An Experimental and Theoretical Investigation of Decision-Making Under Risk

Thesis by

Aldo Lucia

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2024
Defended April 30, 2024

Aldo Lucia
ORCID: 0000-0002-4833-5948

ACKNOWLEDGEMENTS

I feel privileged to have had the unique opportunity to be a Ph.D. student at Caltech, and I am grateful to everyone who made this journey possible and so special.

I am indebted to my advisor, Charles Sprenger, for being a fantastic mentor. Charlie's arrival at Caltech during my second year was truly a turning point in my academic career, and I could not be more grateful for everything he has done for me. For similar reasons, I would like to thank Kirby Nielsen, whose continuous support has been invaluable. As I am about to begin a new academic journey, I cannot thank Kirby enough for being such a wonderful example of the type of scholar I aspire to become. I am also deeply grateful to Luciano Pomatto and Marina Agranov for all the encouragement, support, and advice I received from them over these five years.

The main reason I believe Caltech is a special place is its people. During my five years in the Ph.D. program, I had the chance to interact with and receive feedback from most of the faculty members, for which I am sincerely grateful. I would also like to thank all my peer graduate students. Special thanks go to Yimeng, who helped me in so many ways and became a great friend, despite one minor flaw: a bad taste in football. Shunto, who was probably the first person I met when I first visited Caltech, and later became a friend and co-author of my job market paper. Kate, who made many hours both inside and outside of the office much more enjoyable. I also thank her for trying hard to get me more involved in physical activity, which is probably the only thing she failed to accomplish in these five years.

I would also like to thank my former advisors at Bocconi: Simone Cerreia-Vioglio, Marco Ottaviani, and Pierpaolo Battigalli. Reading Simone's work was how I first got excited about research, and I will always be grateful to him for this. I am also grateful to Marco and Pierpaolo for their support and for everything they taught me.

Last but not least, I want to thank my family: my mother Lucia, my father Pino, my sister Gabriella, my fiancée Arianna, my grandparents, my aunts, and my uncles. Thank you all for your unwavering support and belief in me. You inspire me every day to strive to be a better person.

ABSTRACT

This dissertation comprises three chapters related to the fields of decision theory, game theory, and experimental economics. Chapters 1 and 2 use experimental and structural methods to study individual decision-making in the domain of risk, while Chapter 3 examines decision-making under risk in settings of strategic interaction.

In Chapter 1, co-authored with Shunto Kobayashi, we conduct the first experiment that studies two classical behaviors under risk inconsistent with Expected Utility together: the common ratio effect and preferences for randomization. We show that these two behaviors are strongly positively correlated in a manner inconsistent with the predictions of leading economic models and machine learning algorithms. Motivated by this observation, we develop a novel empirical approach which, unlike machine learning algorithms, imposes some basic assumptions on preferences but does not rely on specific decision models. We further demonstrate that this approach provides more accurate predictions—both inside and outside laboratory settings—compared to leading economic models and machine learning algorithms.

In Chapter 2, I design an experiment testing Expected Utility's central independence axiom and contemporaneously eliciting measures of decision confidence. Recent theoretical work implicates decision confidence as a central component of decision-making under risk, attributing failures of Expected Utility to a lack of confidence. I find that choices characterized by high self-reported levels of decision confidence and low response times are more likely to comply with the independence axiom. Contrary to the common certainty effect rationale for independence violations, I show that subjects predominantly violate Expected Utility by choosing risky lotteries over certain amounts when they are unconfident in their choices.

In Chapter 3, co-authored with Marco Loseto, we study static games in which players have convex preferences. Under convexity, players' preferences admit a conservative multi-utility representation: each utility generates an evaluation for each action, and actions are ranked according to the lowest evaluation. We characterize the set of optimal actions for players with convex preferences and propose an efficiency criterion to rank them. Next, we derive a new class of mixed Nash equilibria that we call "strict" because players strictly prefer randomization. In general, convexity may lead to a multiplicity of mixed Nash equilibria. However, we show that when they exist, only strict equilibria ensure that all mixed actions are efficient.

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

LIST OF TABLES

INTRODUCTION

Economists develop models to make predictions, and these models are particularly valuable when their predictions hold beyond the original setting in which they are derived. For instance, let's consider a financial advisor interested in determining the portfolio composition that best aligns with her client's attitude towards risk. To this end, the financial advisor can conduct an experiment eliciting choices over financial assets in the form of monetary lotteries. Next, the financial advisor can use a model to estimate the client's level of risk aversion, which can then be used to predict the optimal portfolio composition. An important message arising from this example is that a key metric for evaluating any model must be its ability to provide reliable predictions across different settings.

Given this motivation, two important questions arise: What are the tools that we typically use in economics to make predictions, and how effective are they across different settings? The first set of tools consists of economic models that have been developed to rationalize many behavioral observations, for instance in the contexts of time and risk preferences, ambiguity, and loss aversion. However, a growing body of research documents their limited ability to provide reliable predictions across settings. As an alternative approach, machine learning algorithms are gaining attention in the field of economics. These flexible tools allow researchers to make predictions without referencing any economic model, but they have also been criticized for their poor performance across different settings. Overall, poor out-of-sample performance limits the scope of these approaches as credible tools to guide decisions and influence policymaking.

In Chapter 1, co-authored with Shunto Kobayashi, we contribute to the analysis of how well we can rationalize and predict behavior under risk in two ways. First, we conduct an experiment with the main objective of evaluating the predictive performance of economic models and machine learning approaches across different settings. In particular, we look at three main prediction exercises, each characterized by a different level of similarity between the settings involved. In the first exercise, we test the ability to predict choices related to two different tests of Expected Utility. In the second exercise, we attempt to predict certainty equivalents from choices. In the third exercise, we explore the correlation between choices in the experiment and behavior outside of the experiment. The key empirical finding from the experiment is that both economic models and machine learning algorithms generate poor predictions across settings. Motivated by this finding, we develop a new empirical strategy to make predictions and we demonstrate its potential for predicting choices, certainty equivalents, as well as financial habits outside of the experiment.

One interpretation of the empirical approach developed in Chapter 1 is that choices can be classified as either "easy" or "hard". Easy choices can be predicted without committing to a specific decision model, as the predictions of all models align with those of Expected Utility. Conversely, hard choices violate the independence axiom, leading different models to offer divergent behavioral predictions. In Chapter 2, I conduct an experiment in which individuals make choices over lotteries and express how confident they are in their own choices. I show that easy choices, characterized by high self-reported confidence and shorter response times, are more likely to comply with the independence axiom. From an economic point of view, this result is significant because it shows that decision confidence systematically correlates with observed behavior and can thus be used to generate better predictions.

Chapter 3, co-authored with Marco Loseto, turns to the analysis of behavior under risk in settings of strategic interaction. In particular, we study theoretical models that link failures of Expected Utility to a lack of decision confidence in strategic interaction settings. Within these settings, we characterize a new class of Nash equilibria that we call "strict" because players in the equilibrium strictly prefer the mixed actions to the pure actions chosen with positive probability. In these models, randomization may serve as a tool to hedge against uncertainty regarding the value of outcomes, future tastes, or degrees of risk aversion. Our analysis provides testable predictions for models of deliberate randomization in strategic settings that can be explored in future experimental work.

Chapter 1

ROBUST ESTIMATION OF RISK PREFERENCES

1.1 Introduction

Economic models estimated from experimental data have the potential to guide decisions and influence policies when they accurately predict real-world behaviors. For instance, financial advisors can use surveys involving risky choices to predict the portfolio composition that best aligns with their clients' risk attitudes. Similarly, experiments on risk preferences can inform policies aimed at regulating high-risk financial behaviors, such as speculative trading or excessive borrowing. However, even within experimental settings, many economic models that successfully rationalize specific behaviors fail to provide reliable out-of-sample predictions (for discussion and examples see, e.g., Agranov, Healy, and Nielsen, 2023; Chapman et al., 2023, 2023; Dean and Ortoleva, 2015, 2019). As an alternative predictive approach, machine learning algorithms are gaining attention in the field of economics (Hofman et al., 2021). These methods generate predictions without an explicit economic model, but they have also been criticized for their poor out-of-sample performance.[1]

This paper contributes to the important question of how well we can predict behavior under risk in two ways. First, we conduct an experiment to examine two classical behaviors inconsistent with Expected Utility (EU): the common ratio effect and preferences for randomization. These behaviors have mostly been analyzed in isolation in prior experimental work. Our findings reveal that leading economic models and machine learning algorithms produce inaccurate predictions about one non-EU behavior when trained on choices related to the other. Second, motivated by this negative result, we propose a novel empirical approach to predict behavior under risk. This approach can be used to study any preference that is complete, transitive, and continuous—three common assumptions in decision-making models—without committing to a specific economic model. We then demonstrate the potential of this new approach in producing more accurate out-of-sample predictions, both within our experiment and in high-stakes behaviors outside of the lab.

Section 1.2 introduces our novel experimental design, which allows us to assess

[1] See Athey (2017) and Coveney, Dougherty, and Highfield (2016).

the capability of various approaches to concurrently rationalize and predict risky choices in two important settings. The first setting is the common ratio version of the Allais paradox, where subjects can violate EU by either displaying higher or lower risk aversion when one option is certain, compared to when all options are risky. The former behavior is referred to as the common ratio effect, and its counterpart as the reverse common ratio effect. The second setting examines preferences for lottery mixtures. Subjects can violate EU in this setting by showing a strict preference for mixing lotteries, which we call preferences for randomization, or by displaying a strict aversion to mixing lotteries, which we call aversion to randomization. In our experiment, subjects engaged in incentivized binary choice tasks involving monetary lotteries. The study was conducted on Prolific, with a total of 500 subjects recruited on July 28, 2023. We call choice tasks associated with the common ratio version of the Allais paradox "CR-tasks", and those involving mixture lotteries "R-tasks".

We focus on these two tests of EU because they contributed to the development of many non-EU models under risk—see, for instance, Cerreia-Vioglio, Dillenberger, and Ortoleva, 2015; Chew, Epstein, and U. Segal, 1991; Gul, 1991; Loomes and Sugden, 1982, Kahneman and Tversky, 1979; Tversky and Kahneman, 1992—and because of two limitations that we identified in prior experimental work. First, previous experiments have studied these EU tests either in isolation or with a predominant focus on one over the other.[2] As a result, little is known about whether a model's rationale for one non-EU behavior leads to accurate predictions for the other. Second, most prior experiments do not allow for the unambiguous elicitation of aversion to randomization.[3] This is important because aversion to randomization often emerges as the prediction of a popular model like Cumulative Prospect Theory (CPT) under standard parametrizations used to explain behavior in the common ratio

[2]To the best of our knowledge, Agranov and Ortoleva (2017) is the only paper that studies both the common ratio version of the Allais paradox and randomization within the same experiment. Their experiment primarily focuses on randomization, with the common ratio version of the Allais paradox implemented through a single pair of binary choice tasks presented to subjects at the end of the experiment.

[3]There are three approaches that are commonly used in the literature to elicit preferences for mixtures. The first approach requires subjects to choose between two lotteries multiple consecutive times (Agranov and Ortoleva, 2017; Agranov, Healy, and Nielsen, 2023; Dwenger, Kübler, and Weizsäcker, 2018). The second approach allows subjects to delegate their choice to an external randomization device (Agranov and Ortoleva, 2017; Cettolin and Riedl, 2019; Sandroni, Ludwig, and Kircher, 2013; Dwenger, Kübler, and Weizsäcker, 2018). The third approach presents a choice between two lotteries and some mixtures between them (Agranov and Ortoleva, 2023; Dwenger, Kübler, and Weizsäcker, 2018; Feldman and Rehbeck, 2022; Miao and Zhong, 2018; Sopher and Narramore, 2000). Agranov and Ortoleva (2022) provide an overview of these methods. None of them allows for the direct revelation of an aversion to randomization.

version of the Allais paradox.

Our experiment employs a within-subjects design to study the common ratio version of the Allais paradox in CR-tasks and randomization in R-tasks, using a common set of lotteries. Moreover, to enable ourselves to unambiguously identify both preference for and aversion to randomization, we design R-tasks following a procedure proposed by Camerer and Ho (1994). In particular, we evaluate attitudes towards mixtures between any two lotteries s and r through two distinct R-tasks. In the first task, subjects compare lottery s to a 50-50 mixture of lotteries s and r. In the second task, they compare the 50-50 mixture to lottery r. If a subject chooses the 50-50 mixture in both tasks, this indicates a preference for randomization. On the other hand, if the subject avoids the 50-50 mixture in both tasks, we infer an aversion to randomization.

Section 1.3 summarizes the main findings of the experiment and explores the ability of popular economic models and machine learning algorithms to rationalize them. The common ratio effect and preferences for randomization are the two most frequent non-EU behaviors observed in CR-tasks and R-tasks, respectively. In particular, the common ratio effect accounts for around 63% of all non-EU behaviors in CR-tasks, while preferences for randomization account for around 55% of all non-EU behaviors in R-tasks. Moreover, for fixed values of probabilities and prizes of the lotteries, the percentage of non-EU behavior in CR-tasks attributed to the common ratio effect is strongly positively correlated with the percentage of non-EU behavior in R-tasks attributed to preferences for randomization, with a correlation coefficient of 0.63.

After documenting the emergence of and the correlation between non-EU behaviors in the experiment, we turn to analyzing whether popular economic models and machine learning algorithms can accommodate them. We consider EU and CPT as economic models, and gradient boosting trees (GBT) and neural networks (NN) as machine learning algorithms.[4] In particular, we implement two types of out-of-sample exercises. First, we evaluate the ability of various approaches to predict choices in each EU test separately, employing 5-fold cross-validation within CR-tasks and R-tasks. Second, we assess the ability of different approaches to predict the correlation between non-EU behaviors across CR-tasks and R-tasks. To this end, we use choices in CR-tasks as the training set and choices in R-tasks as the test set, and vice versa.

[4] Details about the functional form assumptions we made on economic models and about the training procedures for the machine learning algorithms can be found in Appendix A.2.

Machine learning algorithms predict non-EU behavior within CR-tasks and R-tasks better than economic models. In particular, GBT correctly classifies more than 60% of the observations in cross-validation exercises, while the best-performing economic model, CPT, correctly classifies only around 46% of the observations. However, the ranking between machine learning algorithms and economic models is reversed when attempting to predict one non-EU behavior from another. In these exercises, the predictive accuracy of machine learning models drops by about 30%, raising concerns about their ability to produce generalizable predictions. Moreover, CPT predicts a negative rather than a positive correlation between the common ratio effect and preferences for randomization, performing worse than EU. Overall, EU turns out to be the best model at predicting choices across CR-tasks and R-tasks, despite the fact that non-EU behaviors account for more than 34% of the observations in CR-tasks and more than 47% of the observations in R-tasks.

Our experiment sheds light on a new important instance of a more general problem: economic models and machine learning algorithms often provide poor predictions across settings. Motivated by this observation, Section 1.4 turns to the methodological contribution of the paper, which is the development of a new approach to estimate preferences under risk and make predictions. Intuitively, estimating a model of decision-making under risk requires making a series of assumptions about preferences. Some assumptions are less controversial, as they are embodied in most models. For instance, most models assume that preferences are complete, transitive, and continuous. Other assumptions are more controversial, and this is the case with the independence axiom. The independence axiom is the most empirically challenged assumption of EU, and the common ratio effect and preference for randomization are two examples of behaviors inconsistent with this assumption. Our approach relies on the less controversial assumptions, as it may be more likely that they hold across different settings, while remaining silent about the independence axiom.

Rather than estimating a representation for a preference, which requires making assumptions about the independence axiom, we estimate a representation for its EU core, which is its largest subrelation that satisfies the independence axiom (Cerreia-Vioglio, 2009). Assuming our preference of interest is complete, transitive, and continuous, its EU core can be represented by a set of utility functions.[5] This set generates an upper and lower bound for risk aversion, which becomes wider as the

[5]In particular, Cerreia-Vioglio, Maccheroni, and Marinacci (2017) clarifies that the EU core is represented by the set of "local utilities" introduced by Machina (1982).

preference's inconsistency with the independence axiom increases. We demonstrate that the set of utilities representing the EU core can be easily estimated using standard experimental datasets that examine the independence axiom, and we present the estimation results from our experiment. By estimating mixture models, we observe a significant degree of heterogeneity in terms of risk aversion and adherence to EU. In Section 1.5, we assess whether this heterogeneity can be fruitfully exploited to make reliable predictions across different settings.

We perform various types of out-of-sample exercises, distinguished by the differences in the features of training and test data. Initially, we revisit the two out-of-sample exercises within and across CR-tasks and R-tasks introduced in Section 1.3. Because our empirical approach applies to all models that relax EU by violating the independence axiom, it remains silent about how exactly the independence axiom fails. In other words, we can predict the emergence of EU and non-EU behavior, but we cannot differentiate between specific non-EU behaviors. To establish a fair comparison between our approach and other predictive methods, we evaluate the ability of all methods to predict adherence to EU. Consistent with our findings in Section 1.3, machine learning algorithms exhibit better performance when predicting non-EU behaviors in isolation but fall short when predicting one non-EU behavior using choices related to the other as training data. Importantly, our structural model achieves the best performance in these latter exercises, demonstrating potential advantages in making predictions that are not tied to specific economic models.

To provide an additional setting for testing the predictive performance of various approaches, we elicited certainty equivalents for three binary lotteries at the end of our experiment. Using choices from CR-tasks and R-tasks as training data, our method predicts ranges of possible certainty equivalents, which we demonstrate to encompass most of the observed certainty equivalents. Furthermore, we calculate point predictions for certainty equivalents as the midpoints of these predicted ranges.[6] To compare the accuracy of our approach against other methods, we employ mean squared errors and find that machine learning algorithms underperform compared to economic models, while our approach yields the most accurate out-of-sample predictions.

[6] We abstract away from the question of what theoretical assumptions on preferences justify this aggregation rule. Nevertheless, given that this rule outperforms both the economic models and the machine learning algorithms that we consider, we find its theoretical analysis an interesting direction for future work.

Finally, we explore whether the heterogeneity in risk aversion and adherence to EU, identified through our mixture model using experimental data, correlates with risky behaviors in real-world scenarios. To pursue this, we gathered information on subjects' investment and insurance behaviors. In the context of investments, we assessed whether they had engaged in stock trading and held any cryptocurrencies. Additionally, we verified whether they had insured purchased items, such as mobile phones. This is a classic example of small stakes risk aversion, which can be challenging to rationalize with EU (Rabin, 2000). Our findings reveal that subjects identified as more risk averse in the experiment are less likely to invest, particularly in cryptocurrencies. Moreover, those identified as less consistent with EU are more likely to insure purchased items. Overall, these correlations hint at our structural model as a potentially useful tool for predicting behavior beyond experimental settings.

Our paper has two broad implications for future research, which we expand upon in our concluding Section 1.6. First, our novel experimental design enables us to document the robust positive correlation that exists between the common ratio effect and preferences for randomization. Popular behavioral models, like CPT, cannot rationalize this correlation, and there is a clear need for new research to develop superior predictive approaches. A natural path forward involves considering alternative theories. However, numerous non-EU models have already been developed over the past decades, and their ability to rationalize and predict various behaviors under risk has proven to be somewhat limited. Our paper proposes a different solution, developing a structural estimation approach that is not tied to specific economic models.

Second, our out-of-sample analysis reveals that the optimal approach for making predictions may depend on at least two factors. The first factor concerns the differences between training and test sets. Machine learning algorithms exhibit superior performance when training and test sets contain choices related to the same non-EU behavior. However, their performance significantly deteriorates in all other out-of-sample exercises we examine. The second factor pertains to the analyst's required level of detail in a prediction. Our model generates predictions that, while potentially more accurate, are less detailed than those produced by conventional predictive approaches. For example, it can only predict a range of possible certainty equivalents for a lottery. If the analyst requires more detailed predictions, our approach can still serve as a valuable complement to conventional tools. For instance, it can be used to generate a range of certainty equivalents, after which economic models or machine learning algorithms can be employed to select a value from this range.

1.2 Experimental Design

The key objective of the experiment is to evaluate the predictive performance of economic models and machine learning algorithms across different settings. To achieve this objective, we have designed an experiment involving binary choice tasks between monetary lotteries. The first two settings that we consider are different tests of the independence axiom: one is the common ratio version of the Allais paradox, and the other assesses attitudes toward mixing lotteries. Moreover, as additional settings for predictions, we elicit certainty equivalents for three binary lotteries and collect information about subjects' financial habits.[7]

Binary Choice Tasks

We elicit choices between lotteries over three monetary prizes $L < M < H$. We represent the three-outcome lottery that gives $\$L$ with probability p_L, $\$M$ with probability p_M, and $\$H$ with probability p_H as $(\$L, p_L; \$M, p_M; \$H, p_H)$. Moreover, we denote by δ_X the degenerate lottery that pays $\$X$ with certainty. The independence axiom imposes consistency requirements on choices across two or more binary choice tasks. We first assess the independence axiom via the common ratio version of the Allais paradox, which involves two types of binary choice tasks that we call CR-tasks:

CR1: $\delta_M = (\$M, 1)$ vs. $r = (\$L, 1 - p_H; \$H, p_H)$.

CR2: $0.3\delta_M + 0.7\delta_L = (\$L, 0.7; \$M, 0.3)$ vs. $0.3r + 0.7\delta_L = (\$L, 1 - 0.3p_H; \$H, 0.3p_H)$.

We use the Marschak–Machina (MM) triangle to describe graphically the lotteries in the experiment (Marschak, 1950; Machina, 1982). The left graph in Figure 1.1 illustrates the CR-tasks in the MM triangle. In the MM triangle, the probability of receiving the highest prize H is on the vertical axis, and the probability of receiving the lowest prize L is on the horizontal axis. Therefore, the generic point (p_L, p_H) in the MM triangle represents the lottery $(\$L, p_L; \$M, 1 - p_L - p_H; \$H, p_H)$. Each dashed segment connecting two lotteries indicate that there is a choice task that

[7] We preregistered the experimental design and the analysis plan at the AEA RCT Registry as AEARCTR-0011749 (Kobayashi and Lucia, 2023).

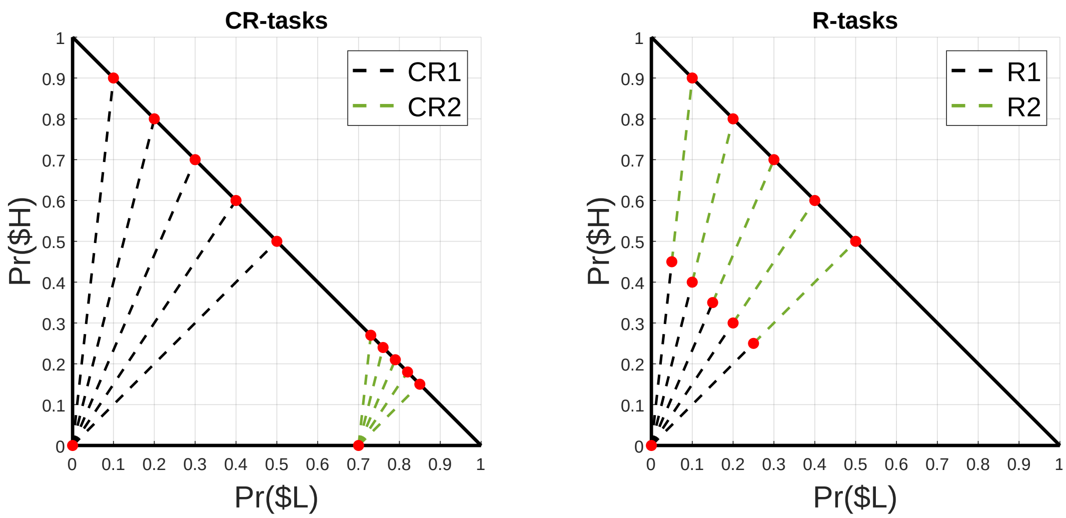

Figure 1.1: Choice Tasks.

involves these lotteries. For instance, the black dashed segments in the left MM triangle of Figure 1.1 represent CR1 choice tasks, while the green dashed segments represent CR2 choice tasks.

There are two possible scenarios in which subjects' choices in CR-tasks are incompatible with the independence axiom. The common ratio effect (CRE) refers to the violation of the independence axiom in which subjects in the experiment choose lottery δ_M in CR1, and lottery $0.3r + 0.7\delta_L$ in CR2. The opposite choices in CR1 and CR2 constitute the other possible violation of the independence axiom, known as the reverse common ratio effect (RCRE).

As an additional assessment of the independence axiom, we also study subjects' attitudes towards mixtures, i.e., randomization. To this end, we consider the following two types of binary choice tasks that we refer to as R-tasks:

R1: $\delta_M = (\$M, 1)$ vs. $0.5\delta_M + 0.5r = (\$L, 0.5(1 - p_H); \$M, 0.5; \$H, 0.5p_H)$.

R2: $r = (\$L, 1 - p_H; \$H, p_H)$ vs. $0.5\delta_M + 0.5r = (\$L, 0.5(1 - p_H); \$M, 0.5; \$H, 0.5p_H)$.

The right MM triangle in Figure 1.1 represents the R1 choice tasks (depicted by black dashed segments) and the R2 choice tasks (depicted by green dashed segments). In studies exploring preferences for randomization, it is common to combine R1 and R2 into a single choice task in which subjects can select either lottery δ_M, lottery r, or a combination of the two.[8] Choosing a mixture of lotteries δ_M and r is

[8]Agranov and Ortoleva (2022) provide an overview of this literature.

typically interpreted as a preference for randomization. However, this approach has a limitation: it does not allow us to observe whether subjects exhibit aversion to randomization, meaning they prefer either of the lotteries δ_M and r over the mixture. By treating R1 and R2 as separate choice tasks, we can observe both preferences for and aversion to randomization.[9] Specifically, subjects in the experiment display a preference for randomization when they consistently choose the lottery $0.5\delta_M + 0.5r$, and aversion to randomization when they consistently reject the lottery $0.5\delta_M + 0.5r$. Both a preference for randomization and an aversion to it are behaviors that violate the independence axiom.

All subjects engage in the CR1, CR2, R1, and R2 choice tasks involving five different prize triplets (L, M, H): (0, 15, 30), (5, 15, 25), (10, 20, 30), (15, 20, 25), and (0, 10, 20). For each of these triplets, subjects undertake all types of tasks with five different probability values for the high prize p_H: 0.5, 0.6, 0.7, 0.8, and 0.9. Thus, subjects face both CR-tasks and R-tasks at precisely the same values of (L, M, H, p_H). In addition to these 100 choice tasks, the experiment includes two choice tasks in which one lottery stochastically dominates the other (referred to as FOSD choice tasks), and three additional types of choice tasks used to elicit certainty equivalents:[10]

MPL1: $(\$X, 1)$ vs. $(\$0, 0.5; \$20, 0.5)$ for $X \in \{3, \dots, 13\}$.

MPL2: $(\$X, 1)$ vs. $(\$5, 0.5; \$25, 0.5)$ for $X \in \{8, \dots, 18\}$.

MPL3: $(\$X, 1)$ vs. $(\$10, 0.5; \$30, 0.5)$ for $X \in \{13, \dots, 23\}$.

We choose not to incorporate choice tasks between certain amounts and a given lottery into a list, as is typically done using the multiple price list (MPL) method. This design decision is made to minimize the amount of instruction that subjects need to understand, retaining binary choice tasks as the sole method for expressing their preferences.[11] We use the certainty equivalents elicited from MPL1, MPL2, and

[9] Camerer and Ho (1994) first uses this approach to distinguish between violations of betweenness, which imposes neutrality over mixing lotteries, and violations of transitivity.

[10] We exclude from the analysis any subjects who violate first-order stochastic dominance more than once. This happens for three subjects only.

[11] Different procedures to elicit risk preferences may result in different observed behavior. Freeman, Halevy, and Kneeland (2019) find that embedding a pairwise choice between a certain monetary amount and a risky lottery in a choice list increases the proportion of subjects choosing the risky lottery.

Table 1.1: Summary of the experimental design.

	Block 1					Block 2		
	CR1	CR2	R1	R2	FOSD	MPL1	MPL2	MPL3
# Tasks	25	25	25	25	2	11	11	11
Order Tasks	Randomized					MPL1, MPL2, MPL3		
Order Blocks	Always First					Always Second		

MPL3 choice tasks to further assess the out-of-sample accuracy of the predictions derived from the empirical analysis of the EU core. Table 1.1 provides a summary of our experimental design. The choice tasks in the experiment are divided into two blocks: Block 1 and Block 2. Block 1 comprises the choice tasks used to test the independence axiom (CR1, CR2, R1, and R2), along with the FOSD choice tasks. The 102 choice tasks in Block 1 are presented to subjects in a randomized order at the beginning of the experiment.

Upon completing Block 1, subjects then proceed to complete the remaining choice tasks in Block 2 (MPL1, MPL2, and MPL3), specifically designed to elicit certainty equivalents. In Block 2, subjects first encounter MPL1 tasks, followed by MPL2 tasks, and ultimately MPL3 tasks. Within each task type in Block 2, the monetary amounts are presented in ascending order.

Recruitment and Experimental Payments

We recruited 500 subjects from Prolific on July 28, 2022. The experiment was conducted using oTree. Our sample consisted of United States citizens who possessed at least a high-school education and maintained a high approval rate on Prolific. For each subject, we collected information about gender, age, income, insurance purchases, and investment behavior.

Each subject received $5.50 upon completing the experiment. Additionally, every subject had a one-in-six chance of being selected to receive an additional bonus payment based on their decisions during the study. Out of the 135 choice tasks, each carried an equal probability of determining the bonus payment amount. Specifically, subjects received the realized amount from the lottery they chose in the randomly selected choice task.[12]

[12] The complete instructions with screenshots from the experiment are presented in Appendix A.5.

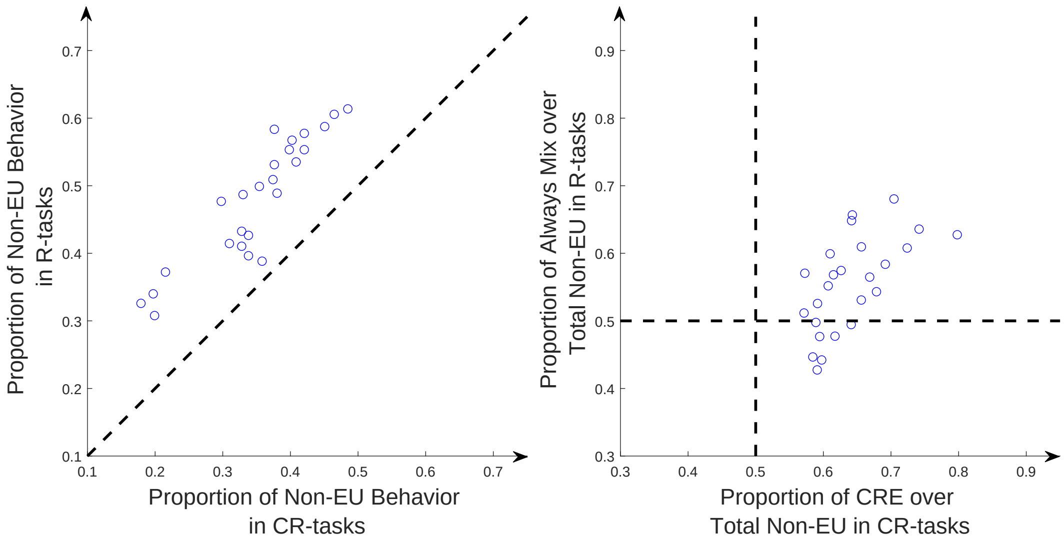

Figure 1.2: Correlation between the overall proportions of non-EU behaviors (left graph) and between the relative proportions of specific non-EU behaviors (right graph) across CR-tasks and R-tasks. Each observation corresponds to a single parameterization (L, M, H, p_H).

1.3 Main Experimental Findings

In this section, we explore the non-EU behaviors observed in CR-tasks and R-tasks, along with the correlation between them. The left graph in Figure 1.2 illustrates the proportion of non-EU choice patterns in CR-tasks (on the x-axis) and R-tasks (on the y-axis) for each pair of lotteries (δ_M, r) in the experiment.[13] Each observation corresponds to a single parameterization (L, M, H, p_H). Given that there are five triplets of prizes in the experiment and five values of p_H for each triplet, there are 25 observations graphed. All observations are situated to the left of the 45-degree dashed line, indicating that non-EU behavior is more prevalent in R-tasks than in CR-tasks for all pairs of lotteries. Moreover, there is a strong positive correlation between non-EU behaviors in CR-tasks and R-tasks, evidenced by a correlation coefficient of 0.9104. This high correlation is noteworthy because it implies that observing EU failures in specific contexts, such as tests of the CRE, may provide insights into other potential EU failures of interest.

Next, we turn to analyzing which of the non-EU behaviors are more likely to be observed in our experiment. For each pair of lotteries (δ_M, r), two potential non-EU choice patterns exist in both CR-tasks and R-tasks. In the context of CR-tasks, these patterns are represented by the CRE and the RCRE. Conversely, within R-tasks,

[13] Appendix A.1 provides a comprehensive description of EU and non-EU choice patterns in CR-tasks and R-tasks.

the two non-EU choice patterns manifest as always choosing the mixture (AM) and never choosing the mixture (NM).

The right graph in Figure 1.2 displays the proportion of CRE choice patterns over the total non-EU behaviors in CR-tasks on the x-axis, and the proportion of AM choice patterns over the total non-EU behaviors in R-tasks on the y-axis. Two important observations arise from Figure 1.2. First, CRE and AM are the most frequent non-EU choice patterns for the majority of pairs of lotteries in the experiment. At the same time, the emergence of NM choice patterns is non-negligible for many pairs of lotteries. This suggests that experiments which do not account for the elicitation of aversion to randomization might neglect a critical dimension in the analysis of attitudes towards randomization. The second fact documented in Figure 1.2 is the strong positive correlation between the prevalence of the CRE and of AM choice patterns, with a correlation coefficient of 0.6311.

Result 1. *Non-EU behaviors in CR-tasks and R-tasks are strongly positively correlated, with a correlation coefficient of 0.9104. Among non-EU behaviors, the CRE and AM are the two most frequently observed ones in CR-tasks and R-tasks, respectively. Moreover, these two non-EU behaviors are strongly positively correlated, with a correlation coefficient of 0.6311.*

Out-of-Sample Predictions: Conventional Toolbox

We now investigate whether leading economic models and machine learning algorithms can predict the observed emergence of the two non-EU behaviors and correlation between them. As popular economic models, we analyze EU model and CPT, which aims to rationalize non-EU behaviors through probability weighting. In addition, we examine the predictive performance of GBT and NN, which are two common machine learning algorithms used for classification tasks. We contrast the performance of EU, CPT, GBT, and NN through two out-of-sample exercises.[14] In the first, we separately analyze behavior in CR-tasks and R-tasks. For each task type, we assess the out-of-sample performance of various methods through cross-validation. In the second exercise, we use choices from either CR-tasks or R-tasks as the training data and aim to predict choices in the alternate tasks.

For all economic models, we estimate mixture models to account for heterogeneity. We also provide information about each subjects to machine learning algorithms

[14] Details on the procedures we followed to estimate economic models and train machine learning algorithms can be found in Appendix A.2. Instead, Appendix A.3 reports the values of the estimated parameters under EU and CPT.

Table 1.2: Out-of-sample exercises within CR-tasks and R-tasks: predict choice patterns.

Exercise	**Pattern**	**Models**			
		EU	CPT	GBT	NN
Train: 80% CR-tasks Test: 20% CR-tasks	Choice Patterns	50.62% (0.82)	46.01% (1.13)	**63.46%** (1.23)	46.30% (0.78)
Train: 80% R-tasks Test: 20% R-tasks	Choice Patterns	40.95% (1.73)	47.06% (1.22)	**60.14%** (0.56)	38.74% (1.06)
Combined	Choice Patterns	45.78% (5.25)	46.54% (1.24)	**61.80%** (1.97)	42.39% (4.22)

Notes: Percentages of correctly classified choice patterns. Standard deviations in parentheses.

including individual fixed effects in the training data. We use as main metric to evaluate these methods the percentages of choice patterns that they correctly classify. Because these methods provide probabilistic predictions, we identify the predicted choice patterns as the ones with the highest associated predicted probability.[15]

Cross-Validation within CR-Tasks and R-Tasks

We employ a fivefold cross-validation, treating data from CR-tasks and R-tasks independently. To divide the available data into five equally sized subsets, we use the following randomized procedure: For each subject and every triplet of prizes in the experiment, we randomly allocate all choice tasks associated with one of the five possible probabilities for the high prize p_H in the risky lottery to each of the five subsets. In doing so, we ensure that each iteration of the cross-validation procedure contains sufficient information about all subjects within both training and test sets. This approach enables accounting for heterogeneity in preferences during the training phase and leverages this heterogeneity to formulate predictions.

Table 1.2 presents the average percentages of choice patterns that are accurately classified across the five iterations of the cross-validation exercises within CR-tasks and R-tasks, with standard deviations in parentheses. The first two rows display the results for the two cross-validation exercises separately, while the last row provides the aggregate results from both exercises. GBT clearly emerges as the winner from

[15]We refer to Appendix A.4 for a more comprehensive probabilistic and deterministic evaluations of all predictive tools. The findings presented in this section are qualitatively similar to the ones obtained in Appendix A.4.

Table 1.3: Out-of-sample exercises across CR-tasks and R-tasks: predict choice patterns.

Exercise	Prediction	Models			
		EU	CPT	GBT	NN
Train: CR-tasks Test: R-tasks	Choice Patterns	40.72%	**42.70%**	35.71%	35.64%
Train: R-tasks Test: CR-tasks	Choice Patterns	**48.59%**	46.29%	34.16%	21.79%
Combined	Choice Patterns	**44.66%** (5.56)	44.50% (2.54)	34.93% (1.10)	28.72% (9.79)

Notes: Percentages of correctly classified choice patterns. Standard deviations in parentheses.

this out-of-sample analysis, delivering markedly superior performance compared to both NN and economic models.

Result 2. *GBT surpasses all other approaches in cross-validation exercises within CR-tasks and R-tasks, achieving an average percentage of correctly predicted choice patterns close to 61%. For context, CPT, the runner-up approach, showcases a percentage of correctly predicted choice patterns that is approximately 15% lower.*

Out-of-Sample Predictions across CR-Tasks and R-Tasks

We conduct two distinct exercises: In the first, CR-tasks serve as the training data and R-tasks as the test data. In the second exercise, conversely, we reverse the roles, employing R-tasks for training and CR-tasks for testing. The percentages of correctly classified choice patterns from these exercises are presented in Table 1.3. The first two rows of Table 1.3 display the results for the two exercises separately, while the last row summarizes the average performance across both exercises.

The descriptive analysis of behaviors in CR-tasks and R-tasks highlight systematic violations of EU. Yet, in these out-of-sample exercises, EU has the best overall out-sample performance. CPT performs worse than EU because it predicts a correlation between non-EU behaviors that is opposite to what we observe in our experiment. In particular, CPT generally rationalizes the CRE through probability weighting. However, probability weighting also implies aversion to randomization for the lotteries in the experiment, while we observe preferences for randomization as the most frequent non-EU behavior.

Machine learning algorithms also demonstrate notably poor performance, with percentages of correctly classified choice patterns around 30% lower than those achieved in cross-validation exercises within CR-tasks and R-tasks. The primary distinction between out-of-sample exercises within a specific type of task and those across tasks is that in the former, lotteries in the training and test sets are identical, while in the latter, they are marginally different. Specifically, choices in R-tasks involve mixture lotteries with three prizes, whereas all lotteries in CR-tasks feature one or two possible prizes. In general, one might expect machine learning algorithms to perform less effectively as the differences between training and test sets increase. However, the strikingly different performance of these approaches in out-of-sample exercises across choices with very similar lotteries raise concerns about their ability to produce generalizable predictions.

The economic models and the machine learning algorithms that we consider fail in the out-of-sample exercises that are most interesting from an economic point of view. If a predictive approach is not capable of predicting different behaviors within the same experiment, then there is little hope that it can be used as guidance for predicting real-world behaviors of interest. The lack of generalizability of conventional predictive approaches calls for the development of new empirical strategies to make predictions.

Result 3. *In cross-validation exercises across CR-tasks and R-tasks, economic models significantly outperform machine learning algorithms. Both EU and CPT allow for approximately 44.5% of choice patterns to be correctly classified across the two out-of-sample exercises. GBT, while achieving the best performance among machine learning algorithms, has a percentage of correctly classified choice patterns that is, on average, around 10% below that of economic models.*

1.4 Empirical Framework

In this section, we begin by reviewing the classic empirical approach used to estimate economic models from choice data. Next, we introduce the notion of the EU core, and we detail the empirical strategy that we developed for estimating it. Finally, we present our estimation results, which will then be used in Section 1.5 to make out-of-sample predictions.

Classic Estimation Approach

In the classic empirical framework proposed by Hey and Orme (1994), a decision-maker chooses lottery p over lottery q if

$$V(p,q) + \epsilon \geq 0,$$

where $V(p,q)$ is a quantity that is greater or equal to zero if the decision-maker prefers lottery p over lottery q, and ϵ is an error term, assumed to be normally distributed with a mean of zero and a variance of σ. The specific functional form of $V(p,q)$ depends on the assumptions made about the decision-making model. For instance, under EU, $V(p,q)$ represents the difference in expected utilities between lottery p and lottery q.

Within this framework, the analyst must select a decision model and also make parametric assumptions within that chosen model. For example, under EU, a common assumption is that the decision-maker operates with a constant relative risk-aversion (CRRA) utility function. With this assumption, the analyst can estimate the free parameter of the CRRA utility function and the variance of the error term using choice data. The estimated model can then be used to undertake counterfactual exercises, such as predicting choices between lotteries in an alternative dataset.

The analysis in Section 1.3 highlights the risks associated with specifying everything about a decision model. For instance, CPT rationalizes the CRE with probability weighting. However, probability weighting also implies aversion to randomization in the experiment, which is inconsistent with what we find. At the same time, the failures of machine learning algorithms in out-of-sample exercises across different types of tasks also emphasize the potential benefit of predictive approaches that rely on an underlying economic structure.

The methodological question we address is: can we generalize the empirical framework of Hey and Orme (1994) to predict behavior under risk without committing to specific decision models? The empirical approach that we propose as an answer to this question builds on the theoretical notion of EU core.

The EU core

Given any reflexive, transitive, and continuous preference relation $\succsim$, its EU core is the subrelation[16] $\succsim^*$ such that for all lotteries p, q, r and for all $\lambda \in (0,1]$,

$$p \succsim^* q \Leftrightarrow \lambda p + (1-\lambda) r \succsim \lambda q + (1-\lambda) r.$$

[16] $\succsim^*$ is a subrelation of $\succsim$ if for all lotteries p and q, $p \succsim^* q$ implies $p \succsim q$.

That is, $p \succsim^* q$ whenever the decision-maker prefers p to q and mixing both lotteries with a third common lottery r does not affect the relative preferences of the decision-maker between p and q.[17] Cerreia-Vioglio (2009) proves that $\succsim^*$ is the greatest subrelation of $\succsim$ that satisfies the independence axiom; that is, if $\succsim^{**}$ is another subrelation of $\succsim$ that satisfies the independence axiom, then $\succsim^{**}$ is a subrelation of $\succsim^*$. If a preference $\succsim$ violates the independence axiom, then its EU core is an incomplete preference relation and admits a multi-utility representation. In particular, there exists a set of utilities $\mathcal{W}$ such that for all lotteries p and q, we have $p \succsim^* q$ if and only if the difference in expected utilities between p and q is non-negative for all utilities within the set $\mathcal{W}$.

Our empirical approach involves obtaining information about a preference by estimating the set of utility functions that represent its EU core.[18] In this way, we can obtain estimates and generate predictions that do not rely on specific decision models.

Econometric Specification

We detail our empirical framework in the context of our experimental design. Our study involves lotteries over a finite set of K monetary prizes $X = \{x_1, \ldots, x_K\}$, with $x_1 < x_2 < \cdots < x_K$. We consider a set of L utility functions $\mathcal{W} = \{v_1, \ldots, v_L\}$, each utility $v_l \colon X \to \mathbb{R}$ is representable as a vector with its k-th component, v_{lk}, being equal to $v_l(x_k)$. We restrict our attention to normalized sets of utilities, setting $v_{11} = \cdots = v_{L1} = 0$ and $v_{1K} = \cdots = v_{LK} = 1$.[19] This means all utilities assign zero to the worst outcome x_1 and one to the best outcome x_K. Moreover, we assume all utilities are weakly increasing; i.e., $v_{l1} \leq v_{l2} \leq \ldots \leq v_{lK}$ for each $v_l \in \mathcal{W}$.

We define $I = \{1, \ldots, N\}$ as a set of subjects in our experiment, $\triangle(X)$ as the set of lotteries over X, and by $\mathcal{D} \subseteq \triangle(X)^2$ as a subset of pairs of lotteries where the subjects express their preferences. An empirical analysis of the EU core requires evaluating whether it holds that $p \succsim_i^* q$ or $q \succsim_i^* p$ for each subject $i \in I$ and each pair of lotteries $(p, q) \in \mathcal{D}$. In the traditional estimation framework, where the objective is to estimate a subject's preferences, the choices made by the subjects

[17] We study the EU core by considering only "one-stage" lottery mixtures, rather than two-stage compound lotteries. In other words, we focus on mixture independence, rather than compound independence, as defined in Uzi Segal (1990).

[18] This set of utility functions is unique up to the closed convex hull. Our empirical approach aims to estimate the extreme points of the convex set of utility functions that represent the EU core.

[19] In our estimation procedure, the number of utilities L is an hyperparameter that can be chosen using standard model selection techniques.

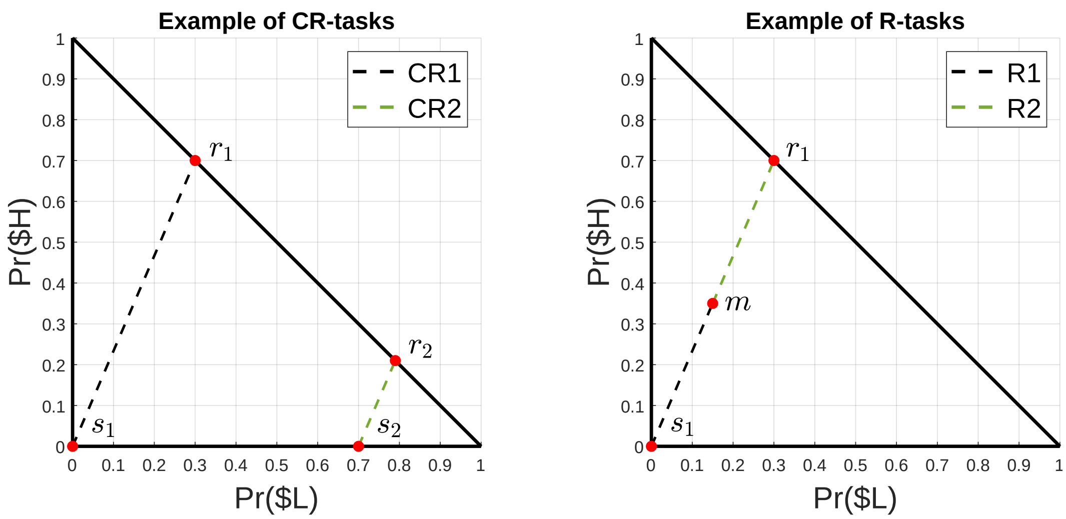

Figure 1.3: Examples of CR-tasks and R-tasks.

can be used directly as inputs for the estimation. However, when the focus of the estimation shifts from a preference relation to its EU core, additional information becomes necessary. Specifically, we need to assess whether the choices made by the subjects signal a violation of the independence axiom.

By observing subjects' choices in experimental settings that test the independence axiom, we construct an index, $Core_i$, for each subject i as follows: for each pair of lotteries $(p, q) \in \mathcal{D}$,

$$Core_i(p, q) := \begin{cases} 3 & \text{if there is no experimental evidence against } p \succsim_i^* q \\ 2 & \text{if there is no experimental evidence against } q \succsim_i^* p \\ 1 & \text{otherwise.} \end{cases}$$

We construct two distinct versions of the index $Core_i$ within the context of our experiment: $Core_i^{CR}$ and $Core_i^{R}$. Each version builds upon different sets of information. Specifically, $Core_i^{CR}$ uses data solely from CR-tasks, while $Core_i^{R}$ is based on R-tasks only. We explain how we construct these indexes using the examples from Figure 1.3. The left and right graphs in Figure 1.3, respectively, depict the CR-tasks and R-tasks associated with a generic pair of lotteries (s_1, r_1) from our experiment.

The absence of experimental evidence contradicting $s_1 \succsim_i^* r_1$ in CR-tasks means that subject i chooses for lottery s_1 in the CR1 task and lottery s_2 in the CR2 task. Observing this choice pattern, we assign the value of 3 to the index $Core_i^{CR}(s_1, r_1)$. In a parallel manner, we find no evidence refuting $r_1 \succsim_i^* s_1$ in the CR-tasks if subject

i chooses lottery r_1 in the CR1 task and lottery r_2 in the CR2 task. In this case, we assign to the index $Core_i^{CR}(s_1, r_1)$ the value of 2. The two remaining choice patterns are the CRE and the RCRE, and are both incompatible with EU. If we observe either the CRE or the RCRE for subject i, we assign to the index $Core_i^{CR}(s_1, r_1)$ the value of 1.

Turning to R-tasks, we find no evidence against $s_1 \succsim_i^* r_1$ if subject i chooses lottery s_1 in the R1 task and lottery m in the R2 task. When this happens, we assign the value of 3 to the index $Core_i^R(s_1, r_1)$. Similarly, if subject i chooses lottery m in the R1 task and lottery r_1 in the R2 task, we have no evidence against $r_1 \succsim_i^* s_1$. In this case, we assign the value of 2 to the index $Core_i^R(s_1, r_1)$. The two remaining non-EU choice patterns are AM and NM. If we observe either AM or NM for subject i, we assign to the index $Core_i^R(s_1, r_1)$ the value of 1.

We define $V(p, q; v_l)$ as the difference in expected utilities between lottery p and lottery q, given a Bernoulli utility function v_l. For each subject $i \in I$, utility $v_l \in \mathcal{W}$, and comparison $(p, q) \in \mathcal{D}$, we associate an error term $\epsilon_{i,l,(p,q)}$. We assume that the vector of error terms $[\epsilon_{i,1,(p,q)}, \ldots, \epsilon_{i,L,(p,q)}]$ across utilities follows a multivariate normal distribution with mean $[0, \ldots, 0] \in \mathbb{R}^L$ and covariance matrix $\Sigma \in \mathbb{R}^{L \times L}$. For any two lotteries p and q, and for any subject i, our empirical framework postulates that

$$Core_i(p, q) = 3 \Leftrightarrow V(p, q, v_l) - \epsilon_{i,l,(p,q)} \geq 0, \text{ for all } l \in \{1, \ldots, L\},$$

and

$$Core_i(p, q) = 2 \Leftrightarrow V(p, q, v_l) - \epsilon_{i,l,(p,q)} < 0, \text{ for all } l \in \{1, \ldots, L\}.$$

In other words, we postulate to find no evidence against $p \succsim_i^* q$ whenever the difference in expected utilities between lotteries p and q, minus an error term, is non-negative for all utilities. Similarly, we expect to find no evidence against $q \succsim_i^* p$ whenever the opposite condition holds.

Our flexible formulation of the error structure extends the normality assumption of the unique error term in Hey and Orme (1994), allowing us to account for potential noise in the $Core_i$ index that might arise from several sources. First, we construct this index by observing the choices of subject i in experimental settings that test the independence axiom. If these choices are noisy, then the resulting $Core_i$ index will also be noisy. Additionally, even in the absence of noise in the choices, the $Core_i$ index might still be noisy due to issues with missing data. For example, we might find no evidence against $p \succsim_i^* q$ simply because we could not observe enough choices involving lotteries p and q.

To account for variation in preferences across subjects, we employ a mixture model and postulate that each subject i belongs to one of C possible different groups (Bruhin, Fehr-Duda, and Epper, 2010). We denote by v_l^c the l-th utility in group c and by Σ_c the covariance matrix in group c, with $c \in \{1, \ldots, C\}$. Within this framework, the probability that we find no experimental evidence against $p \succsim_i^* q$ if subject i belongs to group c is:

$$\Pr(Core_i(p, q) = 3 \mid v_1^c, \ldots, v_L^c, \Sigma_c) = \Phi\left(V(p, q; v_1^c), \ldots, V(p, q; v_L^c); \Sigma_c\right),$$

where Φ represents the cumulative distribution function of the mean-zero multivariate normal distribution. Similarly, the probability that we find no experimental evidence against $q \succsim_i^* p$ if subject i belongs to group c is:

$$\Pr(Core_i(p, q) = 2 \mid v_1^c, \ldots, v_L^c, \Sigma_c) = \Phi\left(-V(p, q; v_1^c), \ldots, -V(p, q; v_L^c); \Sigma_c\right).$$

Given the observed index $Core_i(p, q)$ for all pairs of lotteries $(p, q) \in \mathcal{D}$, we denote by $f(Core_i; v_1^c, \ldots, v_L^c, \Sigma_c)$ the likelihood function for subject i belonging to group c:

$$\begin{aligned} \prod_{(p,q)\in\mathcal{D}} \Big(& \mathbb{1}\left(Core_i(p, q) = 3\right) \cdot \Pr(Core_i(p, q) = 3 \mid v_1^c, \ldots, v_L^c, \Sigma_c) \\ &+ \mathbb{1}\left(Core_i(p, q) = 2\right) \cdot \Pr(Core_i(p, q) = 2 \mid v_1^c, \ldots, v_L^c, \Sigma_c) \\ &+ \mathbb{1}\left(Core_i(p, q) = 1\right) \cdot (1 - \Pr(Core_i(p, q) = 3) - \Pr(Core_i(p, q) = 2))\Big). \end{aligned}$$

Let π_c represent the probability of a subject belonging to group type c. The log-likelihood of the finite mixture model is given by:

$$\sum_{i=1}^{N} \ln \sum_{c=1}^{C} \pi_c f(Core_i; v_1^c, \ldots, v_L^c, \Sigma_c),$$

where the first sum is over subjects and the second sum is over groups.

To sum, for each group of subjects c, structural parameters are utility functions $\{v_1^c, \ldots, v_L^c\}$, covariance matrices Σ_c, and probability of group member ship π_c. We estimate these parameters through maximum likelihood estimation by using the log-likelihood given above.

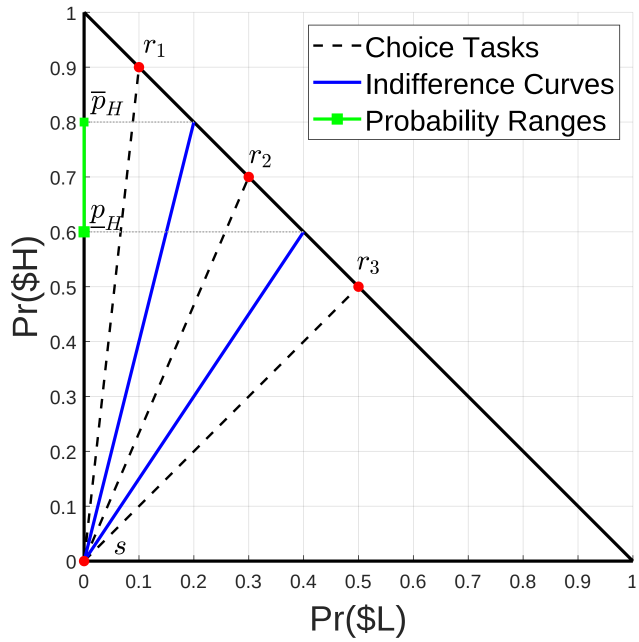

Figure 1.4: Set of utility functions in the MM triangle.

Model Implications: Risk Aversion and Adherence to EU

By using an example, we explain the behavioral implications of our empirical framework. Figure 1.4 illustrates three generic choice tasks in the MM triangle: (s, r_1), (s, r_2), and (s, r_3). To obtain predictions from any decision model in these tasks, one only needs to know the shape of the indifference curve passing through lottery s. Lotteries to the left of this curve will be strictly better than s, while those to the right will be strictly worse.

Under our decision model, each subject makes their decision based on not just one but a set of utility functions. This means that there are multiple indifference curves passing through a given lottery, each corresponding to a different utility function. For instance, the blue solid segments in Figure 1.4 represent two indifference curves passing through lottery s. In this new setting, any lottery r that lies to the left of both indifference curves will be unambiguously preferred to s. Specifically, for every such lottery, we can conclude that $r \gtrsim^* s$, as is the case for lottery r_1 in Figure 1.4. Conversely, any lottery r that lies to the right of both indifference curves will be unambiguously worse than s. In particular, for every such lottery, we can conclude that $s \gtrsim^* r$, as is the case for lottery r_3 in Figure 1.4. For all other lotteries, without making further assumptions, the only conclusion we can draw is that neither $s \gtrsim^* r$ nor $r \gtrsim^* s$ holds.

More generally, given a set of utilities $\mathcal{W}$ and any lottery r on the hypotenuse of the

MM triangle assigning probability p_H to the high prize H and probability $1 - p_H$ to the low prize L, we can compute the range of probabilities $[\underline{p}_H, \overline{p}_H]$, where

$$\underline{p}_H := \max \{p_H \in [0,1] \ : \ v(M) \geq p_H v(H) + (1 - p_H)\, v(L) \ \text{ for all } v \in \mathcal{W}\},$$

and

$$\overline{p}_H := \min \{p_H \in [0,1] \ : \ p_H v(H) + (1 - p_H)\, v(L) \geq v(M) \ \text{ for all } v \in \mathcal{W}\}.$$

These two quantities are well defined because utilities are assumed to be weakly increasing and continuous. In particular, $\underline{p}_H$ is the highest probability of the high prize for which we have $s \gtrsim^* r$. Similarly, $\overline{p}_H$ is the lowest probability of the high prize for which we have $r \gtrsim^* s$. The green squares in Figure 1.4 describe the two points $\underline{p}_H$ and $\overline{p}_H$ given the two indifference curves.

The range of probabilities $[\underline{p}_H, \overline{p}_H]$ is a measure for the extent to which a subject adheres to EU. In general, the narrower the range of probabilities, the more consistent the underlying preference is with EU. Under EU, the range would collapse into a fixed probability p_{EU}. Moreover, the range of probabilities $[\underline{p}_H, \overline{p}_H]$ provides information about subjects' risk attitudes. Specifically, in our experiment, the mid-value prize M is always set as the mean of the high prize H and the low prize L. Consequently, given a triplet of prizes, we can classify an EU subject as risk averse if $\underline{p}_H > 0.5$, risk seeking if $\overline{p}_H < 0.5$, and as neither risk averse nor risk seeking otherwise. In this way, we generalize the empirical analysis of risk attitude under EU to preferences that may violate the independence axiom.[20] In the special case of EU with a unique probability p_{EU}, we would classify a subject as risk averse if $p_{EU} > 0.5$, risk neutral if $p_{EU} = 0.5$, and risk seeking if $p_{EU} < 0.5$.

Estimation Results

We present the estimation results from a mixture model with three groups of subjects and two utilities for each group. The estimates are derived using data from both CR-tasks and R-tasks.[21] The three graphs on the left in Figure 1.5 show the estimated utilities in each group, while the three graphs on the right display the ranges of probabilities $[\underline{p}_H, \overline{p}_H]$ that they induce for the five triplets of prizes in our experiment.

The two utilities in Group 1 are very close to each other, indicating that the behavior of subjects in Group 1 can be accurately described by EU. Furthermore, both utilities

[20] Our classification adopts the aversion to mean-preserving spreads as the primitive notion for risk aversion (Rothschild and Stiglitz, 1970).

[21] In particular, we used both the $Core_i^{CR}$ and $Core_i^{R}$ indices as input for the estimation.

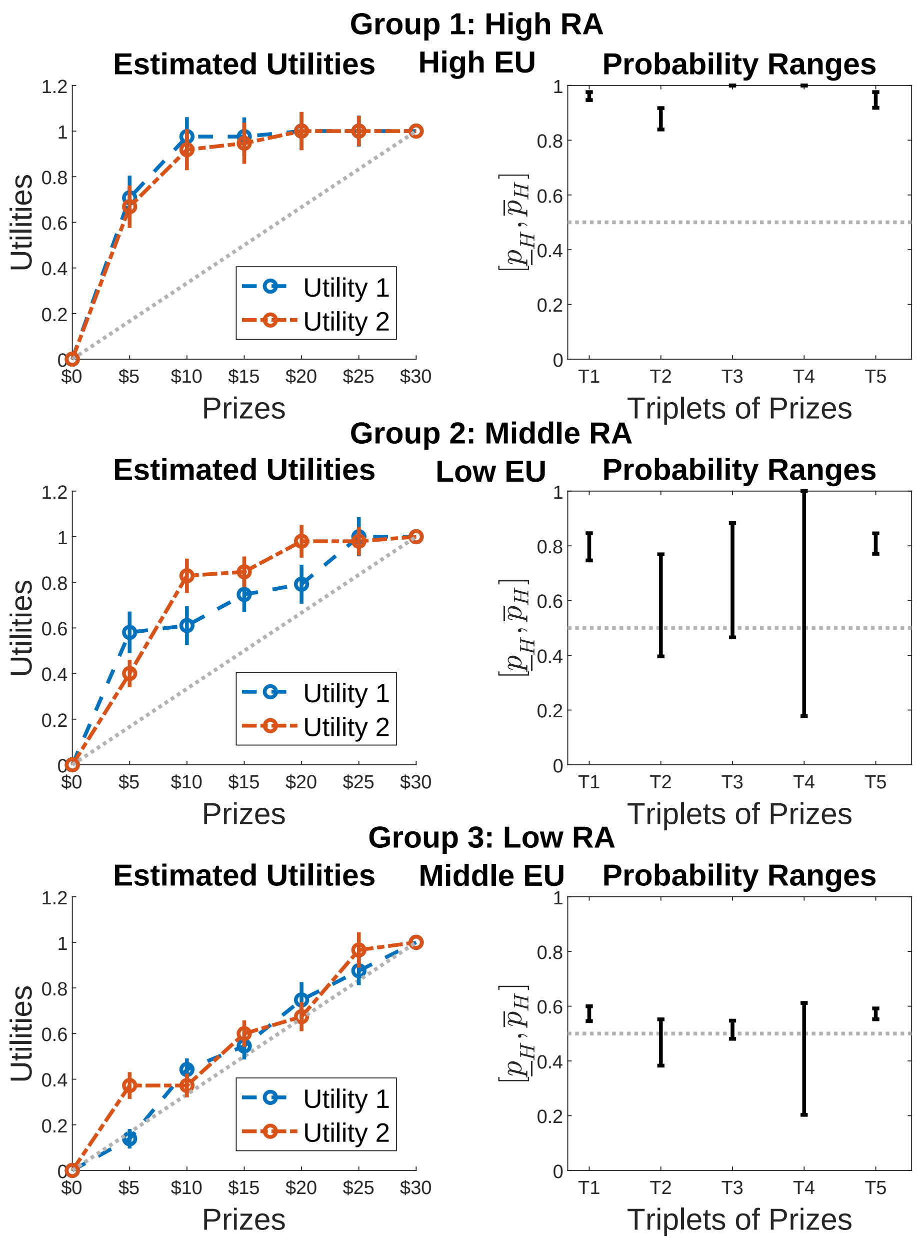

Figure 1.5: Estimated utilities and probability ranges.

Notes: We classify the three groups in terms of risk aversion (RA) and adherence to EU. Estimated group membership probabilities are: 0.2798 for Group 1, 0.4834 for Group 2, and 0.2368 for Group 3. The width of the vertical bars centered around the point estimates in the left graphs indicates the bootstrapped standard errors. We report the values of all estimated parameters in Appendix A.3. In the right graphs, T1 through T5 correspond to prize triplets as follows: T1 = ($0, $15, $30), T2 = ($5, $15, $25), T3 = ($10, $20, $30), T4 = ($15, $20, $25), and T5 = ($0, $10, $20).

in Group 1 are concave and significantly deviate from being linear, pointing to strong levels of risk aversion. This information is reflected in the range of probabilities induced by the two utilities. Specifically, the ranges across the five triplets of prizes are narrow, indicating a high adherence to EU. Additionally, these ranges encompass very high values for the probability of winning the high prize. For instance, in all the triplets of prizes in our experiment, the two estimated utilities both rank the middle prize above a binary risky lottery that offers the high prize with a probability of approximately 0.8, or the low prize with the complementary probability. Therefore, Group 1 reflects subjects in the experiment who systematically opted for the safer available lottery. The estimated proportion of subjects belonging to Group 1 is 0.2798.

Moving to Group 2 and Group 3, the utilities within these groups are more distinct compared to those in Group 1. The difference in utilities is most pronounced in Group 2, which is evident from the broader range of probabilities that they induce. Interestingly, utilities in both Group 2 and Group 3 are neither strictly concave nor convex. As a result, the range of probabilities in these groups spans values both below and above 0.5, which represents the EU threshold for risk aversion in our experiment. This suggests that subjects' behavior in these two groups does not strictly align with either pure risk aversion or risk seeking tendencies. However, the probability ranges in Group 2 consistently register higher values than those in Group 3, indicating relatively more risk averse behavior in comparative terms. The estimated proportion of subjects belonging to Group 2 and 3 are 0.4834 and 0.2368, respectively.

The emergence of non-EU behavior and risk averse tendencies consistently differs within the groups across the five triplets of prizes in the experiment. Notably, subjects consistently exhibit stronger risk aversion when the lowest prize in the triplet is \$0 compared to when it is a positive amount. On the whole, their behavior aligns more closely with EU. These observations are particularly evident in the probability ranges for the triplets $T1 = (\$0, \$15, \$30)$ and $T5 = (\$0, \$10, \$20)$, which are narrower and encompass higher values than other triplets. Among all the triplets, $T4 = (\$15, \$20, \$25)$ stands out as the one with the most pronounced non-EU behavior across all three groups. This triplet is unique in having prizes that are closely spaced, which could be linked to higher noise in responses.[22]

[22]A potential extension of our model might allow the variance of the error term to be influenced by features of the lotteries, such as the prize amounts.

Result 4. *In the mixture model for the EU core with three groups estimated using both CR-tasks and R-tasks, approximately 28% of subjects (Group 1) are extremely risk averse and align closely with EU. The remaining subjects (Group 2 and Group 3) deviate markedly from EU and cannot be categorically classified as either risk averse or risk seeking.*

1.5 Out-of-Sample Predictions: EU core Analysis

In this section, we evaluate whether, and under which conditions, the out-of-sample predictions derived from the analysis of the EU core are more precise than those obtained from specific economic models or machine learning algorithms. We examine various out-of-sample exercises, each involving different degrees of similarity between training and test sets. Initially, we reexamine the two out-of-sample exercises related to choices in CR-tasks and R-tasks, previously discussed in Section 1.3. In these exercises, training and test sets bear a close relationship. Subsequently, we use all choices from CR-tasks and R-tasks as training data and aim to predict the certainty equivalents derived from choices in Block 2 of the experiment. Lastly, we explore the correlation between estimated levels of risk aversion and adherence to EU with financial habits outside of the experiment.

Out-of-Sample Exercises Within and Across CR-Tasks and R-Tasks

Both economic models and machine learning algorithms can be used to predict the probability of all conceivable choice patterns in CR-tasks and R-tasks. On the other hand, the EU core approach only allows us to predict the probability of the possible values of the index $Core$. As a result, we cannot distinguish between non-EU behaviors within CR-tasks and R-tasks. Specifically, we cannot differentiate between CRE and RCRE in CR-tasks, and between AM and NM in R-tasks. To implement a fair comparison that takes into account the different levels of prediction detail attainable by various approaches, we also assess the ability of economic models and machine learning algorithms to predict the index $Core$. Hence, we deem an observed non-EU choice pattern as correctly predicted by these approaches as long as they assign the highest predicted probability to one non-EU behavior, even if the predicted non-EU behavior does not match the observed one.

Table 1.4 shows the percentages of choice patterns with a correctly classified index $Core$ for economic models, machine learning algorithms, and the EU core approach. Consistent with the findings in Section 1.3, GBT achieves the best results in out-of-sample exercises within CR-tasks and R-tasks. Our approach performs significantly

Table 1.4: Out-of-sample exercises: adherence to EU.

Exercise	**Prediction**	**Models**				
		EU	CPT	GBT	NN	EU core
Combined Within Tasks	Index *Core*	45.78% (5.25)	47.73% (1.06)	**64.40%** (1.50)	45.10% (3.93)	60.11% (1.97)
Combined Across Tasks	Index *Core*	44.66% (5.56)	44.94% (2.93)	41.59% (7.09)	42.23% (11.70)	**48.64%** (3.32)

Notes: Percentages of choice patterns with a correctly classified index *Core*. Standard deviations in parentheses. The first row ("Combined Within Tasks") presents the average percentages for the two cross-validation exercises within CR-tasks and R-tasks. For these exercises, we used the same partition of data used in Section 1.3 to assess the predictive accuracy of economic models and machine learning algorithms at predicting specific choice patterns. The second row ("Combined Across Tasks") displays the average percentages for the two out-of-sample exercises across CR-tasks and R-tasks described in Section 1.3.

better than both EU and CPT, while it is around 4% less accurate than GBT. Moreover, in line with the findings in Section 1.3, the ranking between economic models and machine learning algorithms is inverted once we switch from out-of-sample exercises within tasks to across tasks.[23] In this latter type of out-of-sample exercise, the EU core approach outperforms all other methods, yielding approximately 4% higher accuracy than economic models.

Result 5. *The EU core outperforms economic models in out-of-sample exercises within CR-tasks and R-tasks, though it is around 4% less accurate than GBT. Conversely, the EU core approach attains the most accurate results in out-of-sample exercises across CR-tasks and R-tasks, delivering predictions that are approximately 4% more accurate than those of economic models.*

Certainty Equivalents

We use choices from CR-tasks and R-tasks as training data and evaluate the accuracy of different approaches in predicting the certainty equivalents inferred from choices in Block 2 of our experiment. In Block 2, subjects are asked to compare three risky lotteries and various certain prizes. We focus on the subset of observations where subjects shifted their preference between a fixed lottery and a certain amount at most

[23]Here, CPT and NN do relatively better than EU and GBT, respectively. This is a result of the lenient approach we are using in evaluating these methods. For example, we categorize a choice pattern as correctly classified even when we observe a preference for randomization, while CPT predicts an aversion to randomization.

once.[24] For lotteries where subjects make a single switch, we compute the certainty equivalent as the smallest certain amount preferred over the lottery, reduced by \$0.5. If a lottery is chosen over all the certain prizes, its certainty equivalent is computed as the highest prize in the experiment compared to that lottery. Conversely, if the certain prize is always preferred, the certainty equivalent is computed as the smallest prize in the experiment compared to the lottery.

The EU core model enables us to predict a range of certainty equivalents. Specifically, given an estimated set of utilities $\hat{W}$, the certainty equivalent for a lottery p is predicted to lie within:

$$[\min_{v \in \hat{W}} c(p, v), \max_{v \in \hat{W}} c(p, v)],$$

where $c(p, v)$ represents the certainty equivalent of lottery p determined using utility function v.

The estimation results presented in Section 1.4 highlight significant heterogeneity in preferences. In particular, we estimated mixture models for three distinct groups of subjects and ranked these groups in terms of risk aversion and non-EU behavior. The greater the risk aversion, the higher the possible values for the risk premium associated with each lottery.[25] Additionally, increased non-EU behavior implies broader possible ranges of risk premia. Figure 1.6 summarizes with box plots the distribution of risk premia of all lotteries from Block 2 for the three groups of subjects. To construct this graph, we assigned each subject to the group with the highest group membership probability.

Comparing the distribution of risk premia across groups, we observe that the risk premia increase with the predicted level of risk aversion. For instance, the median risk premium in the group with low predicted levels of risk aversion is 0.5, while it is 5.5 in the group with high predicted levels of risk aversion. Furthermore, our model not only accurately predicts the differences in risk premia levels across groups but also the levels within each group. The dashed red lines in Figure 1.6 illustrate the predicted ranges of risk premia in the three groups. In all the groups, the predicted ranges of risk premia are perfectly consistent with the observed ones. In particular, for each of the three lotteries, the predicted range always includes the observed median risk premium.

---

[24] Appendix A.1 summarizes the distribution of risk premia for the three lotteries presented to subjects in Block 2.

[25] The risk premium of a lottery is the difference between the expected value and the certainty equivalent of a lottery.

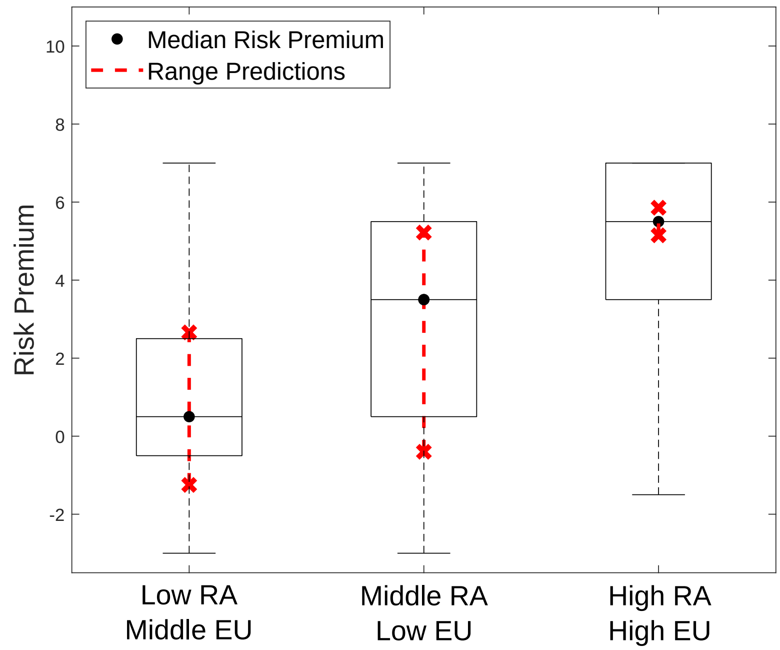

Figure 1.6: Observed risk premia and predicted ranges under the EU core model.

Notes: Box plots depict the distributions of risk premia in the three groups identified using mixture models with the EU core. Groups are classified based on RA and adherence to EU. Observations are considered outliers if they lie more than 1.5 interquartile ranges above the third quartile (75%) or below the first quartile (25%). We identified eleven outliers in the "High RA High EU" group. All outliers are omitted from the graph for clarity.

To benchmark the performance of our approach, we employ both economic models and machine learning algorithms to predict certainty equivalents. EU and CPT directly yield point predictions for the certainty equivalent of every lottery. Similarly, with GBT and NN, we can predict the specific point in a price list where subjects transition from favoring the certain amount to the risky lottery. We then deduce certainty equivalents using the same approach applied to derive certainty equivalents from the observed choices in Block 2.[26] To establish a fair comparison between the EU core approach and the other methods, we compute point predictions of the EU core approach for the certainty equivalents by taking the average values of the predicted ranges.

We use choices from CR-tasks and R-tasks as our training data and compare the mean squared errors of various approaches in predicting certainty equivalents. The results are presented in Table 1.5, where a lower mean squared error indicates superior

[26]In 5.25% of observations, GBT predicts either multiple switches between the certain amount and the risky lottery, or a single switch that is directionally incorrect. In this latter scenario, subjects are forecasted to choose the certain prize when its value is low and opt for the risky lottery when the certain prize value is high. We exclude these observations when evaluating the performance of GBT.

Table 1.5: Out-of-sample exercise across tasks in Block 1 and Block 2.

Exercise	Loss	Models				
		EU	CPT	GBT	NN	EU core
Train: CR-tasks and R-tasks Test: Certainty Equivalents in Block 2	MSE	13.184	16.964	64.148	57.398	**11.444**

Notes: We employ mixture models with three groups to predict risk premia using EU, CPT, and the EU core approach. For predictions with the EU core approach, we use the average risk premium within the predicted range.

model performance. In this out-of-sample evaluation, machine learning algorithms perform considerably worse than other methods. Specifically, GBT predicts that for over 90% of the lotteries, subjects either always prefer the certainty prize, or the risky lottery. Meanwhile, NN yields the analogous prediction for all lotteries. Therefore, both GBT and NN fail to offer reasonable predictions for choices in Block 2. EU and CPT significantly outperform machine learning algorithms in this task, with EU achieving a lower MSE than CPT. Finally, the EU core approach emerges as the top performer.

Result 6. *In the exercise of predicting risk premia from Block 2 using choices in Block 1 as training data, the EU core approach outperforms both economic models and machine learning algorithms. Within economic models, EU performs better than CPT, while machine learning algorithms exhibit the poorest performance in this out-of-sample exercise.*

Investment and Insurance Behaviors Outside the Experiment

In this section, we explore whether the heterogeneity identified in the experiments through the EU core approach, in terms of risk aversion and non-EU behavior, has any correlation with real investment and insurance behaviors. All subjects on Prolific are asked a series of questions about their financial habits when they first enroll on the platform. In the preregistration of the experiment, we chose to evaluate two conjectures.[27] The first conjecture concerns risk aversion, and the second pertains to subjects' adherence to EU.

We posit that individuals who are less risk averse should be more inclined to invest, especially in volatile assets. To assess this conjecture, we focused on two specific questions. The first inquires whether subjects have made investments (either personal or through their employment) in the common stock or shares of a company.

[27] See page 14 of the analysis plan preregistered at the AEA RCT Registry as AEARCTR-0011749 (Kobayashi and Lucia, 2023).

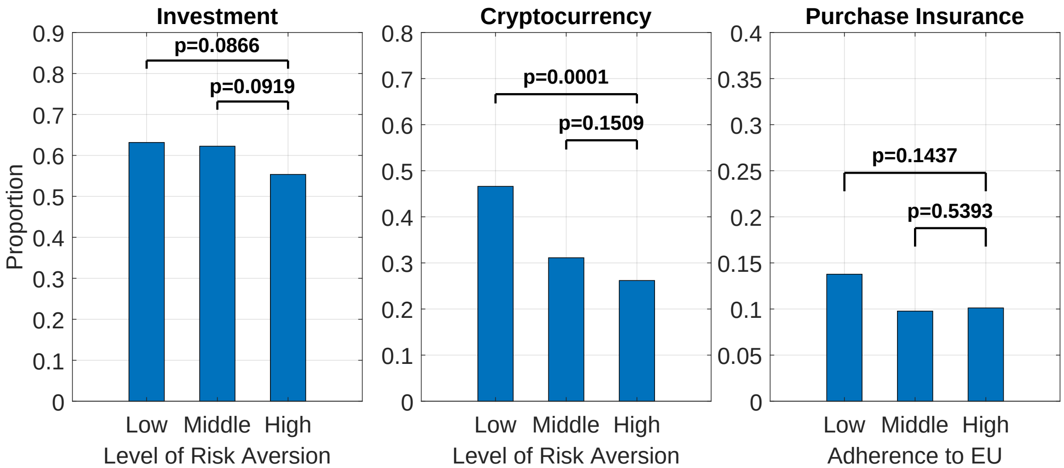

Figure 1.7: Investment and insurance behaviors: proportions.

Notes: We estimated a mixture model for the EU core with three groups, using CR-tasks and R-tasks. Each subject was assigned to the group with the highest group membership probability. The left graph displays proportions based on affirmative responses to the question, "Have you ever made investments (either personally or through your employment) in the common stock or shares of a company?" The middle graph represents proportions based on affirmative answers to the question, "Do you own/hold any cryptocurrencies?" The right graph, meanwhile, focuses on the question, "Do you actively hold any of the following types of insurance policies?" Specifically, it illustrates the proportion of subjects who selected the option "Purchase Insurance (e.g., Mobile Phone)."

Of course, the act of investing in company shares per se does not necessarily correlate with risk aversion. Much depends on the level of risk associated with the specific stocks considered, information we do not possess. Consequently, we also decided to explore whether subjects declared ownership of cryptocurrencies, serving as a more unambiguous proxy for risky behavior. Moreover, to capture non-EU behavior, we examined whether subjects have purchased insurance for items, such as smartphones. This behavior is a classic example of small stakes risk aversion, which can be challenging to rationalize with EU. Therefore, our second conjecture is that subjects who align more closely with EU should be less inclined to purchase this type of insurance.

Figure 1.7 summarizes the responses to the three questions under consideration, across the different groups of subjects identified with the mixture model for the EU core estimated using all data from Block 1. For investments in shares of companies (left graph) and cryptocurrencies (middle graph), we order groups from left to right based on their estimated level of risk aversion. For purchasing insurance (right graph), we order groups based on their estimated adherence to EU.

The left and middle graphs of Figure 1.7 provide evidence in line with our first conjecture: groups of subjects characterized by lower estimated levels of risk aversion have higher proportions of individuals investing in shares of companies or cryptocurrencies. Moreover, the negative correlation between risk aversion and investment behavior is more pronounced for cryptocurrencies than for company shares. This observation supports the idea that investments in cryptocurrencies might be a better proxy for risky behavior. Consistent with our second conjecture, the proportion of subjects purchasing insurance is highest among groups with low predicted adherence to EU.

Result 7. *We classified subjects in terms of risk aversion and adherence to EU using the mixture model for the EU core estimated from choices in CR-tasks and R-tasks. Subjects classified as more risk averse are less likely to invest, particularly in cryptocurrencies. Moreover, subjects classified as less adherent to EU purchase insurance more frequently.*

1.6 Discussion

Our paper offers two main contributions. Empirically, we demonstrate the shortcomings of popular economic models and machine learning algorithms in both rationalizing and predicting two widely documented and influential non-EU behaviors, which were previously mostly analyzed separately in experimental work. Methodologically, we introduce a novel empirical strategy for predicting behavior under risk and showcase its effectiveness through a series of out-of-sample exercises. We conclude by discussing the implications of our results for future research.

A satisfactory model of decision-making under risk should rationalize the strong positive correlation observed between the CRE and preferences for randomization. Conversely, CPT predicts the opposite correlation between these two non-EU behaviors, explaining why this model consistently achieves inferior out-of-sample performance compared to EU in our analysis. A natural direction for future research involves considering alternative models to CPT. Fudenberg, Kleinberg, et al. (2022) demonstrate that adding to CPT a complexity cost, which increases with the number of prizes in a lottery, yields better out-of-sample predictions. In our experiment, adding a complexity cost to CPT would further strengthen the negative relationship between CRE and preferences for randomization, leading to worse out-of-sample predictions.

Of course, many alternative theories to CPT and its generalizations have been proposed, and some of them may be capable of rationalizing our main experimental findings. One important example is the original Prospect Theory (PT; Kahneman and Tversky, 1979). This model drops the rank dependence assumption that characterizes CPT and can rationalize the observed correlation between CRE and preferences for randomization through simple probability weighting. However, there is also abundant experimental evidence demonstrating failures of PT, with CPT being proposed as a solution to the violations of first-order stochastic dominance implied by PT. The overall absence of an economic model that systematically outperforms the others in terms of predictive accuracy generates interest in exploring alternative approaches for making predictions.

Machine learning algorithms offer an alternative approach for predicting behavior under risk, and a growing body of research compares their predictive capabilities with those of economic models.[28] In our analysis, machine learning algorithms outperformed economic models only when the training and test sets included choices over the exact same lotteries. However, their predictive capabilities dropped significantly in all other out-of-sample exercises. Andrews et al. (2022) obtain a similar result when comparing the out-of-sample performance of economic models and machine learning algorithms in the prediction of certainty equivalents. In particular, they observe that the performance of machine learning algorithms are sensitive to which lotteries are included in the training and test sets. Therefore, the sensitivity of these methods to minor differences between training and test sets raises concerns about their ability to produce generalizable predictions that are at least as substantial as those for economic models.

This paper introduces a novel empirical approach to make predictions that retains an underlying economic structure without being tied to specific models. We show that the predictions of our approach are more accurate but at same time less detailed than those produced by fully specified economic models or machine learning algorithms. For instance, our approach does not allow distinguishing between specific non-EU behaviors, or it does not allow one to directly obtain a point prediction for a certainty equivalent. While this paper focuses on choices under risk, we believe our empirical strategy holds promise for extension to choices under uncertainty, exploiting the notion of "unambiguous preference" introduced by Ghirardato, Maccheroni, and Marinacci (2004).

[28] See Andrews et al. (2022), Camerer, Nave, and Smith (2019), Fudenberg and Liang (2019), Noti et al. (2016), Joshua C. Peterson et al. (2021), Plonsky, Erev, et al. (2017), Plonsky, Apel, et al. (2019), and Zhao et al. (2020).

Ultimately, determining the "best" method to predict behavior under risk may depend on at least two factors. If the analyst has access to a training set that is sufficiently close to the test set of interest, machine learning algorithms might be the most suitable alternative. However, what constitutes "sufficiently close" can vary based on the specific nature of the problem at hand. Our analysis underscores how minor discrepancies between training and test sets can lead to substantial declines in the performance of machine learning algorithms. Another critical aspect to consider is the level of detail required in the predictions. For instance, if the analyst aims to estimate measures of risk aversion or adherence to EU, our approach offers a promising alternative to traditional predictive tools. Conversely, if our method does not provide the necessary level of detail in predictions, it can still complement other techniques. For example, if a point prediction for a certainty equivalent is needed, our strategy can first offer a range prediction, which economic models or machine learning algorithms can then refine to pinpoint a value within that range.

References

Agranov, Marina, Paul J Healy, and Kirby Nielsen (June 2023). "Stable Randomisation". In: *The Economic Journal* 133.655, pp. 2553–2579. ISSN: 0013-0133. DOI: 10.1093/ej/uead039. eprint: https://academic.oup.com/ej/article-pdf/133/655/2553/51707894/uead039.pdf. URL: https://doi.org/10.1093/ej/uead039.

Agranov, Marina and Pietro Ortoleva (2017). "Stochastic Choice and Preferences for Randomization". In: *Journal of Political Economy* 125.1, pp. 40–68. ISSN: 00223808, 1537534X. URL: https://www.jstor.org/stable/26549925 (visited on 10/16/2023).

– (May 2022). "Revealed Preferences for Randomization: An Overview". In: *AEA Papers and Proceedings* 112, pp. 426–30. DOI: 10.1257/pandp.20221093. URL: https://www.aeaweb.org/articles?id=10.1257/pandp.20221093.

– (July 2023). "Ranges of Randomization". In: *The Review of Economics and Statistics*, pp. 1–44. ISSN: 0034-6535. DOI: 10.1162/rest_a_01355. eprint: https://direct.mit.edu/rest/article-pdf/doi/10.1162/rest_a_01355/2150363/rest_a_01355.pdf. URL: https://doi.org/10.1162/rest%5C_a%5C_01355.

Andrews, Isaiah et al. (2022). "The Transfer Performance of Economic Models". In: *arXiv preprint arXiv:2202.04796*.

Athey, Susan (2017). "Beyond prediction: Using big data for policy problems". In: *Science* 355.6324, pp. 483–485. DOI: 10.1126/science.aal4321. eprint:

https://www.science.org/doi/pdf/10.1126/science.aal4321. URL: https://www.science.org/doi/abs/10.1126/science.aal4321.

Bruhin, Adrian, Helga Fehr-Duda, and Thomas Epper (2010). "Risk and Rationality: Uncovering Heterogeneity in Probability Distortion". In: *Econometrica* 78.4, pp. 1375–1412. ISSN: 1468-0262. DOI: 10.3982/ECTA7139. (Visited on 04/08/2022).

Camerer, Colin F and Teck-Hua Ho (1994). "Violations of the betweenness axiom and nonlinearity in probability". In: *Journal of risk and uncertainty* 8.2, pp. 167–196.

Camerer, Colin F, Gideon Nave, and Alec Smith (2019). "Dynamic unstructured bargaining with private information: theory, experiment, and outcome prediction via machine learning". In: *Management Science* 65.4, pp. 1867–1890.

Cerreia-Vioglio, Simone (2009). *Maxmin expected utility on a subjective state space: Convex preferences under risk*. Tech. rep. Mimeo, Bocconi University.

Cerreia-Vioglio, Simone, David Dillenberger, and Pietro Ortoleva (2015). "Cautious expected utility and the certainty effect". In: *Econometrica* 83.2, pp. 693–728.

Cerreia-Vioglio, Simone, Fabio Maccheroni, and Massimo Marinacci (2017). "Stochastic dominance analysis without the independence axiom". In: *Management Science* 63.4, pp. 1097–1109.

Cettolin, Elena and Arno Riedl (2019). "Revealed preferences under uncertainty: Incomplete preferences and preferences for randomization". In: *Journal of Economic Theory* 181, pp. 547–585.

Chapman, Jonathan et al. (2023a). "Econographics". In: *Journal of Political Economy Microeconomics* 1.1, pp. 115–161. DOI: 10.1086/723044. eprint: https://doi.org/10.1086/723044. URL: https://doi.org/10.1086/723044.

– (2023b). *Willingness to accept, willingness to pay, and loss aversion*. Tech. rep. National Bureau of Economic Research.

Chew, S. H., L. G. Epstein, and U. Segal (1991). "Mixture Symmetry and Quadratic Utility". In: *Econometrica* 59.1, pp. 139–163. ISSN: 00129682, 14680262. URL: http://www.jstor.org/stable/2938244 (visited on 02/20/2024).

Coveney, Peter V., Edward R. Dougherty, and Roger R. Highfield (2016). "Big data need big theory too". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 374.2080, p. 20160153. DOI: 10.1098/rsta.2016.0153. eprint: https://royalsocietypublishing.org/doi/pdf/10.1098/rsta.2016.0153. URL: https://royalsocietypublishing.org/doi/abs/10.1098/rsta.2016.0153.

Dean, Mark and Pietro Ortoleva (2015). "Is it All Connected? A Testing Ground for Unified Theories of Behavioral Economics Phenomena". In: *ERN: Uncertainty & Risk Modeling (Topic)*. URL: https://api.semanticscholar.org/CorpusID:43571737.

Dean, Mark and Pietro Ortoleva (2019). “The empirical relationship between nonstandard economic behaviors”. In: *Proceedings of the National Academy of Sciences* 116.33, pp. 16262–16267. DOI: 10.1073/pnas.1821353116. eprint: https://www.pnas.org/doi/pdf/10.1073/pnas.1821353116. URL: https://www.pnas.org/doi/abs/10.1073/pnas.1821353116.

Dwenger, Nadja, Dorothea Kübler, and Georg Weizsäcker (2018). “Flipping a coin: Evidence from university applications”. In: *Journal of Public Economics* 167, pp. 240–250. ISSN: 0047-2727. DOI: https://doi.org/10.1016/j.jpubeco.2018.09.014. URL: https://www.sciencedirect.com/science/article/pii/S0047272718301889.

Feldman, Paul and John Rehbeck (2022). “Revealing a preference for mixtures: An experimental study of risk”. In: *Quantitative Economics* 13.2, pp. 761–786.

Freeman, David J, Yoram Halevy, and Terri Kneeland (2019). “Eliciting risk preferences using choice lists”. In: *Quantitative Economics* 10.1, pp. 217–237.

Fudenberg, Drew, Jon Kleinberg, et al. (2022). “Measuring the completeness of economic models”. In: *Journal of Political Economy* 130.4, pp. 956–990.

Fudenberg, Drew and Annie Liang (2019). “Predicting and understanding initial play”. In: *American Economic Review* 109.12, pp. 4112–4141.

Ghirardato, Paolo, Fabio Maccheroni, and Massimo Marinacci (2004). “Differentiating ambiguity and ambiguity attitude”. In: *Journal of Economic Theory* 118.2, pp. 133–173. ISSN: 0022-0531. DOI: https://doi.org/10.1016/j.jet.2003.12.004. URL: https://www.sciencedirect.com/science/article/pii/S0022053104000262.

Gul, Faruk (1991). “A theory of disappointment aversion”. In: *Econometrica: Journal of the Econometric Society*, pp. 667–686.

Hey, John D and Chris Orme (1994). “Investigating generalizations of expected utility theory using experimental data”. In: *Econometrica: Journal of the Econometric Society*, pp. 1291–1326.

Hofman, Jake M et al. (2021). “Integrating explanation and prediction in computational social science”. In: *Nature* 595.7866, pp. 181–188.

Kahneman, Daniel and Amos Tversky (1979). “Prospect Theory: An Analysis of Decision under Risk”. In: *Econometrica* 47.2, pp. 263–291. ISSN: 00129682, 14680262. URL: http://www.jstor.org/stable/1914185 (visited on 10/15/2023).

Loomes, Graham and Robert Sugden (1982). “Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty”. In: *The Economic Journal* 92.368, pp. 805–824. ISSN: 00130133, 14680297. URL: http://www.jstor.org/stable/2232669 (visited on 02/20/2024).

Machina, Mark J. (1982). ""Expected Utility" analysis without the independence axiom". In: *Econometrica* 50.2, pp. 277–323. ISSN: 00129682, 14680262. URL: http://www.jstor.org/stable/1912631 (visited on 02/21/2024).

Marschak, Jacob (1950). "Rational behavior, uncertain prospects, and measurable utility". In: *Econometrica: Journal of the Econometric Society*, pp. 111–141.

Miao, Bin and Songfa Zhong (2018). "Probabilistic social preference: how Machina's Mom randomizes her choice". In: *Economic Theory* 65, pp. 1–24.

Noti, Gali et al. (2016). "Behavior-Based Machine-Learning: A Hybrid Approach for Predicting Human Decision Making". In: *ArXiv* abs/1611.10228. URL: https://api.semanticscholar.org/CorpusID:14606035.

Peterson, Joshua C. et al. (2021). "Using large-scale experiments and machine learning to discover theories of human decision-making". In: *Science* 372.6547, pp. 1209–1214. DOI: 10.1126/science.abe2629. eprint: https://www.science.org/doi/pdf/10.1126/science.abe2629. URL: https://www.science.org/doi/abs/10.1126/science.abe2629.

Plonsky, Ori, Reut Apel, et al. (2019). "Predicting human decisions with behavioral theories and machine learning". In: *arXiv preprint arXiv:1904.06866*.

Plonsky, Ori, Ido Erev, et al. (2017). "Psychological forest: Predicting human behavior". In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 31. 1.

Rabin, Matthew (2000). "Risk Aversion and Expected-Utility Theory: A Calibration Theorem". In: *Econometrica* 68.5, pp. 1281–1292. ISSN: 00129682, 14680262. URL: http://www.jstor.org/stable/2999450 (visited on 10/29/2023).

Rothschild, Michael and Joseph E Stiglitz (1970). "Increasing risk: I. A definition". In: *Journal of Economic Theory* 2.3, pp. 225–243. ISSN: 0022-0531. DOI: https://doi.org/10.1016/0022-0531(70)90038-4. URL: https://www.sciencedirect.com/science/article/pii/0022053170900384.

Sandroni, Alvaro, Sandra Ludwig, and Philipp Kircher (2013). "On the difference between social and private goods". In: *The BE Journal of Theoretical Economics* 13.1, pp. 151–177.

Segal, Uzi (1990). "Two-stage lotteries without the reduction axiom". In: *Econometrica: Journal of the Econometric Society*, pp. 349–377.

Sopher and Narramore (2000). "Stochastic Choice and Consistency in Decision Making Under Risk: An Experimental Study". In: *Theory and Decision* 48.4, pp. 323–349. DOI: 10.1023/a:1005289611789.

Tversky, Amos and Daniel Kahneman (1992). "Advances in prospect theory: Cumulative representation of uncertainty". In: *Journal of Risk and uncertainty* 5.4, pp. 297–323.

Zhao, Chen et al. (2020). “Behavioral neural networks”. In: *Available at SSRN 3633548*.

Chapter 2

WHAT DRIVES VIOLATIONS OF THE INDEPENDENCE AXIOM? THE ROLE OF DECISION CONFIDENCE

2.1 Introduction

A central research question in economics is how individuals make decisions in the presence of uncertainty. Research has demonstrated critical shortcomings in the neoclassical formulation of EU and its central assumption, the independence axiom. In the most famous counterexample of EU proposed by Maurice Allais in 1952, individuals typically violate the independence axiom by showing risk aversion when one certain option is available and risk tolerance when all the available options are uncertain.[1] This tendency, known as the "certainty effect" (Kahneman and Tversky, 1979), is a centerpiece of theoretical alternatives to EU—most notably CPT introduced by Tversky and Kahneman (1992).

Recent research has challenged the empirical regularities related to the certainty effect and the theoretical explanations proposed to rationalize them. In particular, P. Blavatskyy, Ortmann, and Panchenko (2022, 2022) and Jain and Nielsen (2022) document that features of the experimental design can affect the likelihood of the certainty effect. Moreover, Bernheim and Sprenger (2020) find no evidence of the rank dependence assumption on which CPT relies. Motivated by these findings, I design an experiment to study the relevance of certainty and the certainty effect and to investigate an alternative mechanism for violations of EU's independence axiom: lack of confidence when choosing between different lotteries.

This paper provides the first experimental investigation of the independence axiom through the lenses of the EU core, which captures the greatest part of a preference relation that satisfies the independence axiom (Cerreia-Vioglio, 2009). When a decision-maker violates the independence axiom, his EU core is a partial order that is typically interpreted as the subset of his uncontroversial rankings. According to this interpretation, we should observe violations of the independence axiom only in choice problems that are "hard" enough so that the decision-maker feels unconfident about them. Recent non-EU models further appeal to the lack of decision confidence

[1] Independence states that for any three lotteries p, q, and r, and any number $\lambda \in (0,1)$, if p is preferred to q, then $\lambda p + (1-\lambda)r$ is preferred to $\lambda q + (1-\lambda)r$.

as the driving force for violations of the independence axiom. However, little is known empirically about the relationship between decision confidence and violations of the independence axiom that our experiment seeks to explore.

In our experiment, we ask subjects to make incentivized choices between lotteries. After making each choice, subjects report on a scale from zero to 100 how confident they are about their choice. We also collect response times as an additional proxy for confidence. To construct the pairs of lotteries in the experiment, we start with an initial set of *unmixed comparisons*, each consisting of a certain prize and a risky lottery. Next, to test the independence axiom, we create *mixed comparisons* by mixing each of the lotteries in an unmixed comparison with a series of third common lotteries. Evaluating independence by assessing behavior in multiple mixed comparisons allows us to assess the independence axiom—and hence adherence to the EU core—for each choice.

We conduct our experiment on the online platform Prolific.co with 300 subjects. Each subject made binary choices over lotteries in 74 comparisons. We assess adherence to the independence axiom for each choice and relate that adherence to measures of decision confidence. A central result of this paper is that behavior is more likely to comply with the independence axiom when subjects report high confidence and make their decisions quickly. In contrast, behavior systematically deviates from EU when subjects report low confidence and need more time to make their choices.

Our data also show precisely how individuals choose when they are not confident (i.e., when the comparison is less likely to belong to the EU core). Dillenberger (2010) introduces the negative certainty independence (NCI) axiom to rationalize the certainty effect. The NCI axiom can be interpreted as a rule that prescribes what to choose when not confident. In particular, whenever an individual consistent with NCI is not confident that a risky lottery is better than a certain prize, he should prefer the certain prize. Somewhat surprisingly, we find evidence against NCI and its interpretation: risky lotteries are more likely to be chosen in situations of low confidence and to induce independence violations.

Finally, we study whether the mere presence of a certain prize in a decision makes independence violations more likely. Recent works find evidence consistent with the idea that people value certain and uncertain outcomes differently (Halevy, 2008; Andreoni and Harbaugh, 2009; Andreoni and Sprenger, 2010, 2011, 2012). In this paper we first show that a naive analysis may lead to overestimating the role of certainty; there is more data to detect EU violations in choices with a certain alternative than in choices where all alternatives are risky. In the benchmark scenario of an EU decision-maker who makes mistakes, more data to test independence translates into higher expected independence violations. We then adopt the error model of Harless and Camerer (1994) to control for this asymmetry and find that the presence of a certain prize is not always associated with more independence violations.

The first main implication of our results is that the analysis of decision confidence in choices under risk constitutes a promising avenue for deepening our understanding of why subjects deviate from EU in a specific way. Our analysis relates to the recent experimental literature on preferences for randomization. Agranov and Ortoleva (2017) show that individuals strictly prefer to randomize in “hard" comparisons and consequently violate the independence axiom. Arts, Ong, and Qiu (2022) design an experiment to elicit both decision confidence and randomization probabilities in choices under risk. In line with the interpretation of the hard questions in Agranov and Ortoleva (2017), they show that subjects tend to choose randomization probabilities that are close to uniform when they report low confidence measures. In our experiment, we obtain the same relationship between decision confidence and violations of the independence axiom in common ratio questions.

More broadly, this paper relates to the literature in psychology, neuroscience, and to less extent, economics, documenting how behavioral anomalies are often associated with low levels of decision confidence.[2] Enke and Graeber (2019) show that cognitive uncertainty—measured as subjective uncertainty over the ex-ante utility-maximizing decision—is associated with an attenuated relationship between decisions and problem parameters. This attenuated relationship may lead to the emergence of well-known behavioral patterns, such as the fourfold pattern of risk attitudes.

We contribute to this literature by focusing on the choice implications of low decision confidence in the risk domain. Within this domain, we show that subjects tend to adopt a form of “incaution" when reporting low confidence levels by choosing the riskiest available lottery. This finding generates interest in decision models that relies on the positive certainty independence (PCI) axiom (Cerreia-Vioglio, Dillenberger, and Ortoleva, 2020). This axiom rationalizes the behavioral pattern opposite to

[2]We refer to Arts, Ong, and Qiu (2022) for a review of this literature.

the certainty effect and predicts the relationship between decision confidence and preference for risky alternatives we observe in the data.

Moreover, our findings highlight that features of the experimental design may drastically affect the overall amount and, most importantly, the conclusions about what drives violations of the independence axiom. Most experimental works on the independence axiom are variations of the Allais paradox, in which we know that the certainty effect constitutes the modal behavioral pattern.[3] However, our results show that the predominance of the certainty effect may not be robust to richer environments where subjects face a wider variety of unmixed and mixed comparisons.[4]

Another stylized fact in choices under risk that we put under scrutiny is that violations of the independence axiom are less frequent in comparisons with nondegenerate lotteries over common prizes. Camerer (1992) writes that "much as Newtonian mechanics is an adequate working theory at low velocities, EU seems to be an adequate working theory for gambles inside the triangle".[5] Our analysis challenges the generality of this conclusion showing that its validity may depend on the mechanism through which mixed comparisons are constructed.

The paper proceeds as follows. Section 2.2 describes the theoretical framework in the context of our experimental design. Section 2.3 illustrates our experimental design. Section 2.4 presents our main findings and Section 2.5 concludes.

2.2 Theoretical Framework

We describe the theoretical framework in the context of our experimental design. All questions in the experiment involve lotteries over the set of monetary prizes $X = \{\$1, \$7, \$20\}$. We denote the set of lotteries with prizes in X by $\triangle(X)$. We refer to generic prizes in X by x and denote generic lotteries in $\triangle(X)$ by p, q, r and s. We represent the three-outcome lottery, q, giving \$1 with probability $q(1)$, \$7 with probability $q(7)$ and \$20 with probability $q(20)$ as $(\$1, q\,(1)\,; \$7, q\,(7)\,; \$20, q\,(20))$. We write the lottery that gives $\$x$ for sure as δ_x and we refer to generic pairs of lotteries $(s, r) \in \triangle(X)^2$ as comparisons. Moreover, we denote by N be the set of all the subjects in the experiment, and by $\succsim_i$ and $\succ_i$ the weak and strict preference relations of a subject $i \in N$ over $\triangle(X)$.

The preference $\succsim_i$ satisfies the independence axiom if for all lotteries $q, s, r \in \triangle(X)$

[3] We refer to P. R. Blavatskyy (2010) for a review of this literature.

[4] Jain and Nielsen (2022) also find that the certainty effect is not the most common behavioral pattern that violates the independence axiom.

[5] See also Starmer (2000).

and for all $\lambda \in (0, 1]$,

$$s \succsim_i r \Rightarrow \lambda s + (1 - \lambda)q \succsim_i \lambda r + (1 - \lambda)q.$$

The EU core of $\succsim_i$ is the subrelation $\succsim_i^*$ such that for all lotteries $q, s, r \in \triangle(X)$ and for all $\lambda \in (0, 1]$,[6]

$$s \succsim_i^* r \Leftrightarrow \lambda s + (1 - \lambda)q \succsim_i \lambda r + (1 - \lambda)q.$$

That is, $s \succsim_i^* r$ whenever subject i prefers s to r and mixing both lotteries s and r with a third common lottery q does not affect the relative preferences of i between s and r.[7] Cerreia-Vioglio (2009) proves that $\succsim_i^*$ is the greatest subrelation of $\succsim_i$ that satisfies the independence axiom.[8] Therefore, to study how the independence axiom fails, we test for each subject i separately or for all subjects $i \in N$ at the aggregate level the following hypothesis:

$$s \succsim_i^* r \text{ or } r \succsim_i^* s. \tag{EU-CORE}$$

Throughout the paper, we say that hypothesis EU-CORE holds if we find no evidence against it, while we say that it fails otherwise. By correlating the results of hypothesis EU-CORE with the measures of decision confidence that we collect, we test the interpretation of $s \succsim_i^* r$ as individual i being confident that lottery s is better than lottery r. Moreover, we examine whether individuals are more likely to choose the safer or the riskier lottery when hypothesis EU-CORE does not hold and when they declare to be unconfident. If either lottery s or lottery r is degenerate, this analysis will allow us to shed light on the relevance of the certainty effect. Finally, we study whether hypothesis EU-CORE is more likely to hold in comparisons where the two lotteries are risky and have the same support.

2.3 Experimental Design

The rationale behind the experimental design is to create a rich dataset to study for what comparisons individuals are more likely to violate hypothesis EU-CORE and test whether the lack of decision confidence can explain failures of hypothesis EU-CORE. This section first illustrates the comparisons that we consider in

[6] $\succsim_i^*$ is a subrelation of $\succsim_i$ if for all lotteries s and r, $s \succsim_i^* r$ implies $s \succsim_i r$.

[7] In our experiment, we study the EU core by considering only "one-stage" lottery mixtures, rather than two-stage compound lotteries. In other words, we focus on mixture independence, rather than compound independence, as defined in Segal (1990).

[8] That is, if $\succsim_i^{**}$ is another subrelation of $\succsim_i$ that satisfies the independence axiom, then $\succsim_i^{**}$ is a subrelation of $\succsim_i^*$.

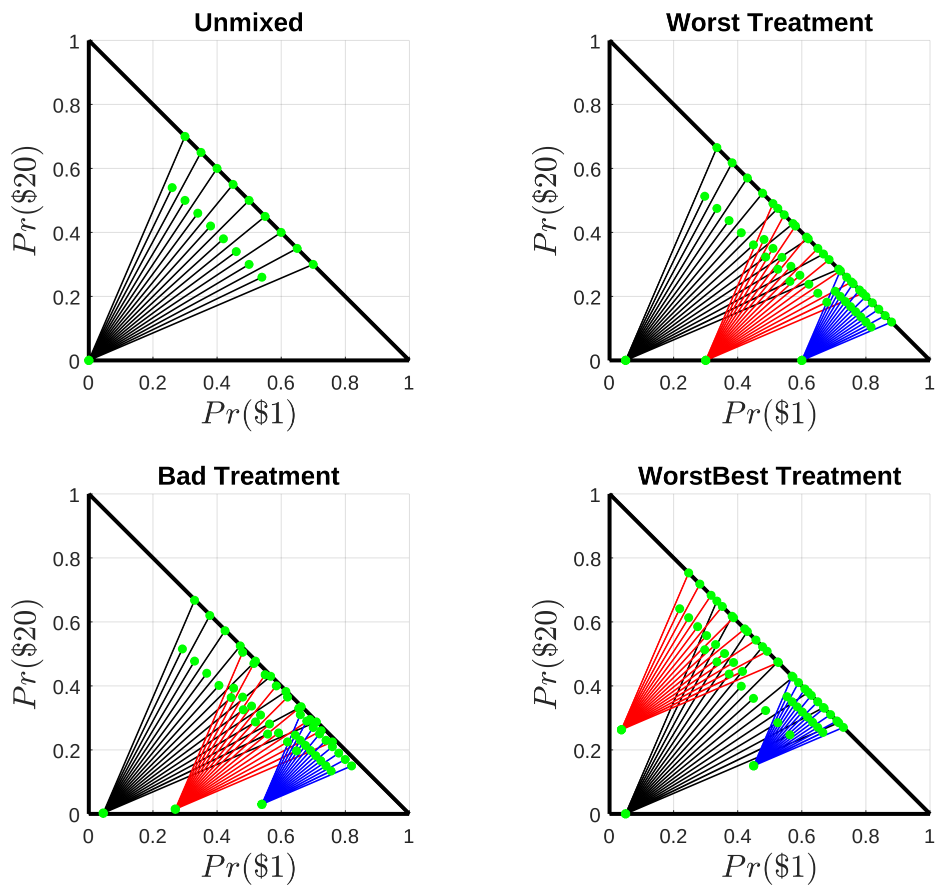

Figure 2.1: Unmixed and mixed comparisons in the three treatments.

the experiment. Next, we describe the questions that subjects answer about each comparison. Finally, we discuss the recruitment procedures and the experimental payments.[9]

Comparisons

There are three treatments: “Worst", “Bad" and “WorstBest". In all treatments, subjects face the same seventeen comparisons that we call *unmixed*, each involving the degenerate lottery δ_7 and a risky lottery. The treatments differ in how we construct the additional comparisons to test the independence axiom. In what follows, we use the MM triangle to describe the lotteries in the experiment (Marschak, 1950; Mark J. Machina, 1982). The top-left graph in Figure 2.1 shows the unmixed comparisons in the MM triangle. In all the graphs of Figure 2.1, the probability of receiving

[9] We preregistered the experimental design and the analysis plan at the AEA RCT Registry website on February 21, 2022. Link to the preregistration: https://www.socialscienceregistry.org/trials/8615.

\$20 is on the vertical axis, and the probability of receiving \$1 is on the horizontal axis. Therefore, the generic point (x, y) in the MM triangle represents the lottery $(\$1, x; \$7, 1-x-y; \$20, y)$. Each segment connecting the degenerate lottery δ_7 with a risky lottery represents an unmixed comparison between these two lotteries.

In order to test hypothesis EU-CORE for an unmixed comparison (δ_7, r), we need at least another comparison (p, q) of the following form:

$$p = \lambda\delta_7 + (1-\lambda)z \text{ and } q = \lambda r + (1-\lambda)z,$$

where both lotteries p and q are constructed by mixing lotteries δ_7 and r with a third common lottery z, using a fixed probability weight $\lambda \in (0, 1)$. We call comparisons that satisfy this property *mixed*. In particular, we omit the dependence on the third common lottery and refer to $(\lambda\delta_7, \lambda r)$ as a λ*-mixed comparison*.

Each treatment includes an equal number of 0.95-mixed, 0.7-mixed, and 0.4-mixed comparisons. Overall, there are 51 mixed comparisons in each treatment. The 0.95-mixed comparisons involve one almost degenerate lottery, i.e., $0.95\delta_7 + 0.05z$, while lotteries in the remaining mixed comparisons are all “far" from being degenerate. Therefore, we can study whether the presence of a degenerate or almost degenerate lottery is the main driver for the violations of hypothesis EU-CORE.

In the Worst treatment, we build λ-mixed comparisons by mixing the lotteries in each of the seventeen unmixed comparisons with the worst lottery $(\$1, 1)$ using the probability weight $\lambda \in \{0.95, 0.7, 0.4\}$. To construct mixed comparisons in the Bad treatment, we repeat the same procedure except for replacing lottery $(\$1, 1)$ with lottery $(\$1, 0.9; \$7, 0.05; \$20, 0.05)$, which is inside the MM triangle. The bottom-left graph of Figure 2.1 represents the mixed comparisons in the Bad treatment. Unlike the Worst treatment, mixed comparisons in the Bad treatment have lotteries with the same support. Therefore, we can study the relevance of this feature by comparing failures of hypothesis EU-CORE in the Worst and the Bad treatments.

Mixed comparisons in the Worst and the Bad treatments cluster in the southeast region of the MM triangle. This concentration may preclude us from detecting violations of hypothesis EU-CORE.[10] To account for this potential concern, we consider an additional treatment that we call WorstBest. The WorstBest treatment shares the same 0.95-mixed comparisons of the Worst treatment. The 0.7-mixed comparisons are constructed by mixing the lotteries in each of the 0.95-mixed

[10]For instance, this is the case if subjects' preferences are consistent with the fanning-out hypothesis (Mark J. Machina, 1982, 1987).

Table 2.1: Summary information of the three treatments.

	Worst	Bad	WorstBest
# Unmixed Comparisons	17	17	17
# Mixed Comparisons	51	51	51
# Dominance Comparisons	6	6	6
Third Common Lottery	Fixed	Fixed	Alternate
Probability Weights	0.95,0.7,0.4	0.95,0.7,0.4	0.95,0.7,0.4
Sample Size	100	100	100

comparisons with lottery $(\$20, 1)$ using $0.7/0.95$ as probability weight. Finally, for 0.4-mixed comparisons, we mix the lotteries in each of the 0.7-mixed comparisons with lottery $(\$1, 1)$ using $0.4/0.7$ as probability weight.[11] The bottom-right graph of Figure 2.1 describes the mixed comparisons in the WorstBest treatment.

Table 2.1 summarizes our experimental design. The large number and the diversity of the comparisons in the experiment enable us to test hypothesis EU-CORE throughout the MM triangle and ensure that systematic and persistent violations of hypothesis EU-CORE are not just a reflection of indifference.[12] To further evaluate the reliability of our data, we also include in each treatment six comparisons involving stochastically dominated lotteries.[13] When presenting our results, we exclude all subjects that chose the stochastically dominated lottery more than once.[14]

Questions

We first ask subjects to indicate the lottery they prefer for each comparison. Next, we ask them to report their confidence level on a scale from zero (not confident at all) to 100 (completely confident). We also collect response times in these answers as an indirect measure of decision confidence.[15] Subjects answer one question at a time. Once subjects select an answer, they can not modify it. Figure 2.2 shows a decision screen from the experiment. In this example, a subject who declared to prefer lottery ticket A over lottery ticket B is asked to report how confident he feels about this choice. The slider always starts at 50. In order to proceed to the next

[11] This approach ensures that differences in expected values between lotteries in unmixed and mixed comparisons are constant across the three treatments. That is, in all treatments, λ-mixed comparisons can be created by mixing the lotteries in unmixed comparisons with some third common lottery z using λ as probability weight.

[12] The “indifference" argument is a common critique for experiments that document preference reversals (P. R. Blavatskyy, 2010).

[13] We report these six comparisons in Appendix B.1.

[14] Overall, 32 out of 300 subjects chose the stochastically dominated lottery more than once.

[15] We report our analysis on response times in Appendix B.5.

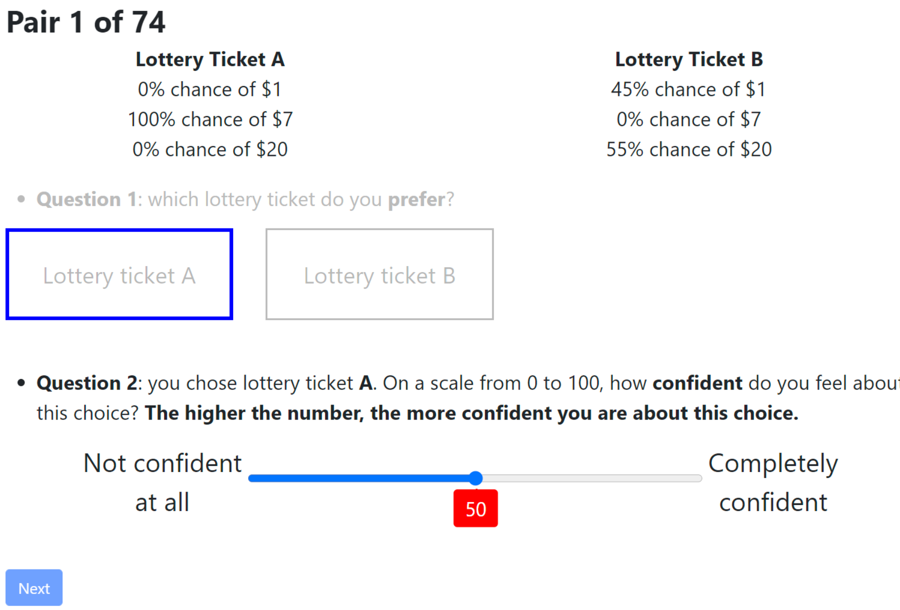

Figure 2.2: Decision screen from the experiment.

question, subjects need to click on the slider at least once.

Recruitment and Experimental Payments

We recruit 300 subjects through the online platform Prolific.co to run the experiment. A total of nine sessions were conducted between February 22 and February 26 of 2022. Overall, we recruited 100 subjects for each treatment. Our sample consists of United States citizens between the ages of 18 and 30 with at least a high-school education. We focus on this sample because most previous experiments involving common ratio questions have been conducted on undergraduate samples. Moreover, given that there are three times more women than men within the population of possible participants that meet these criteria, we ask Prolific to recruit an equal number of men and women.[16]

Each subject receives a fixed payment of $4.75 and has a one out of ten chance of receiving a bonus payment. The software randomly selects one of the 74 comparisons for subjects that receive a bonus payment. The bonus payment consists of the prize that subjects win by playing the lottery they declared to prefer during the experiment in the selected comparisons.[17]

[16] Table B.2 in Appendix B.2 summarizes the demographic information of the participants.

[17] The complete instructions with screenshots from the experiment are presented in Appendix B.6.

Table 2.2: Percentage of comparisons consistent with hypothesis EU-CORE at the individual level. Standard errors in parenthesis.

	Worst	Bad	WorstBest
% hp. EU-CORE holds	57.03%	57.58%	52.08%
	(0.71)	(0.76)	(0.73)
N. of subjects	95	82	91
N. of observations	4845	4182	4641

2.4 Results

Table 2.2 summarizes the fraction of comparisons consistent with hypothesis EU-CORE in the three treatments of the experiment. In the Worst and the Bad treatments, approximately 57% of the observations are consistent with hypothesis EU-CORE. The main difference between mixed comparisons in these two treatments is that lotteries have different support in the Worst treatment while sharing the same support in the Bad treatment. However, given the statistically indistinguishable percentage of comparisons consistent with hypothesis EU-CORE, we find that this difference is inessential. In the WorstBest treatment, the percentage of consistent comparisons goes down by approximately five percentage points. The novel approach used to construct mixed comparisons in the WorstBest treatment allows us to detect more failures of hypothesis EU-CORE. Most importantly, it will enable us in Section 2.4 to shed new light on the role of certainty as a driver for such failures.

The analysis that follows aims to uncover the relevance of decision confidence and the availability of certain alternatives as potential explanations for the violations of hypothesis EU-CORE that we observe. We begin by presenting the estimates from the two linear probability models reported in Table 2.3. In both regressions, the dependent variable is equal to one if the observation is consistent with hypothesis EU-CORE, zero otherwise. The two regressions differ in the variable used to measure decision confidence. Regression (1) uses the collected measures of decision confidence, while regression (2) uses response times. We find a strong positive correlation between decision confidence and the likelihood of being consistent with hypothesis EU-CORE. Moreover, subjects who spend more time choosing the preferred lottery are more likely to violate hypothesis EU-CORE.[18]

Table 2.3 also provides new insights into how subjects are violating the indepen-

[18]In Appendix B.5, we show that response times negatively correlate with decision confidence, confirming the intuition that subjects spend more time when not confident about their choices.

Table 2.3: Linear probability models predicting consistency with hypothesis EU-CORE at the individual level.

	Linear Probability Model					
	Dep. Variable: one if hypothesis EU-CORE holds, zero otherwise					
	(1): Measure of Confidence: Self-Report			(2): Measure of Confidence: Response Time		
	Coefficient	Robust Standard Error	P-value	Coefficient	Robust Standard Error	P-value
$Conf$	0.3611	0.0526	0.000			
$RespTime$				-0.0570	0.0123	0.000
$Risky$	-0.2893	0.0295	0.000	-0.3187	0.0294	0.000
$Treatment$						
Bad	-0.0284	0.0265	0.286	-0.0288	0.0275	0.295
WorstBest	-0.0888	0.0299	0.003	-0.0791	0.0306	0.010
$Type$						
0.95-mixed	0.0870	0.0080	0.000	0.0769	0.0076	0.000
0.7-mixed	0.2620	0.0107	0.000	0.2471	0.0109	0.000
$DistToIndiff$	0.3022	0.0552	0.000	0.3279	0.0555	0.000
Constant	0.2980	0.0440	0.000	0.6667	0.0356	0.000

Notes: The dependent variable is one if hypothesis EU-CORE holds, zero otherwise. $Treatment$ is equal to zero if the observation belongs to the Worst treatment, one if it belongs to the Bad treatment, and two if it belongs to the WorstBest treatment. $Type$ is equal to zero if the comparison is unmixed, one if it is 0.95-mixed, and two if it is 0.7-mixed. $Risky$ is equal to one if the subject chooses the riskier alternative, zero otherwise. $Conf$ is the reported decision confidence divided by 100. $RespTime$ is the logarithm of response times measured in seconds. $DistToIndiff$ is the absolute value difference between the fraction of safer choices in a comparison and 0.5. There are $13,668$ observations. Standard errors are clustered at the individual level (268 clusters).

Table 2.4: Percentage of certainty effect violations (unmixed comparisons only).

	Worst	Bad	WorstBest
# Total Violations of Hp. EU-CORE	907	762	838
% Certainty Effect Violations	35.94%	29.92%	45.94%

dence axiom and what is the role of certainty in driving such violations. Subjects who choose riskier over safer lotteries are approximately a 30% more likely to violate hypothesis EU-CORE. Section 2.4 will further document this observation for unmixed comparisons, concluding that the certainty effect is not the most relevant violation of hypothesis EU-CORE. The analysis of decision confidence in Section 2.4 will also provide a rationale for why this happens, showing that subjects are more likely to prefer the risky lottery over the certain prize when they are less confident.

Moreover, subjects are more likely to violate hypothesis EU-CORE in unmixed comparisons than in mixed comparisons. Taken at face value, this result is consistent with the idea that the availability of certain alternatives plays a role in driving violations of hypothesis EU-CORE. Nevertheless, in Section 2.4, we will show that this conclusion is not entirely robust to a more sophisticated analysis that allows us to control for the different stringency of the requirements that hypothesis EU-CORE imposes on different types of comparisons.

The Prevalence of the Reverse Certainty Effect

Subjects violate hypothesis EU-CORE in line with the certainty effect if they choose the certain alternative in an unmixed comparison and the riskier lottery in one of the associated mixed comparisons. The reverse certainty effect refers instead to the opposite behavioral pattern: subjects choosing the risky lottery in an unmixed comparison and then switching to the safer lottery in one of the associated mixed comparisons. Any failures of hypothesis EU-CORE in unmixed comparisons can be then classified as a certainty effect or a reverse certainty effect violation. In this section, we will study which of the two behavioral patterns is more frequent in our experiment.

Table 2.4 documents that in all treatments of the experiment, certainty effect violations are less frequent than reverse certainty effect violations. However, an important aspect that is worth considering to compare the relevance of these two types of violations is that their emergence in an experiment may be inflated or deflated depending on the relative attractiveness of risky lotteries. For instance, if most subjects prefer risky lotteries over certain prizes in unmixed comparisons, we will have more data

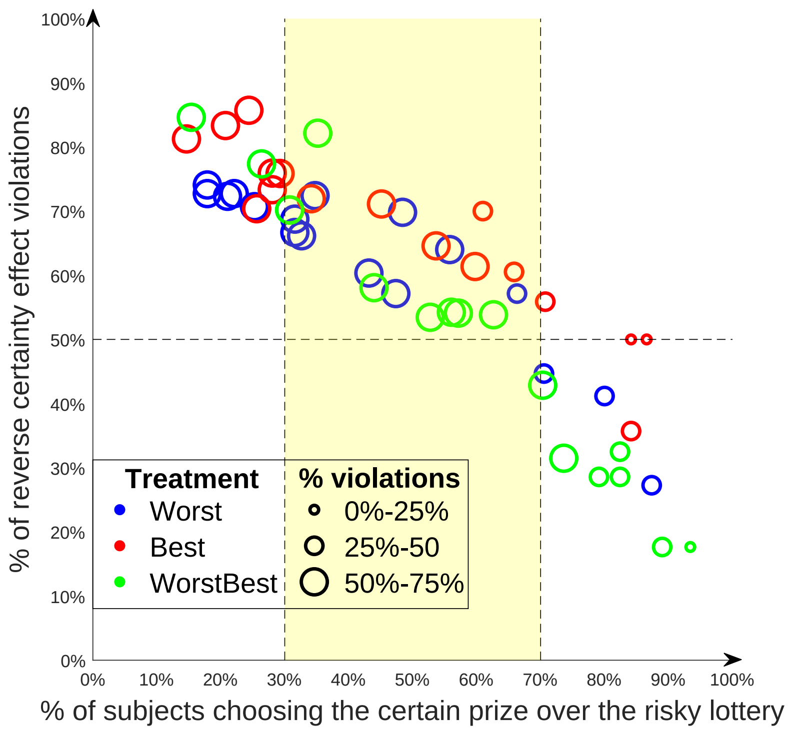

Figure 2.3: Percentage of safe choices and certainty effect violations.

to test certainty effect violations. In the benchmark scenario of an EU subject who makes mistakes, more data to test the certainty effect automatically translates into higher expected certainty effect violations.

For this reason, we examine the emergence of the certainty effect controlling for the fraction of subjects in each unmixed comparison that prefer the risky lottery. Figure 2.3 shows for each unmixed comparison the percentage of subjects choosing the certain prize on the x-axis and the percentage of certainty effect violations over all violations of hypothesis EU-CORE on the y-axis. The blue circles represent unmixed comparisons in the Worst treatment, the red circles in the Bad treatment, and the green circles in the WorstBest treatment. The size of the circles informs about the overall percentages of violations of hypothesis EU-CORE, with bigger circles corresponding to higher ranges of percentages.

The fact that the relevance of the certainty effect increases as the fraction of subjects choosing the safe lottery in unmixed comparisons increases suggests that one has to take seriously the potential bias arising from the unbalance in the available data discussed above. Nevertheless, “fair" comparisons can be made by looking at unmixed comparisons in which both lotteries are chosen by a non-negligible fraction of subjects. For instance, the shaded yellow region in Figure 2.3 includes all the unmixed comparisons for which the percentage of subjects who chose the certain prize is between 30% and 70%. For comparisons in this region, reverse certainty effect violations are always more frequent than certainty effect violations.

The size of the circles in Figure 2.3 provides further evidence against the relevance of the certainty effect. As the fraction of subjects choosing the certain prize increases, the size of the circles tends to be smaller. In other words, the more data we have to observe certainty effect violations of hypothesis EU-CORE, the smaller the total number of hypothesis EU-CORE failures that we observe.[19]

A Possible Mechanism: Decision Confidence

The estimates of the linear probability models in Table 2.3 indicate that hypothesis EU-CORE is more likely to hold when subjects declare to be confident about their decisions. Figure 2.4 provides a graphical representation of this result, describing the empirical cumulative distribution functions of decision confidence for observations in which hypothesis EU-CORE holds (red distribution) and does not hold (blue distribution). The red distribution in Figure 2.4 stochastically dominates the blue one, indicating that conditional on being consistent with hypothesis EU-CORE, subjects in our experiment tend to report higher confidence levels.[20]

The positive correlation between decisions with low confidence and failures of hypothesis EU-CORE summarized in Figure 2.4 on the one hand, and the prevalence of the reverse certainty effect documented in Section 2.4 on the other hand, jointly suggest that subjects tend to prefer risky lotteries over certain prizes when reporting low confidence levels. The PCI axiom theorizes this form of incaution in a way that is easily understandable using the notion of EU core. In its canonical formulation, a subject i satisfies the PCI axiom if for all lotteries $p, r \in \triangle(X)$, prizes $x \in X$ and weights $\lambda \in [0, 1]$,

$$\delta_x \succsim_i p \Rightarrow \delta_x \succsim_i^* p.$$

In words, the PCI axiom precludes subject i from violating the independence axiom in line with the certainty effect but does not impose any constraint on the reverse

[19]In Appendix B.4, we provide additional evidence of the relevance of the reverse certainty effect by studying the implications of the NCI and PCI axioms on risk attitude.

[20]In Appendix B.3 we further explore the relationship between hypothesis EU-CORE and decision confidence using the notion of “indecisiveness" introduced by Cerreia-Vioglio, Dillenberger, and Ortoleva (2015). An individual is more indecisive than another if his EU core is smaller in the sense of set inclusion. Our analysis provides an empirical justification for the use of the term “indecisive" showing that more indecisive individuals tend to report lower levels of decision confidence.

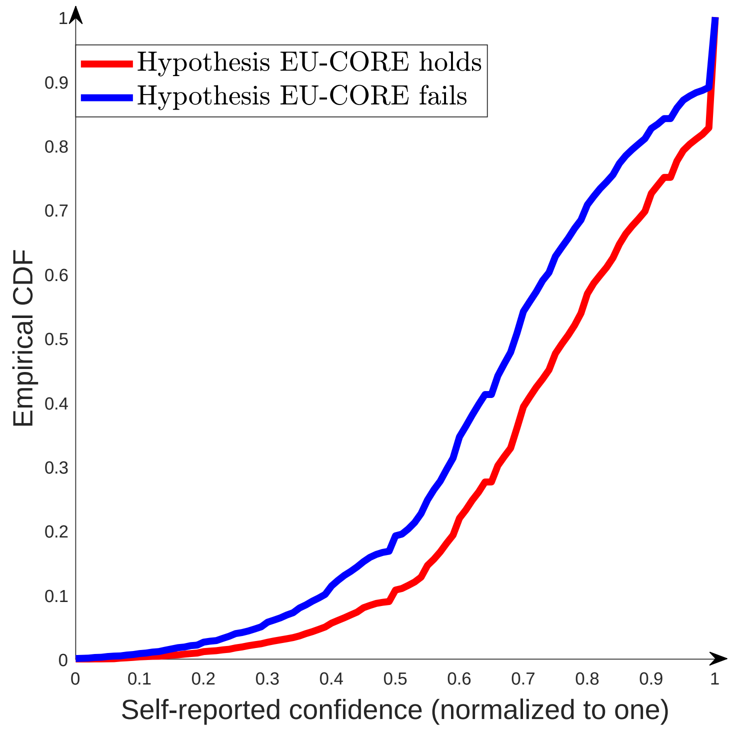

Figure 2.4: Hypothesis EU-CORE and confidence in all treatments.

certainty effect. An equivalent and insightful way to express the PCI axiom is: for all lotteries $p \in \triangle(X)$ and prizes $x \in X$,

$$\neg\left[\delta_x \gtrsim^* p\right] \Rightarrow p \succ \delta_x.$$

Building on the interpretation of the EU core as the subset of uncontroversial comparisons that our analysis supports, the PCI axiom has the following interpretation: when a certain prize is not confidently better than a risky lottery, the risky lottery should be strictly preferred. Motivated by this interpretation, we test whether subjects are more likely to choose the risky lottery or the certain prize in unmixed comparisons when reporting low confidence levels.

Figure 2.5 shows the likelihood of choosing risky lotteries in unmixed comparisons as a function of the reported level of decision confidence. As decision confidence decreases, subjects are more likely to choose risky lotteries over certain prizes. This finding is consistent with the idea hinted by the PCI axiom of individuals being incautious rather than cautious when reporting low levels of decision confidence. Moreover, it provides a rationale for the prevalence of the reverse certainty effect documented in Section 2.4.

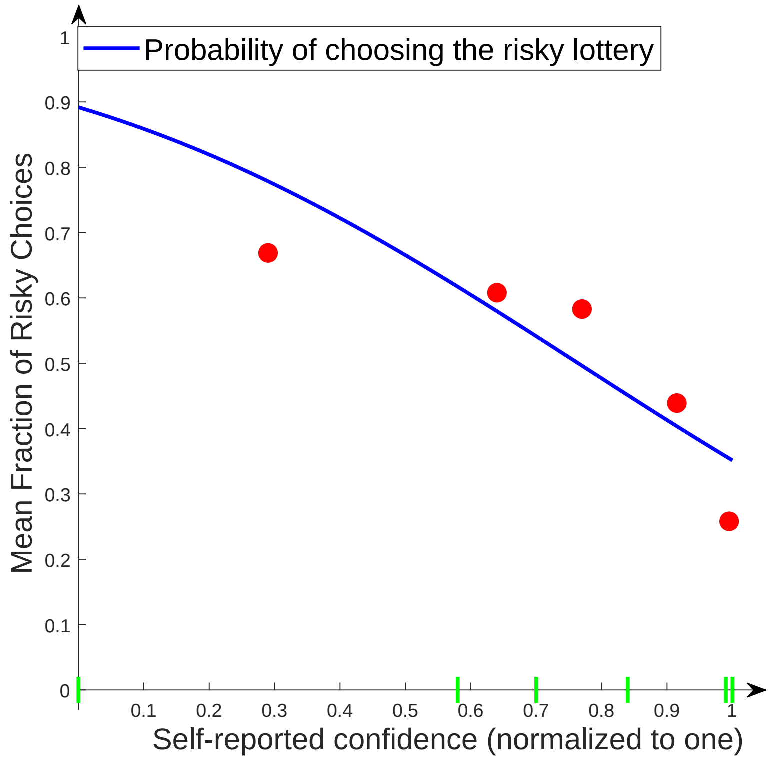

Figure 2.5: Confidence and preference for risky lotteries in unmixed comparisons.

Notes: We partition observations into five categories of equal size based on decision confidence. The red markers on the x-axis denote the thresholds for each category. The red dots placed at the average confidence values in each category represent the fractions of observations in which the risky lottery is preferred over the certain prize. The predicted probabilities of choosing the risky lottery are computed using a probit model.

Certainty Does Not Always Matter

Our analysis of behavior thus far highlights the reverse certainty effect as the most relevant violation of hypothesis EU-CORE in unmixed comparisons. In this section, we take a step back and study whether the mere presence of a certain alternative in a comparison is predictive of more violations of hypothesis EU-CORE. The estimation results in Table 2.3 provide a first positive answer to this question, showing that hypothesis EU-CORE is more likely to fail in unmixed than in mixed comparisons. However, using the same logic adopted to compare certainty effect and reverse certainty effect violations, we now show that this result does not account for the amount of information we have to disprove hypothesis EU-CORE for different categories of comparisons.

Let us consider the four comparisons from the Worst treatment represented in Figure 2.6 and imagine that a subject declared to prefer lottery s_1 over lottery r_1. Disproving

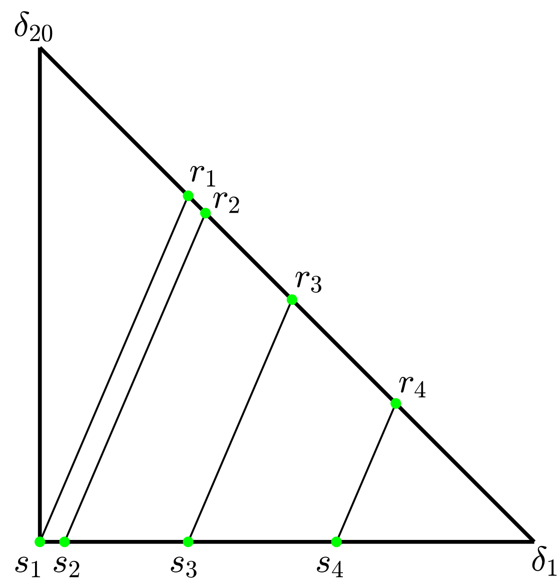

Figure 2.6: Four comparisons from the Worst treatment.

hypothesis EU-CORE for the unmixed comparison (s_1, r_1) amounts to observing a preference for the riskier lottery (r_2, r_3 or r_4) in any of the three associated mixed comparisons. On the contrary, hypothesis EU-CORE for mixed comparisons does not impose any constraint on the preferences expressed in the unmixed comparison. For instance, let us imagine that a subject preferred lottery s_2 over lottery r_2 in the 0.95-mixed comparison (s_2, r_2). Disproving hypothesis EU-CORE for this comparison amounts to observing a preference for the riskier lottery (r_3 or r_4) in any of the remaining mixed comparisons.

Consequently, we have more data to disprove hypothesis EU-CORE for unmixed comparisons than we have for mixed comparisons.[21] To see why this asymmetry can lead to overestimating the role of certainty, let us consider an individual that satisfies EU but with probability 0.1 makes independent mistakes in each of the four comparisons in Figure 2.6 by choosing the least-preferred option. In this case, the probability that this individual satisfies hypothesis EU-CORE in the unmixed comparison (s_1, r_1) is 0.6562, in the 0.95-mixed comparison (s_2, r_2) is 0.73 while in the 0.7-mixed comparison (s_3, r_3) is 0.82. We now describe how we account for this asymmetry by exploiting the error model proposed by Harless and Camerer (1994).

We summarize subjects' choices over the four comparisons in Figure 2.6 by strings

[21]It is worth noting that this observation applies to any experiment that aims at testing the independence axiom.

of chosen lotteries. For instance, choosing lottery s_i over lottery r_i for every index $i \in \{1, 2, 3, 4\}$ corresponds to the string $s_1 s_2 s_3 s_4$. In the error model of Harless and Camerer (1994), subjects have strict preferences over lotteries but can make mistakes choosing the least-preferred lottery. In each comparison, mistakes happen with probability $\epsilon \in (0, 1)$ and are independent across choices. For instance, a subject with true preferences $s_1 s_2 s_3 s_4$ with probability $\epsilon(1 - \epsilon)^3$ makes one error and report $r_1 s_2 s_3 s_4$, $s_1 r_2 s_3 s_4$, $s_1 s_2 r_3 s_4$ or $s_1 s_2 s_3 r_4$. We denote by $x_1 x_2 x_3 x_4$ a generic string of chosen lotteries and define by $p(x_1 x_2 x_3 x_4)$ the fraction of subjects in the experiment for which in the absence of mistakes we would observe $x_1 x_2 x_3 x_4$, where $x_i \in \{s_i, r_i\}$ and $i \in \{1, 2, 3, 4\}$. For instance, $p(s_1 s_2 s_3 s_4)$ is the fraction of subjects preferring lottery s_i over lottery r_i for every index $i \in \{1, 2, 3, 4\}$.

Within this framework, consistency with hypothesis EU-CORE in all comparisons amounts to assuming EU:

$$p(s_1 s_2 s_3 s_4) + p(r_1 r_2 r_3 r_4) = 1. \tag{EU}$$

If, instead, we allow hypothesis EU-CORE to fail in unmixed comparisons, the relaxed model becomes:

$$\sum_{x_1 \in \{s_1, r_1\}} p(x_1 s_2 s_3 s_4) + p(x_1 r_2 r_3 r_4) = 1. \tag{CC}$$

Therefore, no matter how “close to certainty" one of the two lotteries in a comparison is, model CC requires consistency with hypothesis EU-CORE. Finally, allowing for failures of hypothesis EU-CORE in unmixed and 0.95-mixed comparisons leads us to the following model:

$$\sum_{x_2 \in \{s_2, r_2\}} \sum_{x_1 \in \{s_1, r_1\}} p(x_1 x_2 s_3 s_4) + p(x_1 x_2 r_3 r_4) = 1. \tag{AC}$$

In other words, model AC requires consistency with hypothesis EU-CORE only for comparisons in which both lotteries are “away from certainty".

In each of the three model specifications, the fractions of true preferences and the error term can be estimated using maximum likelihood estimation.[22] The unit of observation in this analysis is the pattern of choices in an unmixed comparison and the three associated mixed comparisons. Figure 2.6 shows an example of unmixed and associated mixed comparisons from the Worst treatment. Each treatment has seventeen unmixed comparisons with their own associated three mixed comparisons. Therefore, we estimate our three models in each treatment seventeen times, one for each unmixed and associated mixed comparison.

[22]We refer to Harless and Camerer (1994) for a detailed description of the likelihood function.

Table 2.5: Likelihood ratio tests results.

	Worst	Bad	WorstBest
CC-AC	7	9	0
EU-AC	10	3	1
CC-EU	0	0	0
EU-EU	0	5	16

Notes: Each treatment has seventeen patterns of comparisons. This table classifies each pattern of comparisons into four possible categories. We denote by i-j the category of all patterns in which model $i \in \{\text{EU}, \text{CC}\}$ prevails in a likelihood ratio test between model EU and model CC, while model $j \in \{\text{EU}, \text{AC}\}$ in a likelihood ratio test between model EU and model AC.

To evaluate the relevance of certainty for violations of hypothesis EU-CORE, we perform two likelihood ratio tests. The first test compares model EU with model CC, while the second test model EU with model AC. Table 2.5 summarizes the results of these likelihood ratio tests. In the Worst treatment, model EU is always rejected against either model CC or model AC. In the Bad treatment, accommodating for failures of hypothesis EU-CORE in unmixed comparisons also allows explaining our data significantly better, with the exception of five patterns of comparisons. However, the results in the WorstBest treatment completely overturn this conclusion. In this latter treatment, for sixteen out of seventeen patterns, model EU is never rejected. In other words, allowing for violations of hypothesis EU-CORE in unmixed or 0.95-mixed comparisons does not help to explain our data better.

2.5 Discussion

This study sheds new light on what drives violations of the independence axiom. We conduct an experimental investigation involving choices between risky lotteries. Our main finding is that subjects are more likely to be consistent with the independence axiom when they report high decision confidence levels. In this way, we provide empirical support for the psychological interpretation of the EU core as the subset of the uncontroversial rankings. We believe that exploiting the notion of EU core into experimental works, as we do in this paper, represents a promising direction to expand our understanding of decision-making under risk.

Moreover, we analyze decision-making under low decision confidence. Contrary to the certainty effect rationale for independence violations, subjects are more likely to choose a risky lottery over a certain prize and violate the independence axiom when not confident. Given the extensive evidence on the certainty effect and the impact that this evidence had and still has on new theoretical models, more research is plainly needed to test the robustness of our conclusion. An important insight of our work is that the certainty effect may be less relevant in environments where subjects face a greater variety of lotteries than in the Allais paradox.

Our data also questions the relevance of certainty itself. In the WorstBest treatment, where we construct mixed comparisons alternating the third common lottery, we detect more independence violations. Remarkably, we also find that in this treatment, the presence of certain alternatives does not increase independence violations. To our knowledge, this is the first study that alternates the third common lotteries to build mixed comparisons. Given the different conclusions we obtain in the $WorstBest$ treatment, we believe that exploring new ways to construct mixed comparisons constitutes a promising line of research for studying the independence axiom.

References

Agranov, Marina and Pietro Ortoleva (2017). "Stochastic Choice and Preferences for Randomization". In: *Journal of Political Economy* 125.1, pp. 40–68.

Andreoni, James and William Harbaugh (2009). "Unexpected utility: Experimental tests of five key questions about preferences over risk". In.

Andreoni, James and Charles Sprenger (2010). "Certain and uncertain utility: The allais paradox and five decision theory phenomena". In: *Levine's Working Paper Archive* 926159295.

– (2011). *Uncertainty equivalents: Testing the limits of the independence axiom*. Tech. rep. National Bureau of Economic Research.

– (2012). "Risk preferences are not time preferences". In: *American Economic Review* 102.7, pp. 3357–76.

Arts, Sara, Qiyan Ong, and Jianying Qiu (2022). "Measuring subjective decision confidence". In.

Bernheim, B Douglas and Charles Sprenger (2020). "On the empirical validity of cumulative prospect theory: Experimental evidence of rank-independent probability weighting". In: *Econometrica* 88.4, pp. 1363–1409.

Blavatskyy, Pavlo, Andreas Ortmann, and Valentyn Panchenko (Feb. 2022). "On the Experimental Robustness of the Allais Paradox". In: *American Economic Journal: Microeconomics* 14.1, pp. 143–63. DOI: 10.1257/mic.20190153. URL: https://www.aeaweb.org/articles?id=10.1257/mic.20190153.

Blavatskyy, Pavlo, Valentyn Panchenko, and Andreas Ortmann (2022). "How common is the common-ratio effect?" In: *Experimental Economics*, pp. 1–20.

Blavatskyy, Pavlo R (2010). "Reverse common ratio effect". In: *Journal of Risk and Uncertainty* 40.3, pp. 219–241.

Camerer, Colin F (1992). "Recent tests of generalizations of expected utility theory". In: *Utility theories: Measurements and applications*. Springer, pp. 207–251.

Cerreia-Vioglio, Simone (2009). *Maxmin expected utility on a subjective state space: Convex preferences under risk*. Tech. rep. Mimeo, Bocconi University.

Cerreia-Vioglio, Simone, David Dillenberger, and Pietro Ortoleva (2015). "Cautious expected utility and the certainty effect". In: *Econometrica* 83.2, pp. 693–728.

– (2020). "An explicit representation for disappointment aversion and other betweenness preferences". In: *Theoretical Economics* 15.4, pp. 1509–1546.

Dillenberger, David (2010). "Preferences for one-shot resolution of uncertainty and Allais-type behavior". In: *Econometrica* 78.6, pp. 1973–2004.

Enke, Benjamin and Thomas Graeber (2019). *Cognitive uncertainty*. Tech. rep. National Bureau of Economic Research.

Halevy, Yoram (2008). "Strotz meets Allais: Diminishing impatience and the certainty effect". In: *American Economic Review* 98.3, pp. 1145–62.

Harless, David W and Colin F Camerer (1994). "The predictive utility of generalized expected utility theories". In: *Econometrica: Journal of the Econometric Society*, pp. 1251–1289.

Jain, Ritesh and Kirby Nielsen (2022). "A systematic test of the independence axiom near certainty". In.

Kahneman, Daniel and Amos Tversky (Mar. 1979). "Prospect Theory: An Analysis of Decision under Risk". In: *Econometrica* 47.2, pp. 263–291. URL: https://ideas.repec.org/a/ecm/emetrp/v47y1979i2p263-91.html.

Machina, Mark J (1987). "Choice under uncertainty: Problems solved and unsolved". In: *Journal of Economic Perspectives* 1.1, pp. 121–154.

– (1982). ""Expected Utility" analysis without the independence axiom". In: *Econometrica* 50.2, pp. 277–323. ISSN: 00129682, 14680262. URL: http://www.jstor.org/stable/1912631 (visited on 02/21/2024).

Marschak, Jacob (1950). "Rational behavior, uncertain prospects, and measurable utility". In: *Econometrica: Journal of the Econometric Society*, pp. 111–141.

Segal, Uzi (1990). "Two-stage lotteries without the reduction axiom". In: *Econometrica: Journal of the Econometric Society*, pp. 349–377.

Starmer, Chris (2000). "Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk". In: *Journal of economic literature* 38.2, pp. 332–382.

Tversky, Amos and Daniel Kahneman (1992). "Advances in prospect theory: Cumulative representation of uncertainty". In: *Journal of Risk and uncertainty* 5.4, pp. 297–323.

Chapter 3

DELIBERATE RANDOMIZATION UNDER RISK

3.1 Introduction

Random choices interpreted as the outcome of deliberate randomization are the object of theoretical and experimental works that study decision-making under risk. The recent experimental effort to provide robust evidence about deliberate randomization motivates the growing attention on theoretical models that can rationalize this observed pattern.[1] Yet, while several appealing models belong to this category, the lack of analytical tractability that characterizes a large class of them limits their use in applied research.

Convexity is the axiom that captures preferences for randomization. It requires that if a decision-maker (DM) is indifferent between two lotteries p and q, then any convex combination between p and q is weakly preferred. Preferences that satisfy convexity and few other rationality requirements admit a conservative multi-utility representation: the DM has a set of utility functions and reacts to this multiplicity by evaluating each lottery with the utility function that yields the lowest payoff.[2] One can imagine either a DM with multiple selves or a Rawlsian planner that aggregates the preferences of different individuals.[3]

In many economic problems, a DM chooses an action from a set of available alternatives to maximize his well being. Unfortunately, the conservative multi-utility representation is not differentiable, and consequently, it is not possible to use standard optimization techniques to characterize the properties of the set of optimal actions. To overcome this issue, we provide a general characterization of the DM's optimal action(s) in terms of the strict upper-contour sets of the utilities involved in the representation. In particular, we show that an action maximizes the DM's preferences if and only if the intersection of the strict upper-contour sets of the "worst-off" utilities (i.e., the ones whose evaluation of the DM's action is the lowest) is empty.[4]

---

[1] See, for instance, Agranov and Ortoleva (2017).

[2] Cerreia-Vioglio (2009) first studies the implications of convexity for preferences under risk.

[3] Cerreia-Vioglio (2009) describes a DM with multiple selves which is unsure about one or possibly all of the following: the value of the decision outcome, future tastes, and degree of risk aversion.

[4] Properties of maxmin optima have also been exploited in other contexts. For example, in the

Furthermore, we propose a notion of efficiency that strengthens the requirement of optimality in two ways. We start by calling an optimal action minimal if no other action constitutes a Pareto improvement for the set of worst-off utilities that it induces. We motivate this additional requirement by showing that minimal actions induce the smallest set of worst-off utilities. Next, we define an action efficient if it is minimal and there is no other action that constitutes a Pareto improvement for the set of all utilities.[5] We prove that there is always an efficient action within the set of optimal actions. Consequently, this efficiency notion can always serve as a selection criterion for the case of multiple optimal actions. Moreover, we derive conditions that guarantee the uniqueness of the optimal action.

Our general analysis of the set of optimal actions and their properties lays the groundwork for the second part of our paper, where aiming for higher tractability, we turn to the analysis of deliberate randomization for a DM with two utilities. In this setting, the DM never finds it optimal to select more than two actions with positive probability. The value of this result is twofold. First, it shows that finding optimal actions is easier in the two utility specification than in the generic finite case because it is enough to focus on randomization over at most two actions. Second, it provides a testable implication of this assumption: a DM with two utilities should never be willing to pay any positive monetary amount to pick more than two actions with positive probability.

We call randomization strictly beneficial when it allows the DM to achieve a strictly greater payoff than with any pure action. If the DM is indifferent among all the pure actions, randomization is always strictly beneficial unless preferences do not degenerate to EU. In this scenario, we explicitly derive the support of the optimal random choices. Moreover, we study the case of two pure actions, which is relevant in experimental settings. We then use this result to characterize when an observed preference for randomization can be used to rule out both risk aversion and risk-seeking attitudes in the Cautious Expected Utility (C-EU) model, which is a special case of the class of preferences that we consider.[6]

Finally, we apply our results to non-cooperative game theory, the main analytical tool to build formal economic models. One obstacle in studying games in which

theory of optimal auctions, Chung and Ely (2007) provide sufficient conditions for dominant-strategy mechanisms to have maxmin foundation.

[5]In the auction context, Börgers (2017) refines Chung and Ely (2007)'s criterion in order to exclude dominant-strategy mechanisms that he classifies as "dominated". Our refinement notion is stronger because we require efficient actions to be minimal.

[6]See Cerreia-Vioglio, Dillenberger, and Ortoleva (2015).

players have non-EU preferences is that the notion of Nash equilibrium often needs to be modified. For instance, the Nash requirement of correct conjectures is not well defined in models under uncertainty with multiple beliefs. Instead, the class of convex preferences that we focus on does not feature the same problem. While each player has multiple utility functions, the conjecture is unique. Moreover, all the standard assumptions for the existence of a Nash equilibrium hold for the class of convex preferences that we consider.

Specifically, we study a static game with two players with convex preferences. Each player has two actions and two utility functions. We partition the possible mixed equilibria into three categories: weak, partially strict, and strict. In weak equilibria, players are indifferent in equilibrium between their mixed and pure actions, as in the EU case. Partially strict and strict mixed equilibria instead constitute an element of novelty. In these equilibria, at least one player strictly prefers the equilibrium mixed action to the pure actions in the support. We provide necessary conditions for the existence of these new types of equilibria and we illustrate them in a simple coordination game. In this example, convexity may lead to a multiplicity of mixed Nash equilibria. However, we show that when they exist, only strict mixed Nash equilibria are such that both players play efficient mixed actions.

Related Literature

This paper contributes to the recent theoretical literature that studies stochastic choice as the outcome of deliberate randomization.[7] This strand of contributions builds on the idea first proposed by Machina (1985) that individuals with non-stochastic preferences over lotteries may have an explicit desire to randomize their choices. Battigalli et al. (2017) develop a framework to model random choices under uncertainty. Our paper, instead, focuses on choices under risk, building on the multi-utility representation result obtained by Cerreia-Vioglio (2009) for preferences that satisfy convexity. This representation is appealing because it encompasses several well-known decision criteria under risk, such as the C-EU model of Cerreia-Vioglio, Dillenberger, and Ortoleva (2015) or the maxmin model of Maccheroni (2002).

The premise of this paper is that the multi-utility representation in Cerreia-Vioglio (2009) is not differentiable, so standard optimization techniques to study random choices are not viable. Cerreia-Vioglio, Dillenberger, and Ortoleva (2020) make an analogous remark for betweenness preferences that satisfy Dillenberger's (2010)

[7] See Agranov and Ortoleva (2022) for a recent review of the literature.

NCI axiom, such as the Gul's (1991) model of disappointment aversion. Given that NCI implies convexity, the representation in Cerreia-Vioglio (2009) is more general. At the same time, our focus on a finite set of utilities in practice allows for betweenness violations. For this reason, we see our paper as complementary to Cerreia-Vioglio, Dillenberger, and Ortoleva (2020) for the analysis of preferences in which randomization can be strictly beneficial.

A growing experimental literature supports the hypothesis that subjects make stochastic decisions deliberately. Agranov and Ortoleva (2017) provide evidence in favor of the class of convex preferences that we consider, showing that models of bounded rationality or random preferences cannot rationalize subjects' stochastic behavior in their experiment. Consistently with the conservative multi-utility interpretation of convex preferences, hedging and diversification were the main motivations behind this stochastic behavior. Agranov and Ortoleva (2023) push this observation further, showing not only the existence of questions for which subjects want to randomize but also their prevalence. Our paper provides new testable predictions for models of deliberate randomization, deriving properties of optimal random choices and conditions under which strict preferences for randomization are inconsistent with both risk aversion and risk-seeking attitudes.

Evidence of preferences for randomization extends to strategic settings. Agranov, Healy, and Nielsen (2023) show that randomization is a stable and pervasive feature of several choice environments, including games. In their experiment, a sizable part of individuals displays preferences for randomization in individual decision problems but especially in games. Calford (2021) studies the role of mixed actions for ambiguity averse players with maxmin EU preferences (Gilboa and Schmeidler (1989)). He proves that the set of rationalizable strategies grows larger as preferences for randomization weaken. We also apply our results to static games. However, while each player has multiple utilities in our setting, the conjecture is unique. Consequently, it is not necessary to modify the Nash equilibrium notion, as is the case with models under ambiguity.[8]

Allen and Rehbeck (2021) also study preferences for randomization in settings of strategic interaction by focusing on concave perturbed utility games. In their framework, players' preferences are represented by a general base utility index and an additively separable concave perturbation function. By making different functional form assumptions on the perturbation function, they construct and study properties

[8]See, for instance, Marinacci (2000).

of the best-response functions. Our framework differs because rather than relying on utility perturbations, we start by imposing convexity on players' preferences and exploit the general axiomatic representation in Cerreia-Vioglio (2009) to model preferences for randomization. At the same time, in Section 3.6, we provide a closed-form expression for the best-response function under the assumption that players have maxmin preferences and that randomization is strictly beneficial. We then use this characterization to compute the set of all possible Nash equilibria in a simple coordination problem.

Outline

The rest of the paper is organized as follows. Section 3.2 sets up the decision model. Section 3.3 provides a general characterization of the set of DM's optimal actions. Section 3.4 deals with the efficiency and uniqueness of optimal actions. Section 3.5 studies the implications of deliberate randomization for a DM with two utilities. Section 3.6 applies our results to the analysis of mixed Nash equilibria in a static game where players have convex preferences. Section 3.7 summarizes the main findings and concludes. All the proofs of the statements are collected in the appendix.

3.2 Model

This section begins with the introduction of the decision framework. After that, we describe the conservative multi-utility model of Cerreia-Vioglio (2009) and the additional assumptions we impose on his representation.

Decision Framework

Following Luce and Raiffa (1957),[9] a decision framework is a quartet $\langle A, S, C, \rho\rangle$, where A is a finite set of conceivable pure actions, S is a finite set of states, C is a finite set of consequences and $\rho\colon A \times S \to C$ is the consequence function. Given a generic set X, we denote by $\triangle(X)$ the set of probability distributions over X. The DM can commit his actions to some random devices. We denote by $\mathcal{A} = \triangle(A)$ the set of feasible actions.[10]

The DM has a belief $\mu \in \triangle(S)$ over the states. Every feasible action α, given a

[9] Luce and Raiffa (1957) introduce this framework to study choice under uncertainty. Here, we endow the DM with a subjective belief over the states.

[10] $\triangle(A)$ is the set of conceivable actions. In principle, not all conceivable actions are feasible: $\mathcal{A} \subseteq \triangle(A)$. In this paper, we assume $\mathcal{A} = \triangle(A)$.

belief μ induces a lottery according to the stochastic outcome function:

$$\hat{\rho}\colon \mathcal{A} \times \triangle(S) \to \triangle(C).$$

The specification of the belief is relevant in Section 3.6, where we consider an application of our results to game theory. In all the other sections, we omit the dependence from the belief in the notation because it plays no specific role.

Preferences

We denote by $\mathbb{E}(\alpha, v)$ the EU of action α, with utility $v\colon C \to \mathbb{R}$:

$$\mathbb{E}(\alpha, v) = \sum_{a\in A} \alpha(a) \sum_{s\in S} \mu(s) v\left(\rho(a, s)\right).$$

We also indicate by $\gtrsim_v$ the binary relation representing the preferences of an EU DM with utility v:

$$\alpha \gtrsim_v \beta \Leftrightarrow \mathbb{E}(\alpha, v) \geq \mathbb{E}(\beta, v),$$

with $\alpha, \beta \in \mathcal{A}$. Moreover, we denote by $\succ_v$ and $\sim_v$ the asymmetric and symmetric parts of $\gtrsim_v$, respectively. Given an action $\alpha \in \mathcal{A}$, we denote by $SUCS_v(\alpha)$ the strict upper-contour set of α for utility v. That is,

$$SUCS_v(\alpha) \coloneqq \{\alpha' \in \mathcal{A} \ : \ \alpha' \succ_v \alpha\}.$$

In words, $SUCS_v(\alpha)$ is the set of actions that utility v strictly prefers to α.

When a preference $\gtrsim$ over $\mathcal{A}$ is complete, transitive, continuous and satisfies convexity,[11] Cerreia-Vioglio (2009) shows the existence of a utility function u that represents $\gtrsim$ as follows: there exist a set of normalized utility functions $\mathcal{W}$ and a function $\mathrm{U}\colon \mathbb{R} \times \mathcal{W} \to [-\infty, +\infty]$ such that for all $\alpha \in \mathcal{A}$,[12]

$$u(\alpha) = \inf_{v\in\mathcal{W}} \mathrm{U}\left[\mathbb{E}(\alpha, v), v\right]. \tag{$\star$}$$

For every utility $v \in \mathcal{W}$, the DM computes the EU of action α and then distorts it with the function $\mathrm{U}[\cdot, v]$, which we assume to be strictly increasing in the first argument.[13] Of all possible distorted EU evaluations, the DM adopts a conservative criterion assigning to α the smallest one. We further assume that $\mathcal{W}$ is finite so

[11] The preference relation $\gtrsim$ satisfies convexity if and only if for all $\alpha \in \mathcal{A}$, $\beta \in \mathcal{A}$ and $\lambda \in (0, 1)$, $\alpha \sim \beta \Rightarrow \lambda\alpha + (1 - \lambda)\beta \gtrsim \alpha$.

[12] We fix an arbitrary consequence $c \in C$ and define $\mathcal{W} = \mathcal{W}_1 = \{v \in \mathbb{R}^C \ : v(c) = 1\}$.

[13] Cerreia-Vioglio (2009) proves that $\mathrm{U}[\cdot, v]$ must be increasing in the first argument. The additional requirement that we impose is satisfied in the special cases of Maccheroni (2002) and Cerreia-Vioglio, Dillenberger, and Ortoleva (2015).

that the smallest evaluation is always well defined. We call an action optimal if it maximizes ($\star$).

Because this representation relies on minimal assumptions for $\succsim$, it encompasses several decision models under risk. If $\mathcal{W}$ is a singleton, the representation reduces to EU. When $\mathrm{U}[x, v] = x$ for all $v \in \mathcal{W}$ and $x \in \mathbb{R}$, we obtain the maxmin EU model of Maccheroni (2002). Finally, if $\mathrm{U}[x, v] = v^{-1}(x)$ for all $v \in \mathcal{W}$ and $x \in \mathbb{R}$, we get the C-EU model of Cerreia-Vioglio, Dillenberger, and Ortoleva (2015).

Given an action $\alpha \in \mathcal{A}$, denote by S_α the support of α and by M_α the set of worst-off utilities that α induces:

$$M_\alpha := \operatorname*{arg\,min}_{v \in \mathcal{W}} \mathrm{U}\left[\mathbb{E}(\alpha, v), v\right].$$

Moreover, given a utility function $v \in \mathcal{W}$, denote by P_v the set of pure actions for which v belongs to the induced set of worst-off utilities:

$$P_v := \{a \in A | v \in M_a\}.$$

3.3 Optimal Actions

Our main result characterizes the set of optimal actions in terms of the strict upper-contour sets of the worst-off utilities that these actions induce.

Proposition 1. *Action $\alpha^* \in \mathcal{A}$ is optimal if and only if* $\bigcap_{v \in M_{\alpha^*}} SUCS_v(\alpha^*) = \emptyset$.

Proposition 1 establishes that an action is optimal whenever there is no other action that is strictly better for all the worst-off utilities that the action induces. Suppose that the intersection of the strict upper-contour sets of all the worst-off utilities that action α^* induces is empty. Then, for all actions $\alpha \in \mathcal{A}$, there must exists a utility $v \in M_{\alpha^*}$ such that $\alpha^* \succsim_v \alpha$. Consequently, action α^* is optimal.

For the other direction, suppose that the intersection of the strict upper-contour sets of all the worst-off utilities that action α^* induces is non-empty. Then, there must exist an action α that is strictly better than α^* for all utilities $v \in M_{\alpha^*}$. Given that the set of utilities $\mathcal{W}$ is finite, the payoffs of action α^* for utilities that do not belong to M_{α^*} must be larger than the payoff of the worst-off utilities in M_{α^*} by some finite amount, say $\epsilon > 0$. Because all utilities are continuous, it is possible to mix action α^* with a little bit of α to make all utilities in M_{α^*} better off without rendering anyone outside M_{α^*} worst off by more than ϵ. Therefore, action α^* is not optimal.

Proposition 1 hints at a strategy to verify whether an action α^* is optimal: check whether the intersection of the strict upper-contour sets of all the worst-off utilities in M_{α^*} is empty. The next proposition introduces an indirect tool to simplify this task.

Proposition 2. *Action $\alpha^* \in \mathcal{A}$ is optimal if and only if $\bigcap_{v \in M_{\alpha^*}} SUCS_v(\alpha) = \emptyset$ for all $\alpha \in \mathcal{A}$ with $S_\alpha \subseteq S_{\alpha^*}$.*

According to Proposition 2, an action α^* is optimal whenever there are no actions α and $\tilde{\alpha}$ such that the support of action α^* contains the support of action $\tilde{\alpha}$, and action α is strictly better than action $\tilde{\alpha}$ for all utilities in M_{α^*}. For instance, suppose that for all utilities $v \in M_{\alpha^*}$ and for some pure actions $a \in S_{\alpha^*}$ and $\tilde{a} \in A$, we have $\tilde{a} \succ_v a$. By Proposition 2, we can conclude that α^* is not optimal.

The argument for the proof of Proposition 2 goes as follows. Take an action $\tilde{\alpha}$ with $S_{\tilde{\alpha}} \subseteq S_{\alpha^*}$ and suppose that there exists another action α that is strictly better for all utilities in M_{α^*}. Given that the set of utilities $\mathcal{W}$ is finite, the payoffs of action α^* for utilities that do not belong to M_{α^*} must be larger than the payoff of the worst-off utilities in M_{α^*} by some finite amount, say $\epsilon > 0$. Because all utilities are continuous, it is possible to add a little bit of α and subtract a little bit of $\tilde{\alpha}$ from action α^* to make all utilities in M_{α^*} better off without rendering anyone outside M_{α^*} worse off by more than ϵ. Notice that the resulting action is well defined because $S_{\tilde{\alpha}} \subseteq S_{\alpha^*}$. Therefore, action α^* is not optimal.

Representation in the Marschak–Machina Triangle

We conclude this section with a graphical representation of the results in Propositions 1 and 2. Figure 3.1 shows an example with three pure actions (a, b and c) and three utility functions (v_1, v_2 and v_3) using a revisitation of the MM triangle.[14] Every point in the triangle corresponds to the lottery associated with an action. The figure also includes the indifference curves for the three utilities. Given that the level of the indifference curves matters, we make it explicit through the thickness of the curves. The indifference curves of utilities v_1, v_2, and v_3 passing through $\hat{\alpha}$ have the same thickness and thus achieve the same level of utility. At the same time, the indifference curve of utility v_3 passing through α^* is thicker than the indifference curve passing through $\hat{\alpha}$ because it is associated with a higher level of utility.

[14] The canonical MM triangle represents the set of all lotteries involving three fixed outcomes. Here instead, we represent the set of all lotteries arising from random choices that involve three actions.

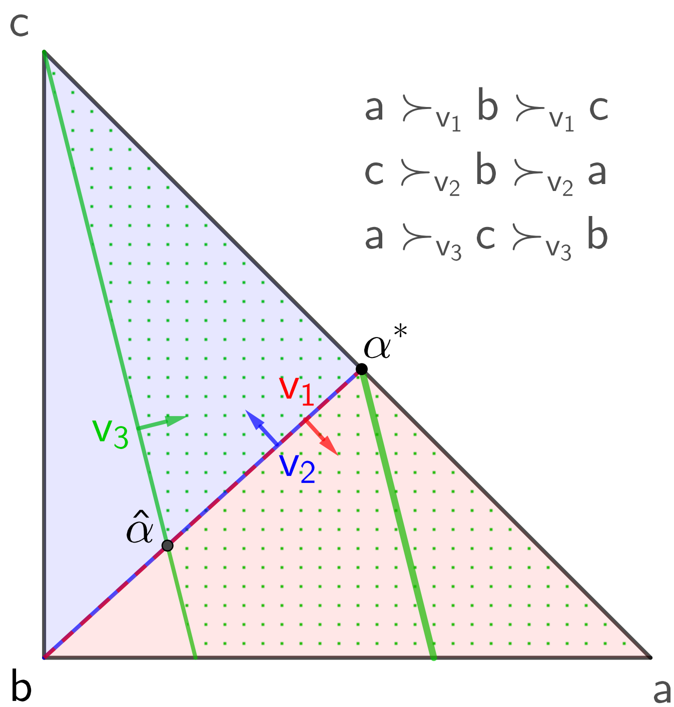

Figure 3.1: Example with $\mathcal{A} = \triangle(\{a, b, c\})$ and $\mathcal{W} = \{v_1, v_2, v_3\}$.

According to Proposition 1, an action α is optimal if there is no other action that is strictly better for all utilities in M_α. For instance, let us consider the mixed action $\hat{\alpha}$ and notice that $M_{\hat{\alpha}} = \{v_1, v_2, v_3\}$. Understanding whether action $\hat{\alpha}$ is optimal amounts to check whether the intersection of the strict upper-contour sets of the three utilities at $\hat{\alpha}$ is empty. In Figure 3.1, the shaded red region is the strict upper-contour set for v_1, the shaded blue region for v_2, and the dotted green region for v_3. As it is clear from the figure, the intersection is empty: to make utility v_2 strictly better off, it is necessary to make utility v_1 strictly worse off and vice versa. Therefore, by Proposition 1 action $\hat{\alpha}$ is optimal.

To appreciate the practical use of Proposition 2, consider a richer decision framework with the same set of utility functions but with a larger set of pure actions A' such that $\{a, b, c\} \subset A'$. Suppose that we are interested in verifying whether an action α with support $S_\alpha = A'$ is optimal. Thanks to Proposition 2, it is still possible to address this task by looking at the MM triangle in Figure 3.1. Indeed, if there are two actions in the triangle such that for all utilities in M_α, one action is strictly better than the other, then by Proposition 2 α is not optimal.

3.4 Uniqueness and Efficiency

Proposition 1 characterizes the set of optimal actions in terms of the strict upper-contour sets of the worst-off utilities. This section explores additional properties that optimal actions might satisfy. Figure 3.1 shows that the set of optimal actions does not need to be a singleton. In the example, this set consists of all actions in the segment with extremes $\hat{\alpha}$ and α^*. The next proposition answers the following question: under what condition is an optimal action unique?

Proposition 3. *Let $\alpha^* \in \mathcal{A}$ be an optimal action. Then α^* is unique if and only if there is no action $\alpha \in \mathcal{A}$, with $\alpha^* \neq \alpha$, such that $\alpha^* \sim_v \alpha$ for all $v \in M_{\alpha^*}$.*

Proposition 3 states that an optimal action α^* is unique whenever there is no action α that is indifferent to α^* for all utilities in M_{α^*}. If such action exists, it is possible to mix α^* with a little bit of α. The resulting new action keeps the set of worst-off utilities fixed to M_{α^*} maintaining the same level of minimum utility. Conversely, suppose that there are two actions α^* and α that are optimal. All utilities in M_{α^*} weakly prefer action α to action α^*. Let us build a new action $\hat{\alpha}$ by mixing action α^* with a little bit of α. The resulting action $\hat{\alpha}$ is still optimal. Furthermore, the set of utilities $M_{\hat{\alpha}}$ coincides with all the utilities in M_{α^*} for which α^* is indifferent to α. Therefore, the action $\hat{\alpha}$ is optimal and all the utilities in $M_{\hat{\alpha}}$ are indifferent between actions $\hat{\alpha}$ and α^*.

If the condition in Proposition 3 fails, the set of optimal actions is not a singleton. To reduce the extent of this multiplicity, we propose an efficiency criterion that refines the set of optimal actions. The set of worst-off utilities M_{α^*} plays a key role in determining whether action α^* is optimal. For this to be the case, there must be no other action α that is strictly better than α^* for all utilities in M_{α^*}. A first natural refinement is then to ask that there is no other action α that Pareto dominates α^* in M_{α^*}. That is, there is no other action α such that α is weakly better than α^* for all utilities $v \in M_{\alpha^*}$, and α is strictly better than α^* for at least one utility $v \in M_{\alpha^*}$.

This Pareto efficiency requirement in the set of worst-off utilities relates to the following question: how large is the set of worst-off utilities? Suppose an optimal action α^* is not Pareto efficient in M_{α^*}. In this case, it is possible to find another action whose set of worst-off utilities is strictly smaller in the sense of set inclusion. The following proposition formalizes this intuition.

Proposition 4. *An optimal action α^* is Pareto efficient in M_{α^*} if and only if $M_{\alpha^*} \subseteq M_\alpha$ for any other optimal action α.*

According to Proposition 4, an optimal action α^* is Pareto efficient in M_{α^*} whenever there is no other optimal action that induces a strictly smaller set of worst-off utilities. For instance, let us come back to the scenario in Figure 3.1. The action $\hat{\alpha}$ is optimal because no action is strictly better for all the utilities. However, from Proposition 4, it is possible to conclude that action $\hat{\alpha}$ is not Pareto efficient in $M_{\hat{\alpha}}$. Indeed, any action α in the interval $(\hat{\alpha}, \alpha^*]$ is also optimal and $M_\alpha = \{v_1, v_2\} \subset \{v_1, v_2, v_3\} = M_{\hat{\alpha}}$. We refer to actions that are Pareto efficient in the induced set of worst-off utilities as minimal and denote by M_{min} the set of worst-off utilities that they induce.

Despite any action in the interval $(\hat{\alpha}, \alpha^*]$ is minimal, the most natural action to pick seems α^*, because utilities v_1 and v_2 are always indifferent, while utility v_3 strictly prefers action α^*. In other words, a sensible selection criterion should also impose an efficiency requirement for utilities that are outside M_{min}. This consideration leads us to our definition of efficiency.

Definition 1. *An action $\alpha^* \in \mathcal{A}$ is efficient if it is minimal and there is no other action that Pareto dominates α^* in $\mathcal{W}$.*

In the example of Figure 3.1, α^* is the only efficient action. The next proposition shows that there is always at least one efficient action.

Proposition 5. *For any finite set of utilities $\mathcal{W}$, there always exists an efficient action.*

The existence of a minimal action follows from Proposition 4 and by the fact that the set of utilities $\mathcal{W}$ is finite. If an optimal action is not minimal, then by Proposition 4 there must exist another optimal action that induces a strictly smaller set of worst-off utilities. Given that $\mathcal{W}$ is finite, there must exist a minimal action that induces the smallest set of worst-off utilities.

At this point, it is not possible to directly establish the existence of an efficient action by solving a maximization problem over the set of minimal actions because this set may not be compact as in the example of Figure 3.1. We circumvent this issue as follows. First, we maximize again ($\star$) over the set of optimal actions using $\mathcal{W} \setminus M_{min}$ as set of utility functions. Second, we show that all the actions that solve the maximization problem must be minimal. Third, within this compact subset of minimal actions, we maximize the sum of the expected utilities over all utilities in $\mathcal{W} \setminus M_{min}$. Finally, we prove that any solution to this latter maximization problem is efficient.

3.5 Deliberate Randomization with Two Utilities

In this section, we study the role of deliberate randomization for a DM with convex preferences and two utilities: $\mathcal{W} = \{v_1, v_2\}$. From an operational point of view, we show that this assumption is appealing because it simplifies the structure of the set of optimal actions. At the same time, it still allows interesting deviations from EU. For instance, in the C-EU model, Cerreia-Vioglio, Dillenberger, and Ortoleva (2015) show that two utilities are enough to rationalize the certainty effect in the Allais' common ratio example.

Our first result is that a DM with convex preferences and two utilities never finds it strictly beneficial to select more than two pure actions with positive probability.

Proposition 6. *If $|\mathcal{W}| = 2$, then*

$$\max_{\alpha \in \mathcal{A}} u(\alpha) = \max_{\alpha \in \{\alpha' \in \mathcal{A} \,:\, |S_{\alpha'}| \leq 2\}} u(\alpha).$$

There are two possibilities for any three pure actions in the support of an optimal mixed action: either both utilities are indifferent among them, or they have opposite preferences. Otherwise, the mixed action would not be optimal. In the case of indifference, it is easy to construct another optimal mixed action with smaller support. In the proof, we show that this is also possible in the scenario of opposite preferences.

To fix ideas, consider the example in Figure 3.1 neglecting the role of utility v_3. Utilities v_1 and v_2 have opposite preferences for the pure actions a, b and c. In particular, $a \succ_{v_1} b \succ_{v_1} c$ and $c \succ_{v_2} b \succ_{v_2} a$. We show that if there is a mixed action inside the triangle that is optimal (for instance, action $\hat{\alpha}$), then there must exist a unique mixed action α^* with support $\{a, c\}$ that is indifferent to the pure action b for both utilities. Therefore, starting from $\hat{\alpha}$, one can reduce to zero the probability weight of action b and increase the probability weights of actions a and c by $\alpha^*(a)\hat{\alpha}(b)$ and $\alpha^*(c)\hat{\alpha}(b)$, respectively. The resulting new mixed action has smaller support and is still optimal.

Proposition 6 provides a testable implication of our restriction on the set of utility functions. Experiments that document deliberate randomization typically focus on binary comparisons. For instance, in Agranov and Ortoleva (2023) subjects can use an external randomization device to choose between two lotteries, exactly as in our theoretical framework. To test our restriction on the number of utilities, one can enlarge the set of available lotteries and add a small cost for selecting more than two lotteries with a positive probability. Subjects consistent with the assumption of two utilities should never be willing to pay any positive amount.

Maintaining the assumption of two utilities, we now characterize the mixed actions that maximize the DM's preferences when there are no optimal pure actions. In this case, we call randomization strictly beneficial.

Definition 2. *Randomization is strictly beneficial if*

$$\exists \alpha \in \mathcal{A} : u(\alpha) > \max_{a \in A} u(a).$$

In what follows, we first look at the case where the DM is indifferent among all the pure actions. Then, we conclude by studying what happens when there are only two pure actions.

Indifference

A non-EU DM with convex preferences always strictly benefits from randomization when indifferent among all the pure actions.

Proposition 7. *Assume that* $\arg\max_{a \in A} u(a) = A$. *For any finite set of utilities* $\mathcal{W}$, *randomization is strictly beneficial if and only if there is no utility* $v \in \mathcal{W}$ *such that* $P_v = A$.

Whenever a utility v always belongs to the set of worst-off utilities, then it is as if the DM had EU preferences with utility v. This result holds regardless of the size of $\mathcal{W}$. The next proposition characterizes the set of optimal mixed actions under indifference.

Proposition 8. *Suppose that* $\mathcal{W} = \{v_1, v_2\}$, $\arg\max_{a \in A} u(a) = A$ *and there is no utility* $v \in \mathcal{W}$ *such that* $P_v = A$. *A mixed action* $\alpha \in \mathcal{A}$ *is optimal if and only if the following conditions hold:*

1. $M_\alpha = \{v_1, v_2\}$.
2. $S_\alpha \subseteq \arg\max_{a \in P_{v_1} \setminus P_{v_2}} U\left[\mathbb{E}(a, v_2), v_2\right] \cup \arg\max_{a \in P_{v_2} \setminus P_{v_1}} U\left[\mathbb{E}(a, v_1), v_1\right].$

The evaluation of the optimal mixed action α must be the same for both utilities. Otherwise, it is always possible to increase the minimum evaluation. Moreover, the optimal mixed action must select with positive probability only pure actions for which the two utilities disagree in the evaluations. That is, each pure action must belong to $P_{v_1} \setminus P_{v_2}$ or $P_{v_2} \setminus P_{v_1}$. In light of Proposition 6, it is enough to consider only mixed actions that assign positive probability to two pure actions, one from each set.

Intuitively, when the two utilities have two different evaluations for a pure action, selecting the action with positive probability is strictly beneficial because it helps the DM hedging against his conservative nature. However, when the two evaluations coincide, no hedging is possible. Because of the indifference assumption, utility v_1 assigns the same value to all the actions in $P_{v_1} \setminus P_{v_2}$. Therefore, among these actions, an optimal mixed action must select only those that maximize the evaluation for utility v_2. An analogous argument applies to actions in $P_{v_2} \setminus P_{v_1}$.

Two Actions

In most experiments that document deliberate randomization, there are only two feasible pure actions for each choice. The setting with binary actions is also interesting in several applications, such as the static game we consider in the next section. We begin characterizing strict benefits from randomization when there are only two pure actions.

Proposition 9. *Assume that $A = \{a, b\}$ and $\mathcal{W} = \{v_1, v_2\}$. Randomization is strictly beneficial if and only if the following are true:*

1. *There is no utility $v \in \mathcal{W}$ such that $P_v = A$.*
2. *Either $a \succ_{v_1} b$ and $b \succ_{v_2} a$, or $b \succ_{v_1} a$ and $a \succ_{v_2} b$.*

The DM can find randomization strictly beneficial even if the two pure actions do not ensure the same minimum evaluation. Therefore, the DM must be able to commit credibly to stick with the indications of the randomization device. As for the case of indifference, the DM's preferences must not degenerate to Expected Utility. Moreover, the two utilities must disagree in ranking the two pure actions. Randomization is a valuable tool only when this internal disagreement is present because it allows the DM to hedge against his pessimistic nature in evaluating pure actions.

Besides shedding light on the drivers that make randomization desirable, Proposition 9 also allows studying the DM's risk attitude. In the C-EU model, Cerreia-Vioglio, Dillenberger, and Ortoleva (2015) show that the DM is risk averse (versus risk seeking) if and only if all the utilities in $\mathcal{W}$ are concave (versus convex). Thanks to Proposition 9, it is possible to rule out both risk attitudes considering two pure actions a and b, where a is a mean-preserving spread of b.[15]

Corollary 1. *Assume that $\mathcal{W} = \{v_1, v_2\}$, $A = \{a, b\}$ and that action a is a mean-preserving spread of action b. If incentives to randomize are strict, then a C-EU DM is neither risk averse, nor risk seeking.*

According to Proposition 9, if incentives to randomize are strict, the two utilities must disagree in the ranking between actions a and b. But this necessarily implies that one utility is convex while the other one is concave. Consequently, a C-EU DM is neither risk averse nor risk seeking.

We conclude with the characterization of the optimal and unique mixed action.

Corollary 2. *Assume that $A = \{a, b\}$ and $\mathcal{W} = \{v_1, v_2\}$ and that randomization is strictly beneficial. The action $\alpha \in \mathcal{A}$ is uniquely optimal if and only if $M_\alpha = \mathcal{W}$.*

When there are two utilities, two actions, and incentives to randomize are strict, the unique optimal mixed action equalizes the payoff of the two utilities.

3.6 Games with Convex Preferences

A normal game form is a mathematical structure $\langle N, (S_i)_{i \in N}, C, g \rangle$, where N is a finite set of players, S_i and is the set of available actions for each player i, C is the set of consequences and $g\colon \times_{i \in N} S_i \to C$ is the outcome function that associates consequences with strategy profiles. Therefore, each player i faces the decision framework $\langle S_i, S_{-i}, C, g \rangle$.[16]

Given a conjecture $\mu_i \in \triangle(S_{-i})$, player i chooses $\alpha_i \in \triangle(S_i)$ to maximize

$$u_i(\alpha, \mu_i) = \min_{v \in \mathcal{W}_i} \mathrm{U}\left[\mathbb{E}_{\mu_i}(\alpha, v), v\right],$$

where we make the dependence from the conjecture explicit. Similarly, we write M_{α_i,μ_i}, P_{v,μ_i}, $\succ_{v,\mu_i}$ and $\sim_{v,\mu_i}$ instead of M_{α_i}, P_v, $\succ_v$ and $\sim_v$. A normal-form game with convex preferences G adds to the normal game form the profile $(\mathcal{W}_i)_{i \in N}$ of sets of utility functions on C. Every normal-form game with convex preferences always

[15] We identify actions with the lotteries that they induce. Formally, given a belief μ, we say that action a is a mean-preserving spread of action b if $\hat{\rho}(a, \mu)$ is a mean-preserving spread of $\hat{\rho}(b, \mu)$. For a definition of mean-preserving spread, we refer to Rothschild and Stiglitz (1970).

[16] We denote by $-i = N \setminus \{i\}$ the set of players different from i.

has a Nash equilibrium because all the standard assumptions for existence hold.[17] In what follows, we characterize the set of all possible Nash equilibria when there are two players, each having two pure actions and two utility functions.

Consider a normal-form game with convex preferences G with $N = \{A, B\}$, $S_A = \{a_1, a_2\}$, $S_B = \{b_1, b_2\}$ and $|\mathcal{W}_A| = |\mathcal{W}_B| = 2$. For convenience, we identify mixed actions for players A and B with the probabilities $\alpha \in (0,1)$ and $\beta \in (0,1)$ that they assign to actions a_1 and b_1. We also denote by v and w generic utilities for players A and B. Given a utility $v \in \mathcal{W}_A$ for player A, we represent by β_v the mixed action of player B such that $a_1 \sim_{v,\beta_v} a_2$. Therefore, player A is indifferent between pure actions a_1 and a_2 when using utility v and thinking that player B chooses the mixed action β_v. Similarly, given a utility $w \in \mathcal{W}_B$ for player B, we denote by α_w the mixed action of player A such that $b_1 \sim_{w,\alpha_w} b_2$. We assume that for all $v \in \mathcal{W}_A$ and $w \in \mathcal{W}_B$, $\alpha_w \in [0,1]$ and $\beta_v \in [0,1]$ are well defined. This amounts to imposing that under no single utility, one player has a dominant action.

Strict Mixed Nash Equilibria

The profile of strategies (α_w, β_v) is the mixed Nash equilibrium that would result in a game in which players A and B maximize subjective EU with utilities v and w. Within the subjective EU framework, every player is indifferent between the mixed action played in equilibrium and all the pure actions in the support. When instead players have convex preferences, our analysis thus far shows that given a fixed conjecture about the other player's action, the incentives to play a mixed action may be strict. We now study under what conditions incentives to randomize extend to mixed Nash equilibria of G.

Let $(\alpha, \beta) \in (0,1)^2$ be a mixed Nash equilibrium of G. First, notice that

$$\alpha \in \bar{A} := [\min_{w \in \mathcal{W}_B} \alpha_w, \max_{w \in \mathcal{W}_B} \alpha_w] \text{ and } \beta \in \bar{B} := [\min_{v \in \mathcal{W}_A} \beta_v, \max_{v \in \mathcal{W}_A} \beta_v].$$

These conditions follow directly from the definition of Nash equilibrium. For instance, if $\beta > \max \bar{B}$, then either $a_1 \succ_{v,\beta} a_2$ or $a_2 \succ_{v,\beta} a_1$ for all $v \in \mathcal{W}_A$. In both cases, the mixed action $\alpha \in (0,1)$ for player A can not be a best reply to action β for player B. The next corollary clarifies that, in equilibrium, incentives to randomize are strict for both players A and B depending on whether actions β and α are boundary or interior points of $\bar{B}$ and $\bar{A}$, respectively.[18]

[17] The set $\times_{i \in N} \triangle(S_i)$ is compact and convex. Moreover, for each player $i \in N$, the function u_i is continuous and quasi-concave since the preferences that it represents satisfy convexity.

[18] Given a generic set X, we denote by X^o the set of interior points of X.

Corollary 3. *Let* $(\alpha, \beta) \in (0, 1)^2$ *be a mixed Nash equilibrium of* G*. The following statements are true:*

1. $\alpha \in \bar{A}^{\circ}$ *if and only if* $u_B(\beta, \alpha) > \max\{u_B(b_1, \alpha), u_B(b_2, \alpha)\}$.
2. $\beta \in \bar{B}^{\circ}$ *if and only if* $u_A(\alpha, \beta) > \max\{u_A(a_1, \beta), u_A(a_2, \beta)\}$.

Suppose first that $\alpha \notin A^{\circ}$. Then player B does not strictly benefit from randomization. Indeed, for any $\alpha \notin A^{\circ}$, at least one of the two utilities in $\mathcal{W}_B$ is indifferent between actions b_1 and b_2. By Proposition 9, incentives to randomize can not be strict. Suppose instead that $\alpha \in A^{\circ}$. In this case, player B strictly benefits from randomization. Given that $\alpha \in A^{\circ}$, it must be that one utility in $\mathcal{W}_B$ strictly prefers action b_1 to action b_2, and the other has opposite preferences. By Proposition 9, it is enough to show that there is no utility in $\mathcal{W}_B$ that belongs to the sets of worst-off utilities induced by both pure actions. This latter condition must hold because otherwise, the mixed action β would not be a best reply to the correct conjecture α. An analogous reasoning can be used to prove the second statement of Corollary 3.

In light of Corollary 3, we classify mixed Nash equilibria as follows.

Definition 3. *Let* $(\alpha, \beta) \in (0, 1)^2$ *be a mixed Nash equilibrium of* G*. We call* (α, β)

- *weak if* $\alpha \notin \bar{A}^{\circ}$ *and* $\beta \notin \bar{B}^{\circ}$.
- *partially strict if either* $\alpha \notin \bar{A}^{\circ}$ *and* $\beta \in \bar{A}^{\circ}$ *or* $\alpha \in \bar{A}^{\circ}$ *and* $\beta \notin \bar{A}^{\circ}$.
- *strict if* $\alpha \in \bar{A}^{\circ}$ *and* $\beta \in \bar{A}^{\circ}$.

When the sets $\bar{A}^{\circ}$ and $\bar{B}^{\circ}$ are empty, Corollary 3 implies that there can not be strict mixed Nash equilibria. One example of a game in which this happens is matching pennies. Indeed, as long as the utilities of the two players are strictly increasing, we have $\bar{A} = \bar{B} = \{0.5\}$. For the remaining part of this section, we assume that $\bar{A}^{\circ}$ and $\bar{B}^{\circ}$ are non-empty so that any mixed Nash equilibrium is a priori possible.

The computation of weak equilibria follows the same logic used to compute equilibria under EU. In equilibrium, each player must be indifferent between the two pure actions. We now turn to the analysis of partially strict and strict equilibria, in which at least one player strictly benefits from randomization. In particular, let us focus on player A and suppose that there exists a subset of conjectures $X \subseteq \bar{B}^{\circ}$ under which player A strictly benefits from randomization. The next corollary characterizes the best reply of player A for all conjectures in X, assuming the maxmin EU model.

Table 3.1: Coordination game: outcome function.

		B	
		b_1	b_2
A	a_1	$3,1$	$0,0$
	a_2	$0,0$	$1,3$

Corollary 4. *Let $W_A = \{v_A, w_A\}$. For all conjectures $\beta \in X$, the unique optimal mixed action $\alpha(\beta)$ in the maxmin EU model satisfies*

$$\alpha(\beta) = \frac{\mathbb{E}_\beta[a_2, w_A] - \mathbb{E}_\beta[a_2, v_A]}{\mathbb{E}_\beta[a_2, w_A] - \mathbb{E}_\beta[a_2, v_A] + \mathbb{E}_\beta[a_1, v_A] - \mathbb{E}_\beta[a_1, w_A]}.$$

Corollary 4 provides a closed-form expression that characterizes the best reply of player A for the subset of conjectures under which randomization is strictly beneficial. One insight that emerges from Corollary 4 is that the optimal probability with which player A chooses action a_1 is increasing in $|\mathbb{E}_\beta[a_2, w_A] - \mathbb{E}_\beta[a_2, v_A]|$ and decreasing in $|\mathbb{E}_\beta[a_1, w_A] - \mathbb{E}_\beta[a_1, v_A]|$. That is, in the maxmin EU model, players dislike actions for which there is high variability in their evaluations.

To illustrate all the possible types of mixed Nash equilibria, we consider the coordination game with the outcome function represented in Table 3.1. We assume that players A and B behave according to maxmin EU criterion. Each player has two utility functions, one CARA and one CRRA.[19] Figure 3.2 represents the best replies for the two players.[20] Every intersection point of the two best replies represents a Nash equilibrium. In this example, there are eleven Nash equilibria: two pure, four weak, four partially strict, and one strict. The two pure Nash equilibria are in light blue, the four weak Nash equilibria are in light green, the four partially strict mixed Nash equilibria are in magenta, and the strict Nash equilibrium is in black.

In this simple and analytically tractable scenario of two utility functions for each player, we obtain starkly different predictions from the EU case. From a numerical point of view, convexity may lead to a multiplicity of mixed equilibria. Most importantly, partially strict and strict mixed Nash equilibria do not have an analog under EU. We now extend the notion of efficiency developed in Section 3.4 to profiles of actions and show that strict mixed Nash equilibria are the only type of equilibria that satisfy this notion.

[19]CARA: $v_A(x) = w_A(x) = 1 - \frac{1}{\alpha}e^{-\alpha x}$, with $x \geq 0$ and $\alpha > 0$. CRRA: $v_B(x) = w_B(x) = x^\gamma$, with $x \geq 0$ and $\gamma \in (0, 1)$.

[20]Parameters: $\alpha = 1.52$ and $\gamma = 0.42$.

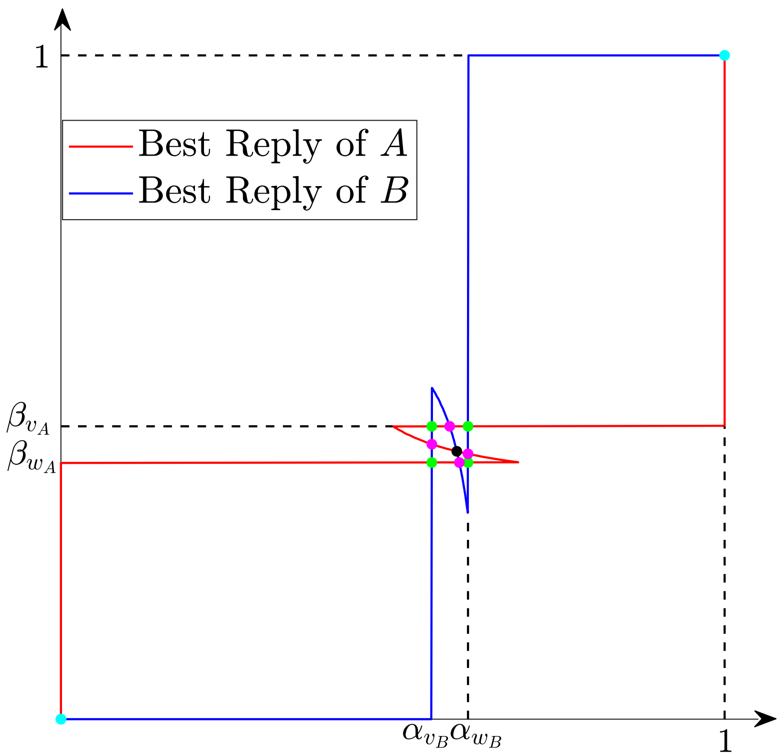

Figure 3.2: Nash equilibria of the coordination game.

Definition 4. *A mixed Nash equilibrium (α, β) of G is efficient if in equilibrium α and β are efficient.*

The next corollary clarifies that only strict mixed Nash equilibria satisfy the efficiency requirement in Definition 4.

Corollary 5. *Let (α, β) be a mixed Nash equilibrium of G and suppose that $\bar{A}^{o}$ and $\bar{B}^{o}$ are non-empty. Then (α, β) is efficient if and only if (α, β) is strict.*

The notion of efficiency introduced in Definition 4 can serve as a selection criterion for settings with multiple mixed Nash equilibria. In the example described in Figure 3.2, there are nine mixed Nash equilibria, but only one is strict and, by Corollary 5, efficient.

Overall, the presence of strict mixed Nash equilibria is the main element of novelty that emerges from our equilibrium analysis in games with convex preferences. We show that, as documented by contemporaneous experimental works, incentives to randomize extend to strategic interaction settings but only under certain conditions. For instance, strict mixed Nash equilibria may arise in a coordination game as the one described in Figure 3.2 but do not exist in matching pennies. Besides being empirically relevant given the evidence of randomization in games, Corollary 5 shows that when they exist, strict mixed Nash equilibria are also normatively appealing because they are the only efficient equilibria.

3.7 Conclusions

Despite the growing theoretical and experimental literature on random choices under risk, the applicability of models that rationalize deliberate randomization is still limited. This paper studies the set optimal actions for a DM whose preferences satisfy convexity, the axiom that makes randomization weakly beneficial. Under convexity, the DM's preferences admit a conservative multi-utility representation: actions are ranked only through the lowest utility valuation they generate.

One drawback of this representation in applications is that it is not differentiable, so standard optimization techniques are not viable. Our main result (Proposition 1) shows that an action is optimal whenever the intersection of the strict upper-contour sets of the worst-off utilities is empty. When more than one action is optimal, we propose Pareto efficiency in the set of worst-off utilities and in the set of all utilities as a selection rule. Proposition 4 clarifies that the first requirement amounts to isolating optimal actions that induce the smallest set of worst-off utilities. The second requirement allows instead to account for utilities outside this set. Proposition 5 guarantees that there is always an optimal action that satisfies both requirements and thus is efficient.

Next, we narrow our attention to random choices for a DM with two utilities. Proposition 6 provides a testable implication of this assumption, proving that a DM with two utilities never finds it strictly beneficial to select more than two actions with positive probability. We then study under what conditions randomization is strictly beneficial and the properties that an optimal random choice must satisfy in two cases: when the DM is indifferent among the pure actions (Proposition 7 and Proposition 8) and when there are only two pure actions (Proposition 9 and Corollary 2).

The binary actions setting recreates the typical environment that subjects face in experiments on randomization under risk. In general, preferences for randomization can coexist with various attitudes towards risk. Our analysis of randomization incentives suggests a new approach to rule out risk aversion and risk seeking in the C-EU model. According to Corollary 1, a C-EU DM is neither risk averse nor risk seeking if incentives to randomize are strict between two actions, one being a mean-preserving spread of the other. Moreover, Corollary 2 shows that when incentives to randomize are strict, the optimal mixed action is unique.

A special case of the decision framework that we study is game theory. We focus on a generic game with two players, each with two actions and two utility functions. The new prediction that arises from our analysis is that strict incentives to randomize extend to strategic interaction settings. Corollary 3 provides necessary conditions for the existence of a new class of mixed Nash equilibria that we call strict because players strictly prefer the equilibrium mixed actions to the pure actions. Corollary 4 derives a closed-form expression of the best-response function for the case in which randomization is strictly beneficial, and players have maxmin preferences. We then exploit this result to compute the mixed Nash equilibria of a simple coordination game. In this example, we find nine mixed Nash equilibria, one of which is strict. Although convexity may lead to a multiplicity of mixed equilibria, we show in Corollary 5 that when they exist, only strict equilibria are such that all the mixed actions are efficient.

References

Agranov, Marina, Paul J Healy, and Kirby Nielsen (2023). “Stable Randomisation”. In: *The Economic Journal* 133.655, pp. 2553–2579.

Agranov, Marina and Pietro Ortoleva (2017). “Stochastic Choice and Preferences for Randomization”. In: *Journal of Political Economy* 125.1, pp. 40–68.

– (2022). “Revealed Preferences for Randomization: An Overview”. In: *AEA Papers and Proceedings* 112, pp. 426–430.

– (July 2023). “Ranges of Randomization”. In: *The Review of Economics and Statistics*, pp. 1–44. ISSN: 0034-6535. DOI: 10.1162/rest_a_01355. eprint: https://direct.mit.edu/rest/article-pdf/doi/10.1162/rest_a_01355/2150363/rest_a_01355.pdf. URL: https://doi.org/10.1162/rest%5C_a%5C_01355.

Allen, Roy and John Rehbeck (2021). “A Generalization of Quantal Response Equilibrium via Perturbed Utility”. In: *Games* 12.1. ISSN: 2073-4336. DOI: 10.3390/g12010020. URL: https://www.mdpi.com/2073-4336/12/1/20.

Battigalli, Pierpaolo et al. (2017). “Mixed extensions of decision problems under uncertainty”. In: *Economic Theory* 63.4, pp. 827–866.

Börgers, Tim (2017). “(No) Foundations of dominant-strategy mechanisms: a comment on Chung and Ely”. In: *Review of Economic Design* 21, pp. 73–82.

Calford, Evan M (2021). “Mixed strategies and preference for randomization in games with ambiguity averse agents”. In: *Journal of Economic Theory* 197, p. 105326.

Cerreia-Vioglio, Simone (2009). *Maxmin expected utility on a subjective state space: Convex preferences under risk*. Tech. rep. Mimeo, Bocconi University.

Cerreia-Vioglio, Simone, David Dillenberger, and Pietro Ortoleva (2015). “Cautious expected utility and the certainty effect”. In: *Econometrica* 83.2, pp. 693–728.

– (2020). “An explicit representation for disappointment aversion and other betweenness preferences”. In: *Theoretical Economics* 15.4, pp. 1509–1546.

Chung, Kim and Jeff Ely (2007). “Foundations of Dominant-Strategy Mechanisms”. In: *Review of Economics Studies* 74.2, pp. 447–476.

Dillenberger, David (2010). “Preferences for one-shot resolution of uncertainty and Allais-type behavior”. In: *Econometrica* 78.6, pp. 1973–2004.

Gilboa, Itzhak and David Schmeidler (1989). “Maxmin expected utility with non-unique prior”. In: *Journal of Mathematical Economics* 18.2, pp. 141–153. ISSN: 0304-4068. DOI: https://doi.org/10.1016/0304-4068(89)90018-9. URL: http://www.sciencedirect.com/science/article/pii/0304406889900189.

Gul, Faruk (1991). “A theory of disappointment aversion”. In: *Econometrica: Journal of the Econometric Society*, pp. 667–686.

Luce, R Duncan and Howard Raiffa (1957). *Games and decisions: Introduction and critical survey*. Wiley, New York.

Maccheroni, Fabio (2002). “Maxmin under risk”. In: *Economic Theory* 19.4, pp. 823–831.

Machina, Mark J (1985). “Stochastic Choice Functions Generated from Deterministic Preferences over Lotteries”. In: *Economic Journal* 95.379, pp. 575–94.

Marinacci, Massimo (2000). “Ambiguous games”. In: *Games and Economic Behavior* 31.2, pp. 191–219.

Rothschild, Michael and Joseph E Stiglitz (1970). “Increasing risk: I. A definition”. In: *Journal of Economic theory* 2.3, pp. 225–243.

Appendix A

APPENDIX TO CHAPTER 1

A.1 Descriptive Analysis

We provide a descriptive analysis of behavior in Block 1 and Block 2 of the experiment.

Block 1

Figure A.1 displays the percentages of various choice patterns observed in CR-tasks (on the left graph) and R-tasks (on the right graph).

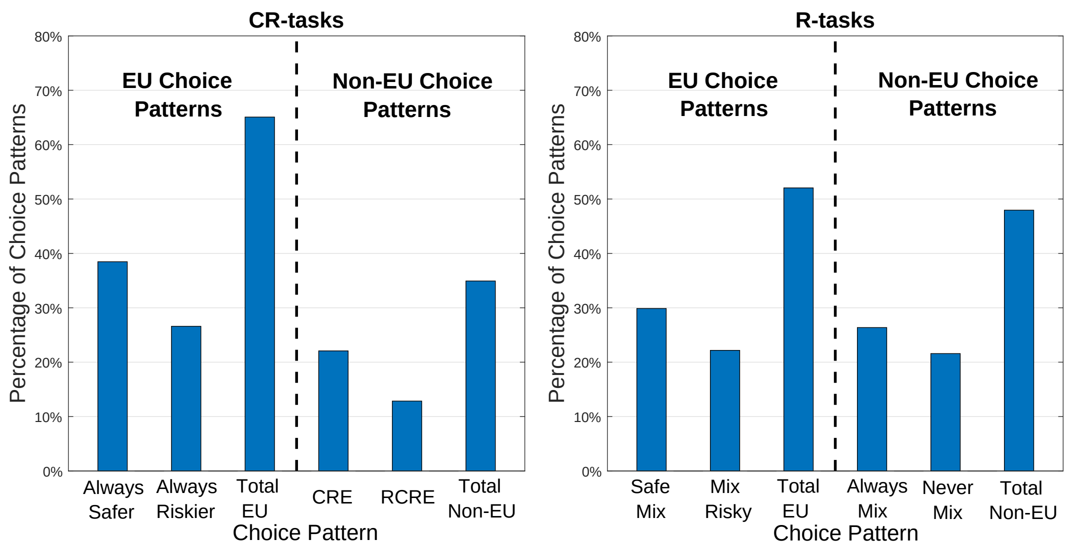

Figure A.1: Percentages of the different choice patterns in CR-tasks and R-tasks.

Notes: There are four possible choice patterns for each pair of lotteries (δ_M, r) in CR-tasks (left graph) and R-tasks (right graph). In CR-tasks, “Always Safer” means consistently selecting the safer lottery, while “Always Riskier” denotes the opposite choice. The CRE and RCRE are the two possible non-EU choice patterns in CR-tasks. Within R-tasks, “Safe Mix” indicates choosing the safe lottery over the mixture, and the mixture over the risky lottery. Vice versa, “Mix Risky” indicates the opposite behavior. For non-EU choice patterns, “Always Mix” indicates always choosing the mixture, while “Never Mix” indicates the opposite behavior. Finally, “Total EU” and “Total Non-EU” indicate, respectively, the aggregate percentages of EU and non-EU choice patterns within CR-tasks (left graph) and R-tasks (right graph).

Within the CR-tasks, two choice patterns are consistent with EU: always choosing the safer lottery ("Always Safer") and always choosing the riskier lottery ("Always

Riskier"), which together make up 65.07% of all choice patterns. Among the non-EU choice patterns in CR-tasks, the CRE is the most frequent, occurring in roughly 9% more cases than the RCRE. In the context of R-tasks, the EU-consistent patterns emerge when subjects either choose the safe lottery over the mixture and then the mixture over the risky one ("Safe Mix"), or vice versa ("Mix Risky"). EU choice patterns in R-tasks account for 52.05% of all choice patterns—a drop of roughly 13% compared to CR-tasks. Turning to non-EU patterns in R-tasks, always choosing the mixture ("Always Mix") constitutes the 26.37% of choice patterns, while never choosing the mixture ("Never Mix") constitutes the 21.58%.

Block 2

Figure A.2 uses box plots to summarize the distribution of risk premia for the three lotteries presented to subjects in Block 2. The distributions of risk premia

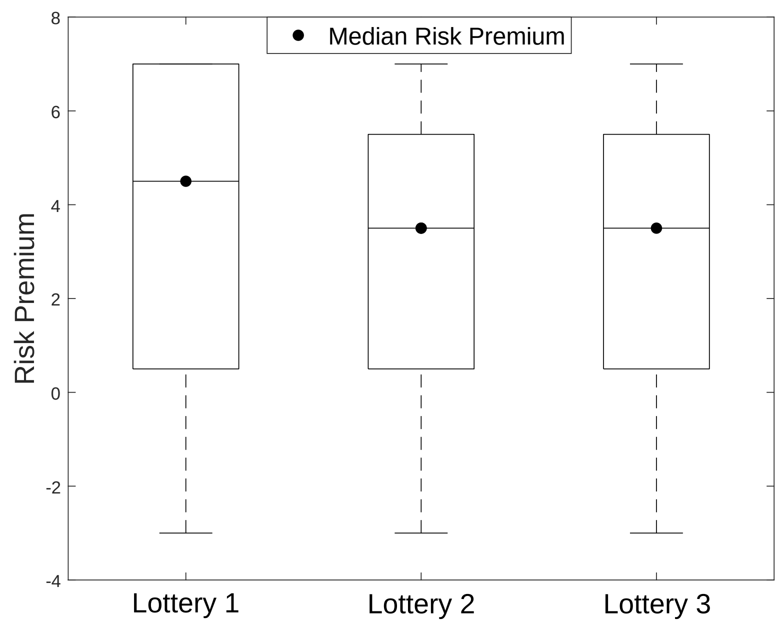

Figure A.2: Distributions of certainty equivalents for the three lotteries considered in Block 2.

Notes: We focus on the subset of observations from Block 2 where subjects shifted their preference between a fixed lottery and a certain amount at most once. The box plots summarize the distributions of risk premia for the three lotteries from Block 2. Lottery 1 pays \$0 or \$20 with equal chance. Lottery 2 pays \$5 or \$25 with equal chance. Lottery 3 pays \$10 or \$30 with equal chance.

are comparable across the three lotteries. For all three lotteries, the median risk premium is around \$4, and more than 75% of subjects have a positive risk premium. Moreover, for each lottery, we observe a significant level of heterogeneity in the risk premia.

A.2 Econometric Procedures

We provide details about the procedures that we follow to estimate economic models and train machine learning algorithms.

Economic Models

Cumulative Prospect Theory (CPT).

The value of lottery $p = (\$L, p_L; \$M, p_M; \$H, p_H)$ under CPT is

$$U_{CPT}(p) = \pi(p_H)u(H) + [\pi(p_H + p_M) - \pi(p_M)]\, u(M) + [1 - \pi(p_H + p_M)]\, u(L),$$

where $v(\cdot)$ is a utility function and $\pi(\cdot)$ is a probability weighting function. For estimation purposes, we consider the functional forms for utility and probability weighting functions proposed by Tversky and Kahneman (1992):

$$u(x) = x^{\alpha};$$
$$\pi(p) = \frac{p^{\gamma}}{[p^{\gamma} + (1-p)^{\gamma}]^{\frac{1}{\gamma}}}.$$

Within the empirical framework proposed by Hey and Orme (1994), a CPT DM chooses lottery p over lottery q if

$$U_{CPT}(p) - U_{CPT}(q) \geq \epsilon,$$

where ϵ is an error term normally distributed with a mean of zero and a variance of $\sigma > 0$.

We define $I = \{1, \ldots, N\}$ as a set of subjects in our experiment, $\triangle(X)$ as the set of lotteries over X, and by $\mathcal{D} \subseteq \triangle(X)^2$ as a subset of pairs of lotteries where the subjects express their preferences. We construct an index, $Choice_i$, for each subject i as follows: for each pair of lotteries $(p, q) \in \mathcal{D}$,

$$Choice_i(p,q) := \begin{cases} 2 & \text{if subject } i \text{chooses lottery } p \text{ over lottery } q \\ 1 & \text{otherwise.} \end{cases}$$

We estimate mixture models with three groups. Within each group c, we estimate the risk aversion coefficient α_c, the probability weighting function coefficient γ_c, and the variance of the error term σ_c. We denote by $f(Choice_i; \alpha_c, \gamma_c, \sigma_c)$ the likelihood function for subject i belonging to group c:

$$\prod_{(p,q)\in\mathcal{D}} \Big(\mathbb{1}\,(Choice_i(p,q)=2)\cdot \Pr(U_{CPT}(p)-U_{CPT}(q)\geq \epsilon \mid \alpha_c,\gamma_c,\sigma_c)$$
$$+\,\mathbb{1}\,(Choice_i(p,q)=1)\cdot \Pr(U_{CPT}(p)-U_{CPT}(q)< \epsilon \mid \alpha_c,\gamma_c,\sigma_c)\Big).$$

Let π_c represent the probability of a subject belonging to group type c. The log-likelihood of the finite mixture model is given by:

$$\sum_{i=1}^{N} \ln \sum_{c=1}^{C} \pi_c f(Core_i;\alpha_c,\gamma_c,\sigma_c),$$

where the first sum is over subjects and the second sum is over groups.

We estimate the utility functions, the parameters of the covariance matrices, and the probabilities of group membership through maximum likelihood estimation. We employ the Global Search algorithm in MATLAB to maximize the log-likelihood. To ensure that the algorithms converges to a global maximum, we employ a multi-start approach initiating multiple searches with 200 different starting points. We also evaluate the robustness of our estimates by estimating the model using an expectation-maximization algorithm (Dempster, Laird, and Rubin, 1977).

Expected Utility (EU). To estimate EU, we repeat the same procedure fixing the value of the probability weighting function to one for all the three groups.

Machine Learning Algorithms

Gradient Boosting Trees (GBT). We employ the LogitBoost algorithm using MATLAB's "fitcensemble" function, a specialized gradient boosting methodology tailored for binary classification. This method uses an ensemble of weak decision tree learners, optimizing the logistic loss to enhance classification accuracy. We allow the algorithm to use the following set of features: probabilities of the lotteries (p, q) and an indicator for each subject. The output of the algorithm is the probability of choosing p over q.

The LogitBoost operates in a stage-wise fashion. For each iteration, the algorithm focuses on the residuals or errors made by the present ensemble, these errors being a product of the logistic loss. Instead of directly approximating the class labels, LogitBoost models the posterior probabilities of the classes. At every step, a new decision tree is trained to fit the current residuals. This tree, once trained, is amalgamated into the ensemble. To optimize performance, we have implemented hyperparameter optimization, adjusting the number of learning cycles, learning rate, and minimum leaf size for the decision trees. The optimization uses a 10-fold cross-validation, parallel computation, and is bounded to a maximum of 150 evaluations with a 4-hour time constraint.

Neural Networks (NN). We employ a neural network classifier using MATLAB's “fitcnet" function. We allow the NN to use the following sets of input variables: probabilities of the lotteries, prizes of the lotteries, observed choices and an indicator for each subject. The outputs of the algorithm are choice probabilities.

The NN follows a structured methodology. Initially, the input data is processed with standardization, ensuring all features have a mean of zero and a variance of one. This pre-processing step aids in stabilizing and speeding up the network's convergence during training. Next, we implement a 10-fold cross-validation strategy to optimize the following hyperparameters of the NN: activation functions, regularization strength and the size of the hidden layers. To accelerate the training, parallel computation is leveraged, and the optimization is constrained to a maximum of 150 evaluations with a 4-hour time limit.

A.3 Estimation Results

We report the estimation results arising from the EU core analysis, as well as those arising from EU and CPT.

EU Core Analysis

Table A.1 presents the estimation results from the EU core analysis, which were obtained using a mixture model with three groups. The estimates are presented with bootstrapped standard errors in parentheses. The row labeled with π presents the estimated sizes of each group. The row labeled with σ shows the standard deviation of the error terms for each utility. The row labeled with ρ indicates the correlation coefficients among the error terms.

EU and CPT

Table A.2 reports the estimation results for EU and CPT derived from mixture models with three groups. For EU, we rank groups of subjects based on risk aversion, which depends solely on the parameter that shapes the curvature of the utility functions (α). For CPT, we rank groups of subjects in terms of risk aversion

Table A.1: EU core mixture model: Estimation results.

Groups	Dataset: CR-tasks and R-tasks					
	High RA High EU		Middle RA Low EU		Low RA Middle EU	
π	0.2798 (0.1355)		0.4834 (0.1128)		0.2368 (0.0973)	
	Utility 1	**Utility 2**	**Utility 1**	**Utility 2**	**Utility 1**	**Utility 2**
\$5	0.7074 (0.1944)	0.6686 (0.1863)	0.5805 (0.1827)	0.4002 (0.1206)	0.1389 (0.0860)	0.3722 (0.1174)
\$10	0.9758 (0.1705)	0.9184 (0.1801)	0.6105 (0.1708)	0.8286 (0.1507)	0.4424 (0.0975)	0.3722 (0.1039)
\$15	0.9758 (0.1684)	0.9467 (0.1807)	0.7466 (0.1553)	0.8461 (0.1343)	0.5455 (0.1181)	0.5997 (0.1158)
\$20	1.0000 (0.1661)	1.0000 (0.1681)	0.7918 (0.1710)	0.9800 (0.1435)	0.7475 (0.1564)	0.6741 (0.1276)
\$25	1.0000 (0.1355)	1.0000 (0.1314)	1.0000 (0.1722)	0.9800 (0.1248)	0.8756 (0.1268)	0.9663 (0.1552)
σ	0.2759 (0.3551)	0.0294 (0.4093)	0.3052 (0.3016)	0.1632 (0.5821)	0.1106 (0.1978)	0.2502 (0.5068)
ρ	0.0034 (0.1364)		0.0000 (0.1110)		0.0090 (0.1261)	

and adherence to EU. We use the predicted proportion of choices in CR-tasks and R-tasks where the safest available lottery was selected as a proxy for risk aversion.[1] Adherence to EU can be directly inferred from the probability weighting function parameter, γ.

A.4 Out-of-Sample Analysis: Additional Results

We develop a more comprehensive probabilistic and deterministic evaluation of all the predictive approaches considered in this paper. We first focus on the ability of economic models and machine learning algorithms to predict choice patterns, thus extending the analysis in Section 1.3. Next, we extend the EU core analysis in Section 1.5.

Predict Choice Patterns. There are four possible choice patterns in CR-tasks and R-

[1] Risk aversion under CPT is influenced by both utility and probability weighting functions. Hence, it is determined by the interaction between α and γ.

Table A.2: EU and CPT mixture models: Estimation results.

	Dataset: CR-tasks and R-tasks					
Model	**Expected Utility (EU)**			**Cumulative prospect theory (CPT)**		
Groups	High RA	Middle RA	Low RA	High RA Low EU	Middle RA Middle EU	Low RA High EU
α	0.0269 (0.2231)	0.4450 (0.3355)	0.9364 (0.2595)	1.0051 (0.1474)	0.2004 (0.2113)	0.7461 (0.1647)
γ				0.1216 (0.2678)	0.9800 (0.2092)	0.9463 (0.1975)
σ	0.1419 (0.1096)	0.5396 (0.5513)	2.0863 (1.1591)	6.0000 (0.7305)	0.2010 (0.0582)	1.3625 (0.3348)
π	0.4935 (0.2209)	0.3463 (0.2183)	0.1602 (0.1410)	0.2401 (0.0544)	0.4166 (0.1377)	0.3433 (0.1085)

Notes: Estimates are presented with bootstrapped standard errors in parentheses. In this table, α denotes the risk-aversion coefficient, γ denotes the probability weighting function coefficient, σ denotes the standard deviation of the error term, and π denotes the estimated size of each group.

tasks. We can thus define a model f as a function that links each lottery pair (δ_M, r) and subject i, with a vector of characteristics X_i, to a vector $f(\delta_M, r; X_i) \in \mathbb{R}^4$. Each component of this vector indicates the estimated probability under model f of one of the four choice patterns. Denoting by $P_i(\delta_M, r) \in \mathbb{R}^4$ the degenerate probability distribution that indicates which choice pattern, associated with the pair of lotteries (δ_M, r) and subject i, was actually observed, we define the loss of a model as follows:

$$L_p(f) := \sum_{(\delta_M, r)\in\mathcal{D}} \sum_{i=1}^{500} l\left(P_i(\delta_M, r), f(\delta_M, r; X_i)\right),$$

where $\mathcal{D}$ is the set of the 25 pair of lotteries (δ_M, r) in our experiment, and $l\colon \mathbb{R}^4 \times \mathbb{R}^4 \to \mathbb{R}$ is a loss function, which we assume to be the Euclidean distance.

Moreover, we analyze the ability of different models to provide accurate deterministic predictions. In particular, we denote by $\tilde{f}(\delta_M, r) \in \mathbb{R}^4$ the vector whose all components are equal to zero, except from the one associated with the choice pattern having the highest predicted probability. To measure deterministic accuracy, we define the deterministic loss of model f as the fraction of choice patterns that are misclassified:

$$L_d(f) := \frac{1}{|\mathcal{D}| \times 500} \sum_{(\delta_M, r)\in\mathcal{D}} \sum_{i=1}^{500} \mathbb{1}\left(P_i(\delta_M, r) \neq f(\delta_M, r)\right),$$

where $|\mathcal{D}|$ denotes the cardinality of the set $\mathcal{D}$.

Table A.3: Choice patterns analysis: deterministic and probabilistic evaluations.

Exercise	**Prediction**	**Loss**	**Models**			
			EU	CPT	GBT	NN
Combined Within Task	Choice Pattern	Det.	6.1869 (1.1196)	6.1136 (1.1023)	**4.3690** **(0.7961)**	6.5877 (1.2214)
		Prob.	3.3573 (0.4031)	3.3824 (0.4277)	**2.5535** **(0.3010)**	3.6912 (0.4520)
Combined Across Task	Choice Pattern	Det.	**1.9933** **(0.0160)**	2.0041 (0.0939)	2.3563 (0.2575)	2.5958 (0.5927)
		Prob.	**1.6767** **(0.1415)**	1.7309 (0.1576)	1.9381 (0.2201)	2.1545 (0.3754)

Notes: Normalized loss functions with standard deviation in parentheses. “Det." stands for “Deterministic”, while “Prob." stands for “Probabilistic”. The smallest in-sample loss was obtained with the GBT algorithm in all out-of-sample exercises.

In general, lower values for both probabilistic and deterministic losses indicate higher probabilistic and deterministic accuracy of a model. However, interpreting the absolute magnitude of these losses can be challenging. To facilitate interpretation, we introduce a normalized loss measure. For a given loss L, its normalized version is defined as:[2]

$$\hat{L}(f) := \frac{L(f)}{L^*},$$

where f is any generic model, and L^* is the lowest possible loss that can be achieved by training machine learning algorithms directly on the test data. Therefore, the normalized loss quantifies how many times greater the loss of a model trained on the training data is compared to the lowest possible loss that can be achieved by training a model directly on the test data.

Table A.3 summarizes the average normalized losses of the various methods for the out-of-sample exercises within and across tasks conducted in Section 1.3 and Section 1.3. The deterministic and probabilistic assessments of all predictive approaches yield the same result. The GBT outperforms EU and CPT in cross-validation exercises within CR-tasks and R-tasks. At the same time, EU achieves the smallest deterministic and probabilistic normalized losses in out-of-sample exercises across tasks. Furthermore, the values of the normalized losses inform us about the magnitude of a model’s loss compared to the smallest achievable loss. Overall, all the normalized losses are significantly greater than one, indicating a substantial cost in

[2] See Fudenberg et al. (2022).

Table A.4: Adherence to EU: deterministic and probabilistic evaluations.

Exercise	Prediction	Loss	Models				
			EU	CPT	GBT	NN	EU core
Combined Within Task	Index *Core*	Det.	6.3195 (1.1297)	6.1190 (1.1742)	**4.1648** **(0.7809)**	6.4227 (1.2513)	4.6769 (0.9830)
		Prob.	3.0751 (0.4194)	3.1845 (0.4557)	**2.4295** **(0.2879)**	3.4511 (0.4736)	2.9069 (0.3836)
Combined Across Task	Index *Core*	Det.	2.0843 (0.0316)	2.0778 (0.0671)	2.2207 (0.4566)	2.2038 (0.6289)	**1.9374** **(0.0401)**
		Prob.	**1.5927** **(0.1874)**	1.6656 (0.1966)	1.8322 (0.3495)	1.8215 (0.4960)	1.6006 (0.1505)

Notes: Normalized loss functions with standard deviation in parentheses. The smallest in-sample loss was obtained with the GBT algorithm in all out-of-sample exercises.

predictive accuracy when transitioning from in-sample to out-of-sample predictions.

EU Core Analysis. When the objective shifts from predicting choice patterns to forecasting the index $Core$, we can define a model as a function that associates each lottery pair (δ_M, r) and subject i, with a vector of characteristics X_i, to a vector $f(\delta_M, r; X_i) \in \mathbb{R}^3$. The j-th component of this vector represents the estimated probability under model f that the index $Core(\delta_M, r)$ assumes the value of $j \in 1, 2, 3$. For each model, we derive deterministic and probabilistic normalized loss functions, replicating the analysis previously described for choice patterns.

Table A.4 summarizes the average normalized losses of the various methods for the out-of-sample exercises within and across tasks conducted in Section 1.5. The GBT outperforms all other predictive approaches in cross-validation exercises within CR-tasks and R-tasks. The probabilistic performances of both EU and the EU core approach are almost identical, with EU achieving a slightly smaller normalized loss. All other predictive approaches have significantly worse probabilistic performance.

A.5 Instructions

General Instructions. You will receive \$4 if you complete the entire study. We anticipate that the study will take about 20 minutes, on average. In addition to this payment, one out of every five participants will be randomly selected to receive a bonus payment. The smallest possible bonus payment is \$0 and the largest possible bonus payment is \$30. You will be informed of how your decisions will influence your bonus payment if you were to be randomly selected.

Description of the Experiment. In this study, we will ask you questions about *lotteries*. A lottery specifies different payments you may receive with different chances.

For example, one lottery might be the following:

Lottery

15% chance of **$17**

25% chance of **$9**

60% chance of **$2**

You can think of this lottery in the following way:

- In 15 out of 100 chances (15% chance) the lottery pays $17.
- In 25 out of 100 chances (25% chance) the lottery pays $9.
- In 60 out of 100 chances (60% chance) the lottery pays $2.

We may allow you to play a lottery and receive the outcome of the lottery as bonus payment. The outcome of the lottery will be determined by the computer using the chances specified. You will learn more about your bonus payment in the following instruction pages.

Check for Understanding. Before we proceed, here is a question to test your understanding. Consider the lottery below:

In how many out of 100 chances does this lottery pay $9?

o 15

o 25

o 60

Lottery

15% chance of **$17**

25% chance of **$9**

60% chance of **$2**

- o 90
- o None of the above

[Subjects are required to provide the correct answer in order to proceed. When they select a wrong answer, we show them the following error message: "The lottery pays $9 with a 25% chance. The correct answer is 25. Please revise your answer."]

Blocks of the Experiment. We will ask you to make choices over lotteries in two different blocks. If you are selected to receive a bonus payment, then we will randomly pick one of these blocks. We will describe within each block how your bonus would be determined if that block were randomly selected. Please click to learn about Block 1.

Block 1. In Block 1, we will show you two lotteries and will ask you to choose between the following two answer choices:

1. I prefer Lottery A
2. I prefer Lottery B

Block 1: Bonus Payment. If Block 1 is selected to determine your bonus payment, how would we pay you? We will randomly select one task from Block 1, and we will let you play the lottery you preferred. The lottery's outcome will be your bonus payment. We ask you to complete a brief training session to check your understanding of the tasks in Block 1.

Please proceed to start the training session!

Training task 1 of 2 - Block 1

We ask you to complete the following training task **assuming that**:

- You prefer **Lottery A over Lottery B**.

Lottery A	**Lottery B**
30% chance of **$40**	**40%** chance of **$30**
70% chance of **$5**	**60%** chance of **$10**

Which lottery do you **prefer**?

I prefer Lottery A ○ I prefer Lottery B ○

Training task 2 of 2 - Block 1

Suppose that this task is selected for bonus payment and your answer is:

Lottery A	**Lottery B**
30% chance of **$40**	**40%** chance of **$30**
70% chance of **$5**	**60%** chance of **$10**

Which lottery do you **prefer**?

I prefer Lottery A ◉ I prefer Lottery B ○

Please answer the following comprehension question:

How do we determine your bonus payment in this example?

- ○ We let you play Lottery A to determine your bonus payment.
- ○ We let you play Lottery B to determine your bonus payment.
- ○ We let you play Lottery A to determine your bonus payment with a 50 in 100 chance, or Lottery B with a 50 in 100 chance.

[Subjects are required to provide the correct answers in both training tasks in order to proceed. When they select a wrong answer in the first training task, we show them the following error message: "We are asking you to answer the question assuming that you prefer Lottery A. Please revise your answer." When they select a wrong answer in the second training task, we show them the following error message: "Given that you preferred Lottery A in this example, we would let you play Lottery A to determine your payment. Please revise your answer."]

Begin Block 1. Thank you for completing the training session. There will be 102 tasks in Block 1. Please answer all the questions thoughtfully to the best of your ability. Remember that there are no right or wrong answers. We are only interested in studying your preferences. Please proceed to start Block 1!

[Subjects complete CR-tasks, R-tasks and FOSD-tasks presented to them in a randomized order.]

Block 2. In Block 2, we will show you two lotteries and will ask you to choose between the following two answer choices:

1. I prefer Lottery A
2. I prefer Lottery B

Block 2: Bonus Payment. If Block 2 is selected to determine your bonus payment, how would we pay you? We will randomly select one task from Block 2, and we will let you play the lottery you preferred. The lottery's outcome will be your bonus payment.

Begin Block 2. There will be 33 tasks in Block 2. Please answer all the questions thoughtfully to the best of your ability. Remember that *there are no right or wrong answers*. We are only interested in studying your preferences.

Please proceed to start Block 2!

Block 2 - Part 1. In the following 11 tasks of Block 2:

- **Lottery A** pays a monetary amount with a **100%** chance that is different in every task. In task 1, the monetary amount paid for sure by Lottery A is **$3**. As you move from one task to the next one, the monetary amount paid for sure by Lottery A increases by $1. For instance, in task 2 is **$4**, in task 3 is **$5**, etc.
- **Lottery B** always pays **$20** with a **50%** chance, or **$0** with a **50%** chance.

[Subjects complete the 11 MPL1 tasks.]

Block 2 - Part 1. In the following 11 tasks of Block 2:

- **Lottery A** pays a monetary amount with a **100%** chance that is different in every task. In task 12, the monetary amount paid for sure by Lottery A is **$8**. As you move from one task to the next one, the monetary amount paid for sure by Lottery A increases by $1. For instance, in task 13 is **$9**, in task 14 is **$10**, etc.
- **Lottery B** always pays **$25** with a **50%** chance, or **$5** with a **50%** chance.

[Subjects complete the 11 MPL2 tasks.]

Block 2 - Part 1. In the following 11 tasks of Block 2:

- **Lottery A** pays a monetary amount with a **100%** chance that is different in every task. In task 23, the monetary amount paid for sure by Lottery A is **$13**. As you move from one task to the next one, the monetary amount paid for sure by Lottery A increases by $1. For instance, in task 24 is **$14**, in task 25 is **$15**, etc.
- **Lottery B** always pays **$30** with a **50%** chance, or **$10** with a **50%** chance.

[Subjects complete the 11 MPL3 tasks.]

Block 2 - Part 1. In the following 11 tasks of Block 2:

- **Lottery A** pays a monetary amount with a **100%** chance that is different in every task. In task 23, the monetary amount paid for sure by Lottery A is **$13**. As you move from one task to the next one, the monetary amount paid for sure by Lottery A increases by $1. For instance, in task 24 is **$14**, in task 25 is **$15**, etc.
- **Lottery B** always pays **$30** with a **50%** chance, or **$10** with a **50%** chance.

[Subjects complete the 11 MPL3 tasks.]

APPENDIX TO CHAPTER 2

B.1 FOSD Questions

Table B.1: Comparisons with dominated lotteries.

	Lottery A			Lottery B		
Questions	Pr($1)	Pr($7)	Pr($20)	Pr($1)	Pr($7)	Pr($20)
1	0.3	0.7	0	0.7	0.3	0
2	0.5	0.2	0.3	0.5	0.1	0.4
3	0	0.8	0.2	0	0.6	0.4
4	0.5	0.4	0.1	0.9	0	0.1
5	0.2	0.6	0.2	0.1	0.5	0.4
6	0.5	0.5	0	0.8	0.2	0

B.2 Demographic Summary

Table B.2: Demographics, overall sample.

Age	
18-24	54.3%
25-30	45.7%
Gender	
Male	50%
Female	50%
Education completed	
High school diploma	50.3%
Undergraduate degree (BA/BSc/other)	40%
Master degree (MA/MSc/MPhil/other)	7%
Doctorate degree (PhD/other)	1.7%
Prefer not to say	1%
Employment	
Unemployed (and job seeking)	22.3%
New job within the next month	0.7%
Part-time	11.3%
Full-time	22.3%
Other / Prefer not to say	43.3%

B.3 Confidence and Indecisiveness

Cerreia-Vioglio, Dillenberger, and Ortoleva (2015) classify one individual as more indecisive than another if his EU core is smaller, in the sense of set inclusion.

Definition 5. *Individual i is more indecisive than individual j if for all lotteries p, q:*

$$p \succsim_i^* q \Rightarrow p \succsim_j^* q.$$

We show that the term “indecisive" presumes a relationship between confidence and independence that we observe in the data: subjects that we classify as more indecisive tend to report lower confidence levels. To this end, Section B.3 details our approach to operationalize the notion of indecisiveness. Section B.3 builds a measure of decision confidence comparable across subjects, while Section B.3 uses this new measure to rank subjects in terms of decision confidence. Finally, Section B.3 shows that more indecisive subjects tend to be less confident.

Pairwise Analysis: Indecisiveness

We denote by t a generic treatment in the experiment and by C_t the set of comparisons in treatment t. For any comparison (s, r), we define the index $Core_i$ for each individual i to be equal to two if there is no evidence against $s \succsim_i^* r$, one if there is no evidence against $r \succsim_i^* s$, and zero otherwise. Using this index, we propose two different criteria to operationalize the notion of indecisiveness proposed by Cerreia-Vioglio, Dillenberger, and Ortoleva (2015) in the context of our experiment:

1. Criterion IND 1: subject i is more indecisive than subject j if for all comparisons $(s, r) \in C_t$,

   $$Core_i(s, r) > 0 \Rightarrow Core_i(s, r) = Core_j(s, r)$$

   and there exists a comparison $(\bar{s}, \bar{r}) \in C_t$ such that $Core_i(\bar{s}, \bar{r}) \neq Core_j(\bar{s}, \bar{r})$.

2. Criterion IND 2: subject i is more indecisive than subject j if

   $$|\{(s, r) \in C_t \; : \; Core_i(s, r) > 0\}| < |\{(s, r) \in C_t \; : \; Core_j(s, r) > 0\}|.$$

Criterion IND 1 amounts to classifying a subject as more indecisive than another only if there is no evidence against it. Because it is very demanding, we show that it does not allow the classification of most pairs of subjects in each treatment. On the contrary, to classify one subject as more indecisive than another according

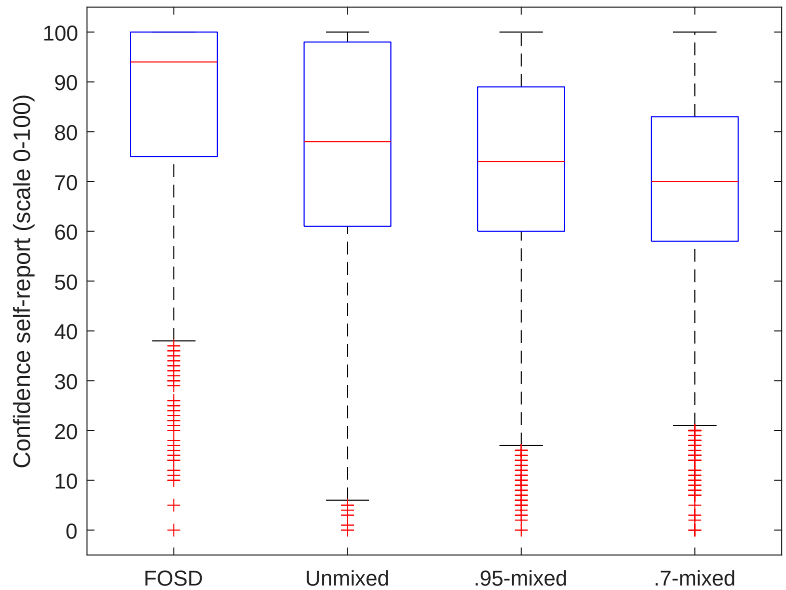

Figure B.1: Distribution of confidence self-reports for different categories of comparisons and all subjects in the experiment.

to criterion IND 2, it is enough to compare the number of comparisons with index $Core$ equal to one. The advantage of criterion IND 2 is that being less demanding, it allows the classification of the vast majority of pairs of subjects. The disadvantage is that it moves away from the original notion of indecisiveness captured in Definition 5. In Section B.3, we show that the qualitative results linking indecisiveness with decision confidence are stable across the two criteria.

Benchmarked Confidence Self-Reports

An important challenge in interpreting the self-reported confidence measures that we collect is that they are subjective and may have different meanings for different subjects. We now describe the approach that we adopt to convert the confidence reports of different subjects into the same unit of measure. To this end, we exploit subjects' confidence statements in comparisons involving stochastic dominance.

Let $FOSD$ be the set of the six comparisons in which one lottery first-order stochastically dominates the other described in Appendix B.1. Figure B.1 shows the distributions of confidence self-reports for different categories of comparisons.[1] A few things emerge from this figure. First, as expected, subjects tend to report higher confidence levels in $FOSD$ (leftmost box plot in Figure B.1) than in any

[1] The red plus signs denote outliers. We classify any observation that is more than 1.5 times the interquartile range away from the first quantile or the third quantile as an outlier.

other category of comparisons. However, while the median value of the confidence self-reports for $FOSD$ comparisons is close to 100, the interquartile range is equal to 25. In other words, there is significant heterogeneity in how subjects express high confidence using numbers. In what follows, we exploit this heterogeneity to benchmark the confidence self-reports in all other categories of comparisons.

We denote by $Conf_i(s,r)$ the confidence self-report of subject i in comparison (s,r) divided by 100. Moreover, we indicate by $\bar{c}_i$ the average confidence self-report in $FOSD$ comparisons for which subject i declares to prefer the dominant lottery. For all comparisons (s,r) not in $FOSD$ and all subjects i, we construct the following benchmarked index of confidence:

$$AdjConf_i(s,r) = \min\left\{\frac{Conf_i(s,r)}{\bar{c}_i}, 1\right\}.$$

Intuitively, $\bar{c}_i$ is an estimate of what number individual i reports to express extreme confidence. In the extreme case in which $Conf_i(s,r)$ is greater or equal to $\bar{c}_i$, we simply assign to the benchmarked index $AdjConf_i(s,r)$ the value of one.[2]

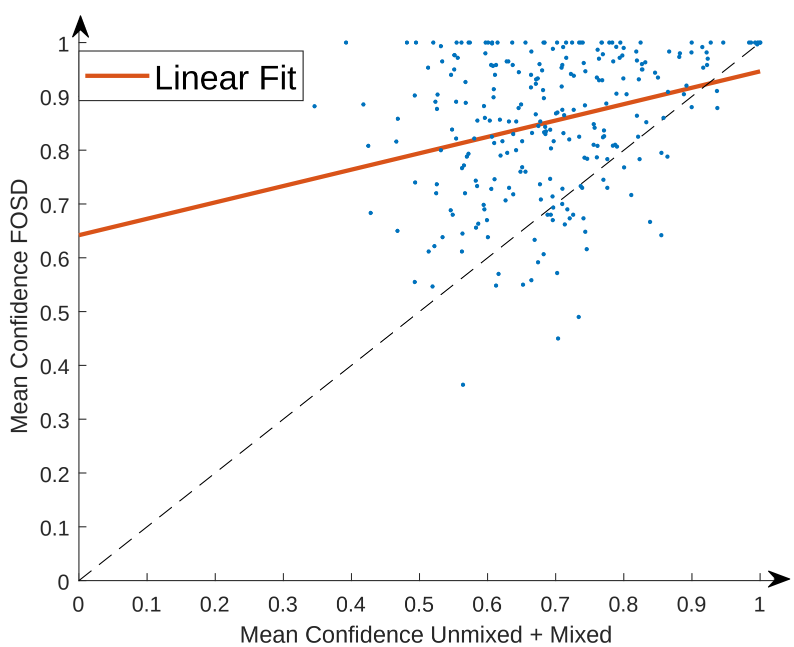

Figure B.2: Mean confidence in FOSD comparisons and in all other comparisons for each subject normalized to one.

The underlying assumption behind this benchmarked confidence index is that subjects who use lower numbers to express extreme confidence in FOSD comparisons

[2]This happens overall for the 24.72% of the self-reported confidence measures.

use lower numbers to express any confidence level. We now test the validity of this assumption. Figure B.2 shows the mean confidence levels in unmixed and mixed comparisons (x-axis) and in FOSD comparisons (y-axis) for each subject.[3] The orange line in Figure B.2 represents the best linear fit (in a least-squares sense). The positive correlation that we observe between these two quantities justifies our benchmarked index: subjects expressing extreme confidence in FOSD comparisons with lower numbers express any confidence with lower numbers.[4] Figure B.2 also allows us to evaluate the reliability of the confidence measures that we collect: for the 83.96% of subjects, the mean confidence in FOSD comparisons is higher than in all other comparisons.

Pairwise Analysis: Confidence

In analogy with the pairwise analysis of indecisiveness in Section B.3, we propose two different criteria to rank pairs of subjects in terms of decision confidence using the benchmarked index $AdjConf_i$:

1. Criterion CONF 1: subject i is less confident than subject j if for all comparisons $(s, r) \in C_t$,

$$AdjConf_i(s, r) \leq AdjConf_j(s, r),$$

   with strict inequality for some comparison $(\bar{s}, \bar{r}) \in C_t$.

2. Criterion CONF 2: subject i is less confident than subject j if for all $x \in [0, 1]$,

$$|\{(s, r) \in C_t \ : \ AdjConf_i(s, r) \leq x\}| \leq |\{(s, r) \in C_t \ : \ AdjConf_j(s, r) \leq x\}|,$$

   with strict inequality for at least one $\bar{x} \in (0, 1)$.

Using criterion CONF 1, we classify subject i as less confident than subject j if the benchmarked confidence self-report of i is lower than the benchmarked confidence self-report of j in all comparisons. As with criterion IND 1 for indecisiveness, criterion CONF 1 is very demanding and does not allow the classification of most pairs of subjects. For this reason, we also consider criterion CONF 2, which generalizes the requirement of criterion IND 2 to the continuous index $AdjConf$.

[3] For FOSD comparisons, we exclude confidence self-reports in pairs of lotteries where the dominated lottery was preferred. Given our restriction on the sample, this can happen at most once for each subject.

[4] Correlation coefficient: 0.303.

Table B.3: Percentage of pairs of subjects in which the more indecisive subject is less confident. Number of classifiable pairs in parenthesis.

Worst			Bad			WorstBest		
	IND 1	IND 2		IND 1	IND 2		IND 1	IND 2
CONF 1	75.00% (4)	66.67% (243)	CONF 1	65.22% (23)	65.27% (334)	CONF 1	54.00% (100)	67.30% (523)
CONF 2	68.85% (61)	59.44% (2,125)	CONF 2	86.49% (111)	60.76% (1,909)	CONF 2	64.22% (218)	56.66% (2,344)

According to criterion CONF 2, subject i is less confident than subject j whenever the empirical cumulative distribution function of $AdjConf_j$ first-order stochastically dominates the empirical cumulative distribution function of $AdjConf_i$.

Results

We now explore the relationship between EU core and decision confidence by using the notion of indecisiveness introduced by Cerreia-Vioglio, Dillenberger, and Ortoleva (2015). In Section B.3 and Section B.3 we propose two criteria to rank subjects in terms of indecisiveness (IND 1 and IND 2) and confidence (CONF 1 and CONF 2). Considering all four possible combinations of these approaches, Table B.3 reports the percentage of pairs of subjects, among those that can be classified according to both criteria, for which the more indecisive subject is less confident. The number in parenthesis below each percentage in Table B.3 represents the total number of classifiable pairs. For instance, in the Worst treatment, 243 pairs of subjects can be classified according to criterion CONF 1 and criterion IND 2. In the 66.67% of these pairs, the more indecisive subject is less confident.

Overall, for all the possible combinations of criteria in all three treatments, the more indecisive subject is less confident in more than half of the classifiable pairs. We see this result as an empirical justification of the term “indecisive" used in Cerreia-Vioglio, Dillenberger, and Ortoleva (2015) to rank subjects’ EU cores.

B.4 Indecisiveness and Risk Aversion

Our analysis shows that hypothesis EU-CORE is less likely to fail when the safer lottery is chosen. Given that the extent to which subjects prefer safer over riskier lotteries positively correlates with their degree of risk aversion, these estimates suggest a negative relationship between indecisiveness and risk aversion. At the same time, Cerreia-Vioglio, Dillenberger, and Ortoleva (2015) show that the opposite relationship holds for preferences that satisfy the NCI axiom. If a subject is more indecisive than another, he must be more risk averse. We now propose a direct test

Table B.4: Percentage of pairs of subjects in which the more indecisive subject is more risk averse. Number of classifiable pairs in parenthesis.

Worst	IND 1	IND 2	**Bad**	IND 1	IND 2	**WorstBest**	IND 1	IND 2
RA 1	13.64% (110)	48.85% (4,274)	RA 1	10.71% (168)	27.33% (3,154)	RA 1	10.11% (356)	31.03% (3,938)
RA 2	12.61% (111)	47.79% (4,340)	RA 2	10.71% (168)	28.38% (3,213)	RA 2	12.40% (363)	33.44% (4,007)

for this prediction.

We adopt two criteria to rank subjects' risk attitudes. The first criterion (RA 1) consists in classifying one subject as more risk averse than another if he chooses the safer over the riskier lottery more often in the experiment. In the second criterion (RA 2), we instead normalize the utility of \$1 to zero, the utility of \$20 to one, and following Hey and Orme (1994), we estimate the utility of \$7 and of the variance of the error term which we assume to be normally distributed with zero mean. The higher the estimated utility, the more risk averse a subject is.

Table B.4 reports the percentage of pairs of subjects, among those that can be classified according to both indecisiveness and risk aversion, for which the more indecisive subject is more risk averse. Below each percentage in Table B.4 is reported the total number of classifiable pairs. For all the possible combinations of criteria in all treatments, more indecisive subjects are less risk averse. In what follows, we prove that preferences satisfying the PCI axiom may explain the correlation between indecisiveness and risk aversion that we observe in Table B.4.

Definition 6. *Let $\succsim$ be a binary relation over the set of lotteries $\triangle(X)$. We say that $\succsim$ satisfies:*[5]

- *Completeness if for each $p, q \in \triangle(X)$, either $p \succsim q$ or $q \succsim p$.*
- *Transitivity if for each $p, q, r \in \triangle(X)$, $p \succsim q$ and $q \succsim r$ imply $p \succsim r$.*
- *Continuity if for each $p \in \triangle(X)$, the sets $\{q \in \triangle(X) : q \succsim p\}$ and $\{q \in \triangle(X) : p \succsim q\}$ are closed.*
- *Strict first-order stochastic dominance if for each $p \in \triangle(X)$, if p first-order stochastically dominates q, then $p \succ q$.*

[5]We denote by $\sim$ and $\succ$ the symmetric and the asymmetric parts of $\succsim$. The set X represents any compact set of monetary prizes.

- *Betweenness if for each $p, q \in \triangle(X)$ and $\lambda \in [0,1]$,*

$$p \sim q \Rightarrow p \sim \lambda p + (1-\lambda)q \sim q.$$

We refer to binary relations that satisfy the five axioms in Definition 6 as betweenness preferences. Next, we introduce the PCI axiom.

Definition 7. *$\succsim$ satisfies the PCI axiom if for each $p, q \in \triangle(X)$, $x \in X$ and $\lambda \in [0,1]$,*

$$\delta_x \succsim p \Rightarrow \lambda\delta_x + (1-\lambda)q \succsim \lambda p + (1-\lambda)q.$$

Let $\succsim_1$ and $\succsim_2$ be the preferences over $\triangle(X)$ for individuals 1 and 2. We use the following standard approach to compare risk attitudes.

Definition 8. *Individual 1 is more risk averse than individual 2 if for each $p \in \triangle(X)$ and $x \in X$,*

$$p \succsim_1 \delta_x \Rightarrow p \succsim_2 \delta_x.$$

We are now ready to show that betweenness preferences satisfying the PCI axiom rationalize what we observe in our experiment: more indecisive subjects tend to be less risk averse. The result follows combining the representation result in Cerreia-Vioglio, Dillenberger, and Ortoleva (2020, Remark 1) with the proof technique used in Cerreia-Vioglio, Dillenberger, and Ortoleva (2015, Proposition 2).

Corollary 6. *Suppose that $\succsim_1$ and $\succsim_2$ are betweenness preferences that satisfy the PCI axiom. If individual 1 is more indecisive than individual 2, then individual 1 is less risk averse than individual 2.*

Proof. We denote by $\mathbb{E}[p,v]$ and $c(p,v)$ the EU and the certainty equivalent of lottery p with utility function v, respectively. By Cerreia-Vioglio, Dillenberger, and Ortoleva (2020), for each individual $i \in \{1,2\}$, there exists a set of utility functions $\mathcal{W}_i$ such that

$$p \succsim_i^* q \Leftrightarrow \mathbb{E}[p,v] \geq \mathbb{E}[q,v] \text{ for all } v \in \mathcal{W}_i,$$

and the functional $u_i \colon \triangle(X) \to \mathbb{R}$ defined by

$$u_i(p) = \sup_{v \in \mathcal{W}_i} c(p,v) \text{ for all } p \in \triangle(X),$$

is a continuous utility representation of $\succsim_i$. Given that individual 1 is more indecisive than individual 2,

$$p \succsim_1^* q \Rightarrow p \succsim_2^* q \Rightarrow p \succsim_2 q.$$

By Cerreia-Vioglio (2009, Proposition 22), $\mathcal{W}_2 \subseteq \mathcal{W}_1$. Therefore, for any risky lottery p,

$$u_1(p) = \sup_{v \in \mathcal{W}_1} c(p, v) \geq \sup_{v \in \mathcal{W}_2} c(p, v) = u_2(p).$$

Consequently, individual 1 is less risk averse than individual 2. $\square$

B.5 Response Times

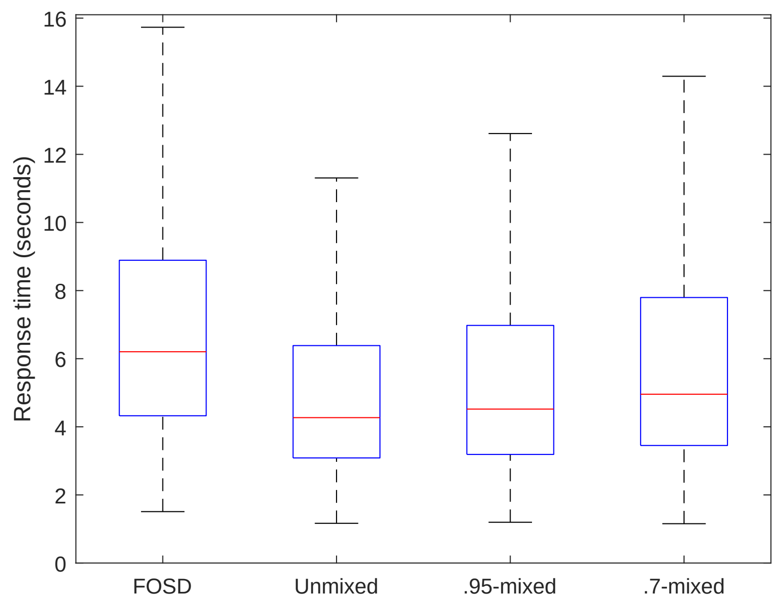

Figure B.3: Distribution of response times for all subjects in the experiment.

Figure B.3 shows the distribution of these response times for different categories of comparisons.[6] Contrary to the confidence self-reports, response times tend to be lower in unmixed than in 0.95-mixed and 0.7-mixed comparisons. Indeed, we observe a negative correlation between confidence self-reports and response times for these three categories of comparisons.[7] At the same time, analogously to confidence self-reports, response times tend to be higher in FOSD comparison than in any other category of comparisons. We suspect that the high response times observed in FOSD comparisons may result from a "too good to be true" effect. Subjects may lose time in double-checking their understanding of questions in FOSD comparisons, given the lack of trade-offs between the two lotteries.

[6] We classify as outlier any observation that is more than 1.5 times the interquartile range away from the first quantile or the third quantile. Overall, there are 1,322 outliers. Given that many outliers correspond to particularly high numbers, representing subjects that most likely took a break during the experiment, we do not report them in Figure B.3 for clarity.

[7] Correlation coefficient: −0.206. P-value less than 0.001.

B.6 Instructions

General Instructions. The study consists of questions about lottery tickets that pay \$1, \$7 or \$20 with some fixed probabilities. Let us highlight from the start that there are no right or wrong answers. We are only interested in studying your preferences. Here is an example of a pair of lottery tickets:

Lottery Ticket A	**Lottery Ticket B**
0% chance of \$1	40% chance of \$1
100% chance of \$7	0% chance of \$7
0% chance of \$20	60% chance of \$20

- Lottery ticket **A** involves no chance at all: it pays \$7 for sure.
- Lottery ticket **B** pays \$20 with probability 60%, or \$1 with probability 40%.

During the experiment, you will encounter 74 pairs of lottery tickets. For each pair, you will answer to two questions.

Notice that after you select an answer to a question and click on Next, you **will not be able to modify it**.

We will now show you these two questions.

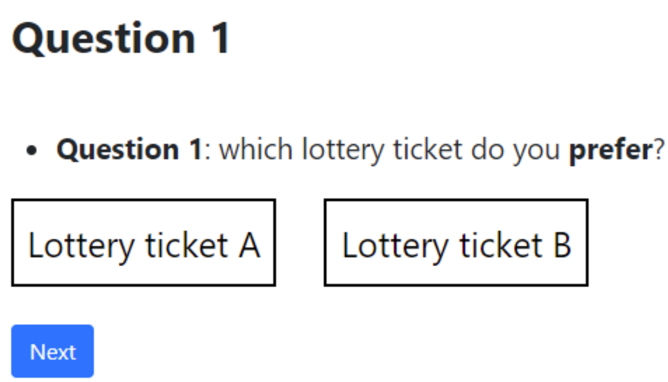

Training session. To familiarize yourself with the setup of the experiment, you will complete a brief training session. The training session consists of **five** pairs of lottery tickets. For each pair of lottery tickets, you will answer the two questions that we described to you. The answers that you give in this training session **do not affect** your monetary compensation. We will describe the details of your monetary

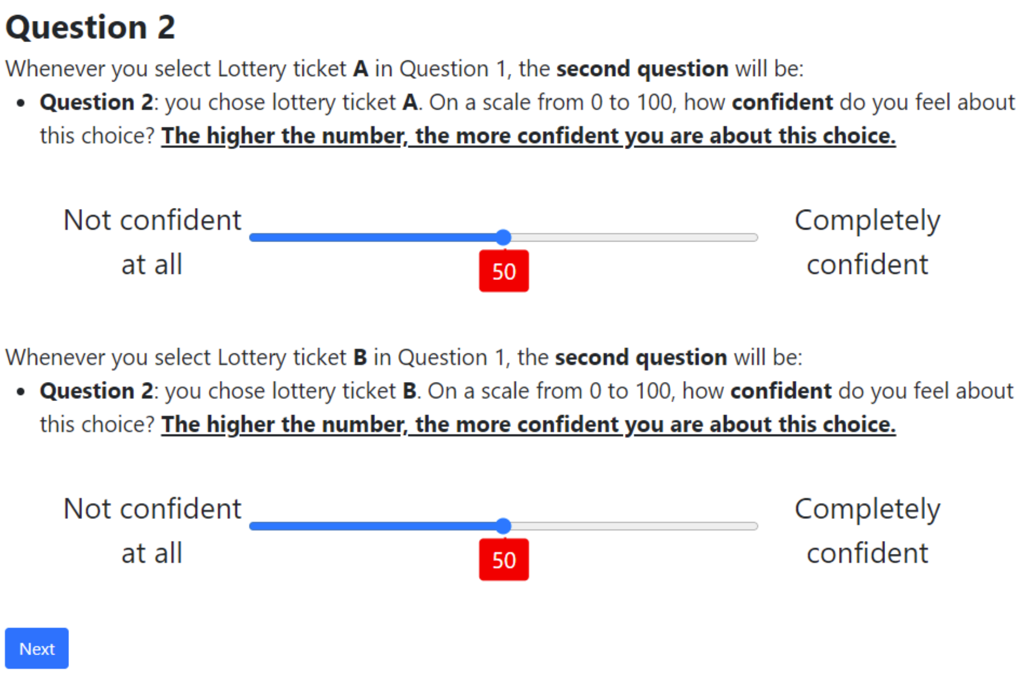

compensation at the end of your training session.

[Subject answers the five comparisons. For each comparison, the decision screen looks like in Figure 2.2.]

Possible rewards. Now that you have familiarized yourself with the questions of the study, you will learn about the details of your compensation. You will receive a **participation fee** of **$4.75** for completing all the questions in the experiment. Moreover, you may also receive a **bonus payment**:

- At the end of the experiment, the computer will randomly select a number between 1 and 10. Each number has an equal probability (**10%**) of being selected.
- If the randomly selected number is 1, you will receive a bonus payment.
- If you are chosen to receive a bonus payment, the computer will randomly pick one **of the 74 pairs of lotteries** that you encountered in the experiment.
- Then, you will be able to **play the lottery ticket from the selected pair that you declared to prefer in Question 1** (which lottery ticket do you prefer?).

That is, the computer will use the probabilities specified in the lottery ticket to select a monetary prize ($1, $7 or $20).

- Your bonus payment will be **the monetary prize that the computer will select**. You will receive your bonus payment together with the participation fee after we review your submission.

Example: Suppose that you are randomly selected for a bonus payment and the randomly picked pair of lottery tickets is:

Lottery Ticket A	**Lottery Ticket B**
0% chance of $1	40% chance of $1
100% chance of $7	0% chance of $7
0% chance of $20	60% chance of $20

- If you chose Lottery ticket **A** in Question 1, you get additional $7 for sure.
- If you chose Lottery ticket **B** in Question 1, you have a 60% chance of getting additional $20 and a 40% chance of getting additional $1.

Begin the experiment. Congratulations, you are now ready to participate in the experiment. If anything is unclear, please let us know through the Prolific anonymized internal messaging service. Otherwise, please click next to begin the experiment.

References

Cerreia-Vioglio, Simone (2009). *Maxmin expected utility on a subjective state space: Convex preferences under risk*. Tech. rep. Mimeo, Bocconi University.

Cerreia-Vioglio, Simone, David Dillenberger, and Pietro Ortoleva (2015). "Cautious expected utility and the certainty effect". In: *Econometrica* 83.2, pp. 693–728.

– (2020). "An explicit representation for disappointment aversion and other betweenness preferences". In: *Theoretical Economics* 15.4, pp. 1509–1546.

Hey, John D and Chris Orme (1994). "Investigating generalizations of expected utility theory using experimental data". In: *Econometrica: Journal of the Econometric Society*, pp. 1291–1326.

Appendix C

APPENDIX TO CHAPTER 3

C.1 Proofs

This appendix contains the proofs of the results stated in Chapter 3.

Proof of Proposition 1

Step 1. *If* $\bigcap_{v \in M_{\alpha^*}} SUCS_v(\alpha^*) = \emptyset$*, then* α^* *is optimal.*

Proof of Step 1. Consider an action α^* such that $\bigcap_{v \in M_{\alpha^*}} SUCS_v(\alpha^*) = \emptyset$. This implies that for all $\alpha \in \mathcal{A}$, there exists $v \in M_{\alpha^*}$ such that

$$u(\alpha^*) = \mathrm{U}\left[\mathbb{E}(\alpha^*, v), v\right] \geq \mathrm{U}\left[\mathbb{E}(\alpha, v), v\right] \geq u(\alpha).$$

Therefore, α^* is optimal.

Step 2. *If* α^* *is optimal, then* $\bigcap_{v \in M_{\alpha^*}} SUCS_v(\alpha^*) = \emptyset$*.*

Proof of Step 2. We show that if there exists an action $\alpha \in \mathcal{A}$ such that $\alpha \succ_v \alpha^*$ for all $v \in M_{\alpha^*}$, then α^* is not optimal. Define a new mixed action $\hat{\alpha}_\lambda$ parametrized by $\lambda \in (0, 1)$ such that for all $a \in A$

$$\hat{\alpha}_\lambda(a) = \lambda\alpha(a) + (1 - \lambda)\alpha^*(a).$$

We now show that there exists a value $\bar{\lambda} \in (0, 1]$ such that for all $\lambda \in (0, \bar{\lambda})$, $u(\hat{\alpha}_\lambda) > u(\alpha^*)$. To this end, consider the function $\Psi\colon [0, 1] \times \mathcal{W} \to \mathbb{R}$ such that $\Psi(\lambda, v) = \mathrm{U}\left[\mathbb{E}(\hat{\alpha}_\lambda, v), v\right] - u(\alpha^*)$, for all $\lambda \in [0, 1]$ and $v \in \mathcal{W}$. As an intermediate step, we prove that for each $v \in \mathcal{W}$, there exists a value $\bar{\lambda}_v \in (0, 1]$ such that for all $\lambda \in (0, \bar{\lambda}_v)$, $\Psi(\lambda, v) > 0$. Take a utility $v \in \mathcal{W}$. There are two possibilities:

1. If $\alpha \succsim_v \alpha^*$, then for all $\lambda \in (0, 1]$

$$\begin{aligned}
\Psi(\lambda, v) &= \mathrm{U}\left[\mathbb{E}(\hat{\alpha}_\lambda, v), v\right] - u(\alpha^*) \\
&= \mathrm{U}\left[\lambda\mathbb{E}(\alpha, v) + (1 - \lambda)\mathbb{E}(\alpha^*, v), v\right] - u(\alpha^*) \\
&\geq \mathrm{U}\left[\lambda\mathbb{E}(\alpha^*, v) + (1 - \lambda)\mathbb{E}(\alpha^*, v), v\right] - u(\alpha^*) \\
&= \mathrm{U}\left[\mathbb{E}(\alpha^*, v), v\right] - u(\alpha^*) \geq 0,
\end{aligned}$$

where at least one of the two weak inequalities holds strict. If $\alpha \succ_v \alpha^*$, the first inequality is strict because the function $\mathrm{U}[\cdot, v]$ is strictly increasing in the first argument. If instead $\alpha \sim_v \alpha^*$, the last inequality is strict because $v \notin M_{\alpha^*}$. For all such v, we let $\bar{\lambda}_v = 1$.

2. If $\alpha^* \succ_v \alpha$, then $\Psi(\cdot, v)$ is strictly decreasing in the first argument because $\mathrm{U}[\cdot, v]$ is strictly increasing in the first argument. $\Psi(\cdot, v)$ is continuous in the first argument because $\mathrm{U}[\cdot, v]$ is continuous in the first argument. If $\Psi(1, v) \geq 0$, the result follows immediately by taking $\bar{\lambda}_v = 1$. Suppose that $\Psi(1, v) < 0$. Notice that $\Psi(0, v) > 0$ because $\alpha^* \succ_v \alpha$ implies that $v \notin M_{\alpha^*}$. Therefore, by the Intermediate Value Theorem, there exists $\bar{\lambda}_v \in (0, 1)$ such that $\Psi(\bar{\lambda}_v, v) = 0$. By $\Psi(\cdot, v)$ strictly decreasing, $\Psi(\lambda, v) > 0$ for all $\lambda \in (0, \bar{\lambda}_v)$.

To conclude the proof, we let $\bar{\lambda} = \min_{v \in \mathcal{W}} \bar{\lambda}_v$. Given that $\mathcal{W}$ is finite, $\bar{\lambda}$ is well defined. By construction, for all $\lambda \in (0, \bar{\lambda})$ and for all $v \in \mathcal{W}$,

$$\Psi(\lambda, v) = \mathrm{U}\left[\mathbb{E}(\hat{\alpha}_\lambda, v), v\right] - u(\alpha^*) > 0.$$

This implies that $u(\hat{\alpha}_\lambda) > u(\alpha^*)$. Consequently, we conclude that α^* is not optimal.

Proof of Proposition 2

Step 1. *If* $\bigcap_{v \in M_{\alpha^*}} SUCS_v(\alpha) = \emptyset$ *for all* $\alpha \in \mathcal{A}$ *with* $S_\alpha \subseteq S_{\alpha^*}$*, then* α^* *is optimal.*

Proof of Step 1. Consider an action $\alpha^* \in \mathcal{A}$. If for all $\alpha \in \mathcal{A}$ with $S_\alpha \subseteq S_{\alpha^*}$, $\bigcap_{v \in M_{\alpha^*}} SUCS_v(\alpha) = \emptyset$, then trivially α^* is optimal by Proposition 1.

Step 2. *If* α^* *is optimal, then* $\bigcap_{v \in M_{\alpha^*}} SUCS_v(\alpha) = \emptyset$ *for all* $\alpha \in \mathcal{A}$ *with* $S_\alpha \subseteq S_{\alpha^*}$.

Proof of Step 2. Take an action $\alpha^* \in \mathcal{A}$. We show that if there exist an action $\tilde{\alpha} \in \mathcal{A}$, with $S_{\tilde{\alpha}} \subseteq S_{\alpha^*}$ and another action $\alpha \in \mathcal{A}$ such that $\alpha \succ_v \tilde{\alpha}$ for all $v \in M_{\alpha^*}$, then α^* is not optimal. Define a new mixed action $\hat{\alpha}_\lambda$ parametrized by $\lambda \in (0, \hat{\lambda})$ such that for all $a \in A$

$$\hat{\alpha}_\lambda(a) = \alpha^*(a) + \lambda\left[\alpha(a) - \tilde{\alpha}(a)\right],$$

where the upper bound $\hat{\lambda}$ is defined as follows:

$$\hat{\lambda} = \max\left\{\lambda \in (0, 1] : \ \forall a \in A,\ \hat{\alpha}_\lambda(a) \geq 0\right\}.$$

Notice that $\hat{\lambda}$ is well defined because $S_{\tilde{\alpha}} \subseteq S_{\alpha^*}$. We now show that there exists a value $\bar{\lambda} \in (0, \hat{\lambda}]$ such that for all $\lambda \in (0, \bar{\lambda})$, $u(\hat{\alpha}_\lambda) > u(\alpha^*)$. To this end, consider the function $\Psi\colon [0, \hat{\lambda}] \times \mathcal{W} \to \mathbb{R}$ such that $\Psi(\lambda, v) = \mathrm{U}\left[\mathbb{E}(\hat{\alpha}_\lambda, v), v\right] - u(\alpha^*)$, for all $\lambda \in [0, \hat{\lambda}]$ and $v \in \mathcal{W}$. As an intermediate step, we prove that for each $v \in \mathcal{W}$, there exists a value $\bar{\lambda}_v \in (0, \hat{\lambda}]$ such that for all $\lambda \in (0, \bar{\lambda}_v)$, $\Psi(\lambda, v) > 0$. Take a utility $v \in \mathcal{W}$. There are two possibilities:

1. If $\alpha \succsim_v \tilde{\alpha}$, then for all $\lambda \in (0, \hat{\lambda}]$

$$\begin{aligned}\Psi(\lambda, v) &= \mathrm{U}\left[\mathbb{E}(\hat{\alpha}_\lambda, v), v\right] - u(\alpha^*) \\ &= \mathrm{U}\left[\mathbb{E}(\alpha^*, v) + \lambda\left[\mathbb{E}(\alpha, v) - \mathbb{E}(\tilde{\alpha}, v)\right], v\right] - u(\alpha^*) \\ &\geq \mathrm{U}\left[\mathbb{E}(\alpha^*, v), v\right] - u(\alpha^*) \geq 0,\end{aligned}$$

   where at least one of the two weak inequalities holds strict. If $\alpha \succ_v \tilde{\alpha}$, the first inequality is strict because the function $\mathrm{U}[\cdot, v]$ is strictly increasing in the first argument. If instead $\alpha \sim_v \tilde{\alpha}$, the last inequality is strict because $v \notin M_{\alpha^*}$. For all such v, we let $\bar{\lambda}_v = \hat{\lambda}$.

2. If $\tilde{\alpha} \succ_v \alpha$, then $\Psi(\cdot, v)$ is strictly decreasing in the first argument because $\mathrm{U}[\cdot, v]$ is strictly increasing in the first argument. $\Psi(\cdot, v)$ is continuous in the first argument because $\mathrm{U}[\cdot, v]$ is continuous in the first argument. If $\Psi(\hat{\lambda}, v) \geq 0$, the result follows immediately by taking $\bar{\lambda}_v = \hat{\lambda}$. Suppose that $\Psi(\hat{\lambda}, v) < 0$. Notice that $\Psi(0, v) > 0$ because $\tilde{\alpha} \succ_v \alpha$ implies that $v \notin M_{\alpha^*}$. Therefore, by the Intermediate Value Theorem, there exists $\bar{\lambda}_v \in (0, \hat{\lambda})$ such that $\Psi(\bar{\lambda}_v, v) = 0$. By $\Psi(\cdot, v)$ strictly decreasing, $\Psi(\lambda, v) > 0$ for all $\lambda \in (0, \bar{\lambda}_v)$.

To conclude the proof, we let $\bar{\lambda} = \min_{v \in \mathcal{W}} \bar{\lambda}_v$. Given that $\mathcal{W}$ is finite, $\bar{\lambda}$ is well defined. By construction, for all $\lambda \in (0, \bar{\lambda})$ and for all $v \in \mathcal{W}$,

$$\Psi(\lambda, v) = \mathrm{U}\left[\mathbb{E}(\hat{\alpha}_\lambda, v), v\right] - u(\alpha^*) > 0.$$

This implies that $u(\hat{\alpha}_\lambda) > u(\alpha^*)$ and α^* is not optimal.

Proof of Proposition 3

Suppose that $\alpha^* \in \mathcal{A}$ is optimal.

Step 1. *If there is no $\alpha \in \mathcal{A}$ with $\alpha^* \neq \alpha$, such that $\alpha^* \sim_v \alpha$ for all $v \in M_{\alpha^*}$, then α^* is unique.*

Proof of Step 1. We show that if α^* is not unique, then there is an action $\alpha \in \mathcal{A}$, with $\alpha^* \neq \alpha$, such that $\alpha^* \sim_v \alpha$ for all $v \in M_{\alpha^*}$. Take an optimal action α with $\alpha \neq \alpha^*$. First, notice that for all $v \in M_{\alpha^*}$, it holds that $\alpha \succsim_v \alpha^*$. Consider the set $\{v \in M_{\alpha^*} : \alpha \sim_v \alpha^*\}$. By Proposition 1, this set is non-empty. If it coincides with M_{α^*}, then the proof is completed: $\alpha \sim_v \alpha^*$ for all $v \in M_{\alpha^*}$. Otherwise, define a new mixed action $\hat{\alpha}_\lambda$ parametrized by $\lambda \in (0, 1)$ such that for all $a \in A$

$$\hat{\alpha}_\lambda(a) = \lambda\alpha(a) + (1 - \lambda)\alpha^*(a).$$

We now show that there exists a value $\bar{\lambda} \in (0, 1]$ such that for all $\lambda \in (0, \bar{\lambda})$, $\hat{\alpha}_\lambda$ is optimal and for all $v \in M_{\hat{\alpha}_\lambda}$, $\hat{\alpha}_\lambda \sim_v \alpha^*$. To this end, consider the function $\Psi\colon [0, 1] \times \mathcal{W} \to \mathbb{R}$ such that $\Psi(\lambda, v) = \mathrm{U}\,[\mathbb{E}(\hat{\alpha}_\lambda, v), v] - u(\alpha^*)$, for all $\lambda \in [0, 1]$ and $v \in \mathcal{W}$. As an intermediate step, we prove that for each $v \in \mathcal{W}$, there exists a value $\bar{\lambda}_v \in (0, 1]$ such that for all $\lambda \in (0, \bar{\lambda}_v)$, $\Psi(\lambda, v) \geq 0$, with equality holding only if $v \in \{v \in M_{\alpha^*} : \alpha \sim_v \alpha^*\}$. Take a utility $v \in \mathcal{W}$. There are two possibilities:

1. If $\alpha \succsim_v \alpha^*$, then for all $\lambda \in (0, 1]$

$$\begin{aligned}\Psi(\lambda, v) &= \mathrm{U}\,[\mathbb{E}(\hat{\alpha}_\lambda, v), v] - u(\alpha^*)\\ &= \mathrm{U}\,[\lambda\mathbb{E}(\alpha, v) + (1 - \lambda)\mathbb{E}(\alpha^*, v), v] - u(\alpha^*)\\ &\geq \mathrm{U}\,[\lambda\mathbb{E}(\alpha^*, v) + (1 - \lambda)\mathbb{E}(\alpha^*, v), v] - u(\alpha^*)\\ &= \mathrm{U}\,[\mathbb{E}(\alpha^*, v), v] - u(\alpha^*) \geq 0,\end{aligned}$$

   where both weak inequalities hold as equalities only if $v \in \{v \in M_{\alpha^*} : \alpha \sim_v \alpha^*\}$. For all v with $\alpha \succsim_v \alpha^*$, we let $\bar{\lambda}_v = 1$.

2. If $\alpha^* \succ_v \alpha$, then $\Psi(\cdot, v)$ is strictly decreasing in the first argument because $\mathrm{U}\,[\cdot, v]$ is strictly increasing in the first argument. $\Psi(\cdot, v)$ is continuous in the first argument because $\mathrm{U}\,[\cdot, v]$ is continuous in the first argument. If $\Psi(1, v) \geq 0$, the result follows immediately by taking $\bar{\lambda}_v = 1$. Suppose that $\Psi(1, v) < 0$. Notice that $\Psi(0, v) > 0$ because $\alpha^* \succ_v \alpha$ implies that $v \notin M_{\alpha^*}$. Therefore, by the Intermediate Value Theorem, there exists $\bar{\lambda}_v \in (0, 1)$ such that $\Psi(\bar{\lambda}_v, v) = 0$. By $\Psi(\cdot, v)$ strictly decreasing, $\Psi(\lambda, v) > 0$ for all $\lambda \in (0, \bar{\lambda}_v)$.

To conclude the proof, we let $\bar{\lambda} = \min_{v \in \mathcal{W}} \bar{\lambda}_v$. Given that $\mathcal{W}$ is finite, $\bar{\lambda}$ is well defined. By construction, for all $\lambda \in (0, \bar{\lambda})$ and for all $v \in \mathcal{W}$,

$$\Psi(\lambda, v) = \mathrm{U}\,[\mathbb{E}(\hat{\alpha}_\lambda, v), v] - u(\alpha^*) \geq 0,$$

where the inequality holds as an equality for $v \in \{v \in M_{\alpha^*} : \alpha \sim_v \alpha^*\}$. Therefore, $u(\alpha^*) = u(\hat{\alpha}_\lambda)$ and given that action α^* is optimal, action $\hat{\alpha}_\lambda$ must be optimal as well. Moreover, by construction, $M_{\hat{\alpha}_\lambda} = \{v \in M_{\alpha^*} : \alpha \sim_v \alpha^*\} \subset M_{\alpha^*}$ and consequently $\hat{\alpha}_\lambda \sim_v \alpha^*$ for all $v \in M_{\hat{\alpha}_\lambda}$.

Step 2. *If α^* is unique, then there is no $\alpha \in \mathcal{A}$ with $\alpha^* \neq \alpha$, such that $\alpha^* \sim_v \alpha$ for all $v \in M_{\alpha^*}$.*

Proof of Step 2. Suppose that there exist two actions α^* and α, with $\alpha^* \neq \alpha$, such that α^* is optimal and $\alpha^* \sim_v \alpha$ for all $v \in M_{\alpha^*}$. We show that α^* is not unique. To this end, define a new mixed action $\hat{\alpha}_\lambda$ parametrized by $\lambda \in (0, 1)$ such that for all $a \in A$

$$\hat{\alpha}_\lambda(a) = \lambda\alpha(a) + (1 - \lambda)\alpha^*(a).$$

Replicating the same argument used for Step 1 of Proposition 3, we can show that there exists a value $\bar{\lambda} \in (0, 1]$ such that for all $\lambda \in (0, \bar{\lambda})$, $\hat{\alpha}_\lambda$ is optimal and $M_{\hat{\alpha}_\lambda} = M_{\alpha^*}$. This proves that the optimal action α^* is not unique.

Proof of Proposition 4

Step 1. *For any optimal action α^*, if $M_{\alpha^*} \subseteq M_\alpha$ for any other optimal action α, then α^* is efficient.*

Proof of Step 1. Take an action α^* that is optimal. We show that if α^* is not efficient, then there exists another optimal action $\hat{\alpha}$ such that $M_{\hat{\alpha}} \subset M_{\alpha^*}$. Given that α^* is not efficient, there exists $\alpha \in \mathcal{A}$ such that $\alpha \gtrsim_v \alpha^*$ for all $v \in M_{\alpha^*}$ and $\alpha \succ_v \alpha^*$ for some $v \in M_{\alpha^*}$. Also notice that by Proposition 1 there exists at least one $v \in M_{\alpha^*}$ with $\alpha \sim_v \alpha^*$ because α^* is optimal. Define a new mixed action $\hat{\alpha}_\lambda$ parametrized by $\lambda \in (0, 1)$ such that for all $a \in A$

$$\hat{\alpha}_\lambda(a) = \lambda\alpha(a) + (1 - \lambda)\alpha^*(a).$$

Replicating the same argument used for Step 1 of Proposition 3, we can show that there exists a value $\bar{\lambda} \in (0, 1]$ such that for all $\lambda \in (0, \bar{\lambda})$, $\hat{\alpha}_\lambda$ is optimal and $M_{\hat{\alpha}_\lambda} = \{v \in M_{\alpha^*} : \alpha \sim_v \alpha^*\} \subset M_{\alpha^*}$.

Step 2: *If α^* is efficient, then $M_{\alpha^*} \subseteq M_\alpha$ for any other optimal action α.*

Proof of Step 2. Take an optimal action α^*. We show that if there exists another optimal action α such that it is not true that $M_{\alpha^*} \subseteq M_\alpha$, then action α^* is not efficient. First, notice that for all $v \in M_{\alpha^*}$,

$$\mathrm{U}\left[\mathbb{E}(\alpha, v), v\right] \geq u(\alpha) = u(\alpha^*) = \mathrm{U}\left[\mathbb{E}(\alpha^*, v), v\right].$$

In particular, there is a utility $\bar{v} \in M_{\alpha^*}$ such that $\bar{v} \notin M_\alpha$. For this utility $\bar{v}$, the weak inequality holds strict. Therefore, α^* is not efficient.

Proof of Proposition 5

Step 1. *There exists an optimal action α^* that is Pareto efficient in M_{α^*}.*

Proof of Step 1. By contradiction, suppose that there is no optimal action α^* that is Pareto efficient in M_{α^*}. Take an optimal action α. By Proposition 4, there exists another optimal action $\hat{\alpha}$ such that $M_{\hat{\alpha}} \subset M_\alpha$. Given that $\mathcal{W}$ is finite, after a finite number of iterations of this argument, we obtain an optimal action α^* such that $M_{\alpha^*} = \{v\}$, for some utility $v \in \mathcal{W}$. By assumption, α^* is not Pareto efficient in M_{α^*}. This implies that there exists another action $\tilde{\alpha} \succ_v \alpha^*$. By Proposition 1, α^* is not optimal, a contradiction. Consequently, the set of minimal actions is non-empty. Before proceeding with Step 2 of the proof, we introduce some additional notation.

Let us denote by O the set of optimal actions, by $E \subseteq O$ the set of minimal actions and by M_{min} the set of worst-off utilities induced by minimal actions. By Step 1, the sets E and M_{min} are well defined. Let us consider the following optimization problem:

$$\max_{\alpha \in O} \min_{v \in \mathcal{W} \setminus M_{min}} \mathrm{U}\left[\mathbb{E}(\alpha, v), v\right]. \tag{C.1}$$

The set of optimal actions O is non-empty and compact by the maximum theorem. Moreover, the objective function is continuous and quasi-concave. Consequently, the set of solutions for problem (C.1) that we denote by E_1 is non-empty and again compact by the maximum theorem.

Step 2: $E_1 \subseteq E$.

Proof of Step 2. Take any optimal action α that is not minimal. We show that $\alpha \notin E_1$. Take any other action $\tilde{\alpha} \in E$. We have

$$\begin{aligned}
\min_{v \in \mathcal{W} \setminus M_{min}} \mathrm{U}\left[\mathbb{E}(\alpha, v), v\right] &= \min_{v \in \mathcal{W}} \mathrm{U}\left[\mathbb{E}(\alpha, v), v\right] \\
&= \min_{v \in \mathcal{W}} \mathrm{U}\left[\mathbb{E}(\tilde{\alpha}, v), v\right] \\
&< \min_{v \in \mathcal{W} \setminus M_{min}} \mathrm{U}\left[\mathbb{E}(\tilde{\alpha}, v), v\right],
\end{aligned}$$

where the first equality holds because α is not minimal and consequently $M_{min} \subset M_\alpha$, the second equality holds because both actions α and $\tilde{\alpha}$ are optimal, the third strict inequality holds because $\tilde{\alpha}$ is efficient and consequently $M_{\tilde{\alpha}} = M_{min}$.

Consider now the following maximization problem:

$$\max_{\alpha \in E_1} \sum_{v \in \mathcal{W} \setminus M_{min}} \mathbb{E}(\alpha, v). \tag{C.2}$$

We denote by E_2 the set of actions that solve (C.2). By the maximum theorem, E_2 is non-empty and compact.

Step 3: *Any $\alpha \in E_2$ is efficient.*

Proof of Step 3. Take any actions $\alpha \in E_2$ and $\tilde{\alpha} \in \mathcal{A}$. If $\tilde{\alpha} \notin O$, then given that α is optimal, there must exist a utility $v \in \mathcal{W}$ such that $\alpha \succ_v \tilde{\alpha}$. If instead $\tilde{\alpha} \in O \setminus E$, then $M_{\tilde{\alpha}} \setminus M_{min} \neq \emptyset$. For any utility $v \in M_{\tilde{\alpha}} \setminus M_{min}$, we have $\alpha \succ_v \tilde{\alpha}$. Let us assume now that $\tilde{\alpha} \in E$. Notice that $E_2 \subseteq E_1 \subseteq E$. We consider two different cases:

- **Case 1**: $\tilde{\alpha} \in E_1$. Given that $\tilde{\alpha} \in E_1$ and $\alpha \in E_2$, we have

$$\sum_{v \in \mathcal{W} \setminus M_{min}} \mathbb{E}(\alpha, v) \geq \sum_{v \in \mathcal{W} \setminus M_{min}} \mathbb{E}(\tilde{\alpha}, v).$$

  Consequently, if there exists a utility $v \in \mathcal{W} \setminus M_{min}$ such that $\tilde{\alpha} \succ_v \alpha$, then there must also exist a utility $\bar{v} \in \mathcal{W} \setminus M_{min}$ such that $\alpha \succ_{\bar{v}} \tilde{\alpha}$. Therefore, no action $\tilde{\alpha} \in E_1$ Pareto dominates $\alpha \in E_2$ in $\mathcal{W}$.

- **Case 2**: $\tilde{\alpha} \in E \setminus E_1$. In this case, for at least one utility $\bar{v} \in \mathcal{W} \setminus M_{min}$, we have

$$\begin{aligned} \mathrm{U}\left[\mathbb{E}(\alpha, \bar{v}), \bar{v}\right] &\geq \min_{v \in \mathcal{W} \setminus M_{min}} \mathrm{U}\left[\mathbb{E}(\alpha, v), v\right] \\ &> \min_{v \in \mathcal{W} \setminus M_{min}} \mathrm{U}\left[\mathbb{E}(\tilde{\alpha}, v), v\right] \\ &= \mathrm{U}\left[\mathbb{E}(\tilde{\alpha}, \bar{v}), \bar{v}\right], \end{aligned}$$

  where the first weak inequality holds for any utility $\bar{v} \in \mathcal{W} \setminus M_{min}$, the second strict inequality holds because $\alpha \in E_1$ and $\tilde{\alpha} \notin E_1$, the third equality holds for at least one utility $\bar{v} \in \mathcal{W} \setminus M_{min}$ because $\mathcal{W}$ is finite. For such utility $\bar{v}$, we have $\alpha \succ_v \tilde{\alpha}$. Therefore, no action $\tilde{\alpha} \in E \setminus E_1$ Pareto dominates $\alpha \in E_2$ in $\mathcal{W}$.

Proof of Proposition 6

Let $\mathcal{W} = \{v_1, v_2\}$ and suppose that the action $\alpha \in \mathcal{A}$ with $|S_\alpha| > 2$ is optimal. We show that there exists another optimal action $\hat{\alpha} \in \mathcal{A}$ with $|S_{\hat{\alpha}}| \leq 2$. If $u(\alpha) = u(a)$ for some $a \in A$, the statement follows. Suppose that $u(\alpha) > u(a)$ for all $a \in A$.

Step 1. *If α is optimal, then $M_\alpha = \{v_1, v_2\}$.*

Proof of Step 1. Suppose that $|M_\alpha| = 1$. Without loss of generality, let $M_\alpha = \{v_1\}$. Given that α is optimal, by Proposition 1 $\alpha \gtrsim_{v_1} a$ for all $a \in S_\alpha$. Therefore, $\alpha \sim_{v_1} a$ for all $a \in S_\alpha$. Given that $u(\alpha) > u(a)$ for all $a \in A$, it must be that $M_a = \{v_2\}$ for all $a \in S_\alpha$. Therefore, we have

$$\mathrm{U}\left[\mathbb{E}(\alpha, v_1), v_1\right] = \mathrm{U}\left[\max_{\tilde{a}\in S_\alpha} \mathbb{E}(\tilde{a}, v_1), v_1\right] > \mathrm{U}\left[\max_{\tilde{a}\in S_\alpha} \mathbb{E}(\tilde{a}, v_2), v_2\right] \geq \mathrm{U}\left[\mathbb{E}(\alpha, v_2), v_2\right],$$

where the equality holds because $\alpha \sim_{v_1} a$ for all $a \in S_\alpha$, the strict inequality because $M_a = \{v_2\}$ for all $a \in S_\alpha$ and the weak inequality because $\mathrm{U}[\cdot, v_2]$ is a strictly increasing function. Consequently, $\mathrm{U}[\mathbb{E}(\alpha, v_1), v_1] > \mathrm{U}[\mathbb{E}(\alpha, v_2), v_2]$, which contradicts $M_\alpha = \{v_1\}$.

Step 2. *If α is optimal, for all $a, a' \in S_\alpha$, we have $a \sim_{v_1} a'$ if and only if $a \sim_{v_2} a'$.*

Proof of Step 2. Suppose that there exist actions $a, a' \in S_\alpha$ such that $a \sim_{v_1} a'$ and $a \succ_{v_2} a'$. We show that α is not optimal. Consider a new action $\tilde{\alpha}$ such that $\tilde{\alpha}(a) = \alpha(a) + \alpha(a')$, $\tilde{\alpha}(a') = 0$ and $\tilde{\alpha}(a'') = \alpha(a'')$ for all $a'' \in S_\alpha \setminus \{a, a'\}$. Therefore, we have $\tilde{\alpha} \sim_{v_1} \alpha$ and $\tilde{\alpha} \succ_{v_2} \alpha$. Given that by Step 1 $M_\alpha = \{v_1, v_2\}$, then $M_{\tilde{\alpha}} = \{v_1\}$ and $u(\alpha) = u(\tilde{\alpha})$. However, Step 1 also implies that $\tilde{\alpha}$ is not optimal, which in turn implies that α is not optimal.

Whenever there are two actions $a, a' \in S_\alpha$, such that $a \sim_{v_1} a'$ and $a \sim_{v_2} a'$, it is possible to construct another optimal mixed action α' with $S_{\alpha'} \subset S_\alpha$. It is enough take α' such that $\alpha'(a) = \alpha(a) + \alpha(a')$, $\alpha'(a') = 0$ and $\alpha'(a'') = \alpha(a'')$ for all $a'' \in S_\alpha \setminus \{a, a'\}$. From now on, assume that there are no actions $a, a' \in S_\alpha$, such that $a \sim_{v_1} a'$ and $a \sim_{v_2} a'$, but still $|S_\alpha| > 2$.

Step 3. *If α is optimal, for all $a, a' \in S_\alpha$, we have $a \succ_{v_1} a'$ if and only if $a' \succ_{v_2} a$.*

Proof of Step 3. Suppose that there exist actions $a, a' \in S_\alpha$ such that $a \succ_{v_1} a'$ and $a \gtrsim_{v_2} a'$. We show that α is not optimal. By Step 2 we have $a \succ_{v_2} a'$, otherwise we could conclude $a \sim_{v_1} a'$. Therefore, by Proposition 2 we conclude that α is not optimal.

Step 4. *If α is optimal and $|S_\alpha| = n > 2$, there exists another optimal action $\tilde{\alpha}$ with $|S_{\tilde{\alpha}}| = n - 1$.*

Proof of Step 4. Consider three actions $a, a', a'' \in S_\alpha$. Without loss of generality, assume $a \succ_{v_1} a' \succ_{v_1} a''$. By Step 3, $a'' \succ_{v_2} a' \succ_{v_2} a$. Define a new mixed action α_λ

parametrized by $\lambda \in [0, 1]$ such that $\alpha_\lambda(a) = \lambda$ and $\alpha_\lambda(a'') = 1 - \lambda$. Consider the sets

$$S_1 \coloneqq \{\lambda \in [0,1] : \alpha_\lambda \succsim_{v_1} a'\} \quad \text{and} \quad S_2 \coloneqq \{\lambda \in [0,1] : \alpha_\lambda \succsim_{v_2} a'\}.$$

These sets are non-empty because $1 \in S_1$ and $0 \in S_2$. Moreover, they are closed because $\succsim_{v_1}$ and $\succsim_{v_2}$ are EU preferences and satisfy continuity. Therefore, let $k_1, k_2 \in (0,1)$ such that $S_1 = [k_1, 1]$ and $S_2 = [0, k_2]$. We have that $\alpha_{k_1} \sim_{v_1} a'$ and $\alpha_{k_2} \sim_{v_2} a'$. In what follows, we show that if α is optimal, then $k_1 = k_2$. If $k_1 < k_2$, for $\lambda \in (k_1, k_2)$ we have $\alpha_\lambda \succ_{v_1} a'$ and $\alpha_\lambda \succ_{v_2} a'$. By Proposition 2, α is not optimal. If $k_1 > k_2$, for $\lambda \in (k_2, k_1)$ we have $a' \succ_{v_1} \alpha_\lambda$ and $a' \succ_{v_2} \alpha_\lambda$. By Proposition 2, α is not optimal. Therefore, it must be that $k_1 = k_2 = k$. Define a new mixed action $\hat{\alpha}$ such that

- $\hat{\alpha}(a') = 0$.
- $\hat{\alpha}(a) = \alpha(a) + \alpha(a')k$.
- $\hat{\alpha}(a'') = \alpha(a'') + \alpha(a')(1 - k)$.
- $\hat{\alpha}(\hat{a}) = \alpha(\hat{a})$ for all $\hat{a} \in A \setminus \{a, a', a''\}$.

Notice that $\hat{\alpha}$ is optimal because $\hat{\alpha} \sim_{v_1} \alpha$ and $\hat{\alpha} \sim_{v_2} \alpha$. Moreover, $|S_{\hat{\alpha}}| = |S_\alpha| - 1$.

Therefore, if $|S_\alpha| = n > 2$, iterating $n - 2$ times Step 4 we can obtain an optimal action $\tilde{\alpha}$ with $|S_{\tilde{\alpha}}| = 2$.

Proof of Proposition 7

Take any finite set of utilities $\mathcal{W}$ and assume that $\arg\max_{a \in A} u(a) = A$.

Step 1. *If randomization is strictly beneficial, then there is no utility $v \in \mathcal{W}$ such that $P_v = A$.*

Proof of Step 1. Suppose that randomization is strictly beneficial. Then, there exists $\alpha \in \mathcal{A}$ such that $u(\alpha) > u(a)$ for all $a \in A$. By contradiction, suppose that there exists $v \in \mathcal{W}$ such that $P_v = A$. Then, for all $a \in A$, we have

$$\mathrm{U}\left[\mathbb{E}(a, v), v\right] = \mathrm{U}\left[\mathbb{E}(\alpha, v), v\right] \geq u(\alpha) > u(a) = \mathrm{U}\left[\mathbb{E}(a, v), v\right],$$

where the first equality holds because $P_v = A$ and $\arg\max_{a \in A} u(a) = A$, the second weak inequality by definition of $u(\cdot)$, the third strict inequality by assumption and the last equality because $P_v = A$. Therefore, we obtained a contradiction.

Step 2. *If there is no utility $v \in \mathcal{W}$ such that $P_v = A$, then randomization is strictly beneficial.*

Proof of Step 2. Take an action $\alpha \in \mathcal{A}$ with $S_\alpha = A$. Notice that for all $\tilde{a} \in A$, and for all $v \in \mathcal{W}$,

$$\mathrm{U}\left[\mathbb{E}(\alpha, v), v\right] \geq \mathrm{U}\left[\min_{a \in A} \mathbb{E}(a, v), v\right] \geq u(\tilde{a}),$$

where the first weak inequality holds because the function $\mathrm{U}[\cdot, v]$ is strictly increasing in the first argument and the second weak inequality because $\arg\max_{a \in A} u(a) = A$ and by definition of $u(\cdot)$. Moreover, if the first inequality holds as equality, then the second inequality is strict. Otherwise, we would have $P_v = A$. Therefore, incentives to randomize are strict.

Proof of Proposition 8

Suppose that $\mathcal{W} = \{v_1, v_2\}$, $\arg\max_{a \in A} u(a) = A$ and there is no utility $v \in \mathcal{W}$ such that $P_v = A$.

Step 1. *If action $\alpha \in \mathcal{A}$ is optimal, then $M_\alpha = \{v_1, v_2\}$.*

Proof of Step 1. See Step 1 in the proof of Proposition 6.

Step 2. *If action $\alpha \in \mathcal{A}$ is optimal, then*

$$S_\alpha \subseteq \arg\max_{a \in P_{v_1} \setminus P_{v_2}} \mathrm{U}\left[\mathbb{E}(a, v_2), v_2\right] \cup \arg\max_{a \in P_{v_2} \setminus P_{v_1}} \mathrm{U}\left[\mathbb{E}(a, v_1), v_1\right].$$

Proof of Step 2. Suppose that there is an action $a \in S_\alpha$ such that

$$a \notin \arg\max_{a \in P_{v_1} \setminus P_{v_2}} \mathrm{U}\left[\mathbb{E}(a, v_2), v_2\right] \cup \arg\max_{a \in P_{v_2} \setminus P_{v_1}} \mathrm{U}\left[\mathbb{E}(a, v_1), v_1\right].$$

We show that action α is not optimal. Assume that $M_\alpha = \{v_1, v_2\}$, otherwise by Step 1 the statement follows. There are three cases:

1. $a \in P_{v_1} \cup P_{v_2}$. Consider a new mixed action $\hat{\alpha}$ such that $\hat{\alpha}(a) = 0$ and for all $a' \neq a$, $\hat{\alpha}(a') = \alpha(a') + \alpha(a)(|A| - 1)^{-1}$. Notice that for all actions $a' \in A$, $a' \succsim_{v_1} a$ and $a' \succsim_{v_2} a$ because $\arg\max_{a \in A} u(a) = A$. Moreover, given that that there is no utility $v \in \mathcal{W}$ such that $P_v = A$, there must be two actions a_1 and a_2 such that $a_1 \succ_{v_1} a$ and $a_2 \succ_{v_2} a$. Therefore, it follows that $\hat{\alpha} \succ_{v_1} \alpha$ and $\hat{\alpha} \succ_{v_2} \alpha$, concluding that α is not optimal.

2. $a \in P_{v_1} \setminus P_{v_2}$. Take an action $a' \in P_{v_1} \setminus P_{v_2}$ such that $a' \succ_{v_2} a$. By assumption, such action exists. Consider a new mixed action $\hat{\alpha}$ such that $\hat{\alpha}(a) = 0$, $\hat{\alpha}(a') = \alpha(a) + \alpha(a')$ and for all $a'' \in A \setminus \{a, a'\}$, $\hat{\alpha}(a') = \alpha(a)$. Given that $a' \succ_{v_2} a$ and $a' \sim_{v_1} a$, it must be that $u(\alpha) = u(\hat{\alpha})$ and $M_{\hat{\alpha}} = \{v_1\}$. By Step 1, there exists another action $\tilde{\alpha}$ such that $u(\tilde{\alpha}) > u(\hat{\alpha}) = u(\alpha)$. Therefore, action α is not optimal.

3. $a \in P_{v_2} \setminus P_{v_1}$. Take an action $a' \in P_{v_2} \setminus P_{v_1}$ such that $a' \succ_{v_1} a$. By assumption, such action exists. Consider a new mixed action $\hat{\alpha}$ such that $\hat{\alpha}(a) = 0$, $\hat{\alpha}(a') = \alpha(a) + \alpha(a')$ and for all $a'' \in A \setminus \{a, a'\}$, $\hat{\alpha}(a') = \alpha(a)$. Given that $a' \succ_{v_1} a$ and $a' \sim_{v_2} a$, it must be that $u(\alpha) = u(\hat{\alpha})$ and $M_{\hat{\alpha}} = \{v_2\}$. By Step 1, there exists another action $\tilde{\alpha}$ such that $u(\tilde{\alpha}) > u(\hat{\alpha}) = u(\alpha)$. Therefore, action α is not optimal.

Step 3. *Action $\alpha \in \mathcal{A}$ is optimal if the following statements are true:*

(1) $M_\alpha = \{v_1, v_2\}$

(2) $S_\alpha \subseteq \underset{a \in P_{v_1} \setminus P_{v_2}}{\arg\max}\, \mathrm{U}\left[\mathbb{E}(a, v_2), v_2\right] \cup \underset{a \in P_{v_2} \setminus P_{v_1}}{\arg\max}\, \mathrm{U}\left[\mathbb{E}(a, v_1), v_1\right].$

Proof of Step 3. Consider any other mixed action $\tilde{\alpha}$ and suppose by contradiction that $u(\tilde{\alpha}) > u(\alpha)$. By (1), it must be that $\tilde{\alpha} \succ_{v_1} \alpha$ and $\tilde{\alpha} \succ_{v_2} \alpha$. By (2), $\tilde{\alpha} \succ_{v_1} \alpha$ implies that

$$\sum_{a \in P_{v_2} \setminus P_{v_1}} \tilde{\alpha}(a) > \sum_{a \in P_{v_2} \setminus P_{v_1}} \alpha(a).$$

Similarly, by (2), $\tilde{\alpha} \succ_{v_2} \alpha$ implies that

$$\sum_{a \in P_{v_1} \setminus P_{v_2}} \tilde{\alpha}(a) > \sum_{a \in P_{v_1} \setminus P_{v_2}} \alpha(a).$$

Therefore, it must be that

$$\sum_{a \in P_{v_2} \setminus P_{v_1}} \tilde{\alpha}(a) + \sum_{a \in P_{v_1} \setminus P_{v_2}} \tilde{\alpha}(a) > \sum_{a \in P_{v_2} \setminus P_{v_1}} \alpha(a) + \sum_{a \in P_{v_1} \setminus P_{v_2}} \alpha(a) = 1,$$

where the last equality holds by (2). Consequently, we obtained a contradiction.

Proof of Proposition 9

Assume that $A = \{a, b\}$ and $\mathcal{W} = \{v_1, v_2\}$.

Step 1. *If randomization is strictly beneficial, then there is no utility $v \in \{v_1, v_2\}$ such that $P_v = \{a, b\}$.*

Proof of Step 1. See Step 1 in the proof of Proposition 7.

Step 2. *If randomization is strictly beneficial, then either $a \succ_{v_1} b$ and $b \succ_{v_2} a$, or $b \succ_{v_1} a$ and $a \succ_{v_2} b$.*

Proof of Step 2. Suppose that it is not true that either $a \succ_{v_1} b$ and $b \succ_{v_2} a$, or $b \succ_{v_1} a$ and $a \succ_{v_2} b$. If $a \succsim_{v_1} b$ and $a \succsim_{v_2} b$, for all mixed actions $\alpha \in \mathcal{A}$, and for all $v \in \{v_1, v_2\}$, we have

$$\mathrm{U}\left[\mathbb{E}(a, v), v\right] \geq \mathrm{U}\left[\mathbb{E}(\alpha, v), v\right] \geq u(\alpha),$$

where the first inequality holds because $a \succsim_{v_1} b$ and $a \succsim_{v_2} b$, the second inequality by definition of $u(\cdot)$. Consequently, $u(a) \geq u(\alpha)$ and randomization is not strictly beneficial. If instead $b \succsim_{v_1} a$ and $b \succsim_{v_2} a$, for all mixed actions $\alpha \in \mathcal{A}$, and for all $v \in \{v_1, v_2\}$, we have

$$\mathrm{U}\left[\mathbb{E}(b, v), v\right] \geq \mathrm{U}\left[\mathbb{E}(\alpha, v), v\right] \geq u(\alpha),$$

where the first inequality holds because $b \succsim_{v_1} a$ and $b \succsim_{v_2} a$, the second inequality by definition of $u(\cdot)$. Consequently, $u(b) \geq u(\alpha)$ and randomization is not strictly beneficial.

Step 3. *Randomization is strictly beneficial if the following statements are true:*

(1) *There is no utility $v \in \{v_1, v_2\}$ such that $P_v = \{a, b\}$.*

(2) *Either $a \succ_{v_1} b$ and $b \succ_{v_2} a$, or $b \succ_{v_1} a$ and $a \succ_{v_2} b$.*

Proof of Step 3. By (1), either $M_a = \{v_1\}$ and $M_b = \{v_2\}$, or $M_a = \{v_2\}$ and $M_b = \{v_1\}$. Without loss of generality, assume $M_a = \{v_1\}$ and $M_b = \{v_2\}$. Then it must be that $b \succ_{v_1} a$. Otherwise, by (2) we have $a \succ_{v_1} b$ and $b \succ_{v_2} a$. This implies that

$$\mathrm{U}\left[\mathbb{E}(b, v_2), v_2\right] > \mathrm{U}\left[\mathbb{E}(a, v_2), v_2\right] > \mathrm{U}\left[\mathbb{E}(a, v_1), v_1\right] > \mathrm{U}\left[\mathbb{E}(b, v_1), v_1\right],$$

where the first strict inequality holds because $b \succ_{v_2} a$, the second strict inequality because $M_a = \{v_1\}$, and the third strict inequality because $a \succ_{v_1} b$. However,

$\mathrm{U}[\mathbb{E}(b, v_2), v_2] > \mathrm{U}[\mathbb{E}(b, v_1), v_1]$ contradicts $M_b = \{v_2\}$. Therefore, it must be that $b \succ_{v_1} a$ and by (2) $a \succ_{v_2} b$. Without loss of generality, assume that $u(a) \geq u(b)$. Define a new mixed action α_λ parametrized by $\lambda \in (0,1)$ such that $\alpha_\lambda(a) = 1 - \lambda$ and $\alpha_\lambda(b) = \lambda$. Notice that for any $\lambda \in (0,1)$, $\alpha_\lambda \succ_{v_1} a$. Moreover, for λ small enough, $M_{\tilde{\alpha}} = \{v_1\}$. Therefore,

$$u(\alpha_\lambda) = \mathrm{U}[\mathbb{E}(\alpha_\lambda, v_1), v_1] > \mathrm{U}[\mathbb{E}(a, v_1), v_1] = u(a) \geq u(b),$$

proving that randomization is strictly beneficial.

Proof of Corollary 1

If incentives to randomize are strict, by Proposition 9 either $a \succ_{v_1} b$ and $b \succ_{v_2} a$, or $b \succ_{v_1} a$ and $a \succ_{v_2} b$. Given that a is a mean-preserving spread of b, then either v_1 is strictly concave and v_2 is strictly convex, or vice versa. In any case, a C-EU DM is neither risk averse, nor risk seeking.

Proof of Corollary 2

By Proposition 6, if α is optimal, then $M_\alpha = \{v_1, v_2\}$. Suppose that $M_\alpha = \{v_1, v_2\}$. Given that randomization is strictly beneficial, either $a \succ_{v_1} b$ and $b \succ_{v_2} a$, or $b \succ_{v_1} a$ and $a \succ_{v_2} b$. Without loss of generality, assume $a \succ_{v_1} b$ and $b \succ_{v_2} a$. Consider any other mixed action $\tilde{\alpha} \neq \alpha$. If $\tilde{\alpha}(a) > \alpha(a)$, then $\alpha \succ_{v_2} \tilde{\alpha}$ and $u(\alpha) > u(\tilde{\alpha})$. If instead $\alpha(a) > \tilde{\alpha}(a)$, then $\alpha \succ_{v_1} \tilde{\alpha}$ and $u(\alpha) > u(\tilde{\alpha})$. Therefore α is optimal and unique.

Proof of Corollary 3

Let $\mathcal{W}_A = \{v_1, v_2\}$ and $\mathcal{W}_B = \{w_1, w_2\}$. Assume that (α, β) is a mixed Nash equilibrium of G.

Step 1. *If $u_B(\beta, \alpha) > \max\{u_B(b_1, \alpha), u_B(b_2, \alpha)\}$, then $\alpha \in \bar{A}^{\mathrm{o}}$.*

Proof of Step 1. Suppose that $\alpha \notin \bar{A}^{\mathrm{o}}$. Then, either $b_1 \succsim_{w,\alpha} b_2$ for all utilities $w \in \mathcal{W}_B$ or vice versa. In both cases, by Proposition 9, incentives to randomize for player B are not strict.

Step 2. *If $\alpha \in \bar{A}^{\mathrm{o}}$, then $u_B(\beta, \alpha) > \max\{u_B(b_1, \alpha), u_B(b_2, \alpha)\}$.*

Proof of Step 2. Suppose that $\alpha \in \bar{A}^{\mathrm{o}}$. Then, either $b_1 \succ_{w_1,\alpha} b_2$ and $b_2 \succ_{w_2,\alpha} b_1$, or $b_2 \succ_{w_1,\alpha} b_1$ and $b_1 \succ_{w_2,\alpha} b_2$. Without loss of generality, assume that the first case holds. By Proposition 9, it is enough to show that there is no utility $w \in \mathcal{W}_B$

such that $P_{w,\alpha} = S_B$. By contradiction, suppose that $P_{w_1,\alpha} = S_B$. Then, we have

$$\begin{aligned} u_B(b_1,\alpha) &= \mathrm{U}\left[\mathbb{E}_\alpha(b_1,w_1),w_1\right] \\ &> \mathrm{U}\left[\mathbb{E}_\alpha(\beta,w_1),w_1\right] \\ &\geq u_B(\beta,\alpha), \end{aligned}$$

where the first equality holds because $P_{w_1,\alpha} = S_B$, the second strict inequality because $b_1 \succ_{w_1,\alpha} b_2$ implies $b_1 \succ_{w_1,\alpha} \beta$, and the third weak inequality by definition of $u_B(\cdot,\alpha)$. However, given that (α,β) is a mixed Nash equilibrium of G, this is not possible.

Step 3. *$\beta \in \bar{B}^{\mathrm{o}}$ if and only if then $u_A(\alpha,\beta) > \max\{u_A(a_1,\beta), u_A(a_2,\beta)\}$.*

Proof of Step 3. It follows from the same arguments that we use in Steps 1 and 2.

Proof of Corollary 4

Let $X \subseteq B^{\mathrm{o}}$ and assume that for all $\beta \in X$, incentives to randomize are strict for player A. By Corollary 2, the unique optimal mixed action α of player A satisfies $M_\alpha = \mathcal{W}_A$. This implies that α satisfies $\mathbb{E}_\beta[\alpha, v_A] = \mathbb{E}_\beta[\alpha, w_A]$. Solving this equation for α yields the desired result.

Proof of Corollary 5

Let $\mathcal{W}_A = \{v_1, v_2\}$ and $\mathcal{W}_B = \{w_1, w_2\}$. Assume that (α,β) is a mixed Nash equilibrium of G.

Step 1. *If (α,β) is strict, then it is efficient.*

Given that (α,β) is a strict mixed Nash equilibrium, given the correct conjectures, incentives to randomize are strict for both players, and actions α and β are optimal. By Corollary 2, α and β are the unique optimal actions and therefore they are efficient. Consequently, the equilibrium (α,β) is efficient.

Step 2. *If (α,β) is efficient, then it is neither weak nor partially strict provided that $\bar{A}^{\mathrm{o}}$ and $\bar{B}^{\mathrm{o}}$ are non-empty.*

Suppose that (α,β) is a weak or partially strict mixed Nash equilibrium. We show that (α,β) is not efficient. Given that (α,β) is not a strict equilibrium, there exists one player for which randomization is not strictly beneficial in equilibrium. Without loss of generality, assume that this is true for player A. That is, given the correct conjecture β, we have

$$u_A(\alpha,\beta) = \max\{u_A(a_1,\beta), u_A(a_2,\beta)\}.$$

By Corollary 3, it must be that either $\beta = \min(\bar{B})$ or $\beta = \max(\bar{B})$. Given that $\bar{B}^{\mathrm{o}}$ is non-empty, these two quantities are distinct. In both cases, it must be the case that one utility in $\mathcal{W}_A$ is indifferent between the pure actions a_1 and a_2, while the other utility in $\mathcal{W}_A$ strictly prefers one of the two actions. Consequently, one of the two pure actions Pareto dominates action α in $\mathcal{W}_A$, proving that action α is not efficient. Therefore, also the mixed Nash equilibrium (α, β) is not efficient.