms separately (p>0.38 for any of the four axioms).

In the light of earlier experimental studies (Battalio et al., 1990; Harless, 1992), the share of subjects who violate a given axiom in Stage 2 might be expected to be higher under the matrix display than under the verbal display as the former facilitates within-state outcome comparisons and, thus, might sensitize subjects more strongly for the potential experience of ex-post regret. While these earlier studies do not address the impact of the display format on how subjects revise initial decisions, the above reasoning suggests that, if deliberate regret avoidance is more prevalent under the matrix display than under the verbal display, more subjects should shy away from ultimately endorsing the axiom under the former display format than under the latter. In consequence, if anything, one might expect axiom selection rates in Stage 3 to be lower and axiom violation rates in Stage 3 to be higher under the matrix display than under the verbal display.

These conjectures, which are inspired by earlier experimental findings, are not confirmed by the data. When comparing axiom violation rates in Stage 2 for subjects that had at least one axiom violation, there is no significant difference between the two display formats (p=0.55), and the overall tendency goes against the hypothesized direction.[30] Also when split up by axioms, there are no differences regarding initial violation rates between display formats (FOSD: p=0.94, TRANS: p=0.55, CONS: p=0.28, BRANCH: p=0.2). Likewise, in Stage 3, there is no significant difference between the two display formats when comparing revised axiom endorsement (p=0.94) and revised axiom violation rates (p=0.25). Again, this result holds when the data is split up by axioms (axiom endorsement: FOSD: p= 0.99, TRANS: p=0.44, CONS: p=0.80, BRANCH p=0.74; axiom violation: FOSD: p= 0.94, TRANS: p=0.55, CONS: p=0.28, BRANCH p=0.20).

**Observation 1.** *There is no systematic difference between the matrix display and the verbal display regarding original lottery violation rates, revised axiom selection rates and revised axiom violation rates.*

[30] This result does not change when we additionally control for axiom choices in Stage 1 in the probit regression.

Overall, we do not find a systematic effect on decision behavior in either Stage 2 or Stage 3. Therefore, in our classification of behavioral patterns we pool the observations from the matrix display and the verbal display.