

A THEORY OF ROBUST EXPERIMENTS FOR CHOICE UNDER UNCERTAINTY

S. GRANT*, J. KLINE, I. MENEGHEL, R. TOURKY, AND J. QUIGGIN

ABSTRACT. Thought experiments are commonly used in the theory of behavior in the presence of risk and uncertainty to test the plausibility of proposed axiomatic postulates. The prototypical examples of the former are the Allais experiments and of the latter are the Ellsberg experiments. Although the lotteries from the former have objectively specified probabilities, the participants in both kinds of experiments may be susceptible to small deviations in their subjective beliefs. These may result from a variety of factors that are difficult to check in an experimental setting: including deviations in the understanding and trust regarding the experiment, its instructions and its method. Intuitively, an experiment is robust if it is tolerant to small deviations in subjective beliefs in models that are in an appropriate way close to the modeler's model. We characterize robust experiments in a theoretical framework and give a number of recipes for the robustification of experiments and their elicited preferences.

THE AUSTRALIAN NATIONAL UNIVERSITY AND THE UNIVERSITY OF QUEENSLAND.

E-mail addresses: `simon.grant@anu.edu.au`, `j.kline@uq.edu.au`,

`idione.meneghel@anu.edu.au`, `j.quiggin@uq.edu.au`, `rabeetourky@gmail.com`.

Date: March 2, 2015.

*Corresponding author.

We thank George Mailath for discussions in the early stages of the development of this project. We also thank Steven Roberts for introducing us to stable selection criteria in statistics and their relation to this work. The comments of Jim Cox, Jürgen Eichberger and Glenn Harrison have greatly improved the exposition. Finally, we thank Mamoru Kaneko, Ludovic Renou, Matthew Ryan and the seminar audiences at Bielefeld, Georgia State and Paris-Dauphine.

1. INTRODUCTION

The development of decision theory has been driven, in large measure, by thought experiments questioning the core postulates of the expected utility model, axiomatized for choice under risk by von Neumann and Morgenstern (1944) and for uncertainty by Savage (1954). The various forms of the expected utility model are theories that summarize the psychological motivations and the behavior of individuals by attributing to them axiomatic interpretations in the language of probability theory.

The first thought experiment to present a serious challenge to expected utility theory was the common consequence problem proposed by Allais (1953) and often referred to as the Allais ‘paradox.’ The problem consisted of a pair of choices, such that the most plausible pair of choices appeared inconsistent with any expected utility representation. The point of the common consequence problem was to challenge the intuitive plausibility of axiomatic systems in which common consequences occurring with equal probability could be disregarded in comparing choices. This feature of the von Neumann-Morgenstern system was gradually formalized as the ‘independence axiom’ (see Fishburn and Wakker (1995) for the history of this process).

Similarly, Ellsberg (1961) proposed thought-experiments for which intuitively plausible choices were inconsistent with the existence of any well-defined subjective probabilities. Discussion of Ellsberg’s paradox has mostly focused on the idea of ambiguity. However, the construction of the one-urn experiment can also be seen as a sharp test of Savage’s ‘sure-thing’ principle. The thought experiments of Allais and Ellsberg are examples of a large class of choice problems designed to test certain postulates on individual behavior in the presence of risk and uncertainty.

The criticisms of Allais and Ellsberg had little impact on the rise of expected utility to theoretical dominance, largely displacing alternatives such as the mean-variance theory of Markowitz (1952). Moreover, at the time the criticism was made, experimental economics was in its infancy.

The 1970s saw a revival of interest in criticisms of expected utility theory, and of attempts to test the predictions of expected utility theory in laboratory experiments, most notably those of Kahneman and Tversky (1979). In a typical experiment choices between a finite number of uncertain bets are elicited from participants. If

a significant number of sampled choices are anomalous, that is, inconsistent with the postulates, then the analyst would like to conclude that these postulates need revision. The work of Kahneman and Tversky (1979) supported by other early experimental results, such as MacCrimmon and Larsson (1979), gave influential support to this view.

It is necessary, however, to consider the alternative view that the postulates are descriptively valid, and that the analyst's interpretation of stated choices as an expression of non-EU preferences is subject to error. For example, it has been suggested that observed choices inconsistent with the predictions of expected utility theory may arise from "mistakes, carelessness, slips, inattentiveness" (Hey (1995)), repetition-inconsistent choices (Neugebauer and Schmidt (2007)), mistrust of the experimenter (Kadane (1992), Quiggin (1993, pp. 42-43)), specific disutility from ambiguity (Smith 1969) or from inappropriate use of heuristics (Tversky and Kahneman (1974), Al-Najjar and Weinstein (2009)).

Projecting these kinds of explanations to their extremes one may even conclude that absent restrictions on the epistemic perspectives of participants, subjective expected utility theory is, in a Popperian sense, untestable. Indeed, we can rationalise any elicited choices in any experiment as arising from subjective expected utility preferences in a probability model that differs greatly from the model anticipated by the experimental designer. These wild deviations from the modeler's model can even be articulated in terms of the Karni and Schmeidler (1991) notion of conceivable states (Grabiszewski (2014)) to basically explain all preferences as arising from expected utility maximization.

To obtain a testable version of expected utility it is necessary to impose some restrictions on the extent to which elicited choices may be inconsistent from what is deemed an admissible expected utility preference ordering in the experiment. Since the 1990s, this approach has been explored in numerous papers, with conclusions more favourable to the 'EU + error' interpretation than those of earlier work. Hey and Orme (1994) used the Akaike information criterion for which expected utility was 'pipped-at-the-post' by rank-dependent utility with overweighting of extreme outcomes. Hey (1995) nonetheless concludes that:

It may be the case that these further explorations may alter the conclusion to which I am increasingly being drawn: that one can explain experimental analyses of decision making under risk better (and simpler) as EU plus noise – rather than through some higher level functional – as long as one specifies the noise appropriately.

Along similar lines Harrison (1994) argues that many experiments have proposed tests which fail criteria for good experimental design, notably that the rewards corresponding to the null hypothesis should be “perceptibly and motivationally greater” than the rewards corresponding to the alternative hypothesis.

In this paper we focus on a systematic analysis of this problem, arising in the design of experimental tests based on thought experiments, such as those of Allais and Ellsberg, and in more general experiments. Intuitively appealing experimental tests of theoretical postulates may be systematically fragile in the sense that arbitrarily small errors may be sufficient to produce violations of those postulates.

A well-conceived thought experiment, such as Ellsberg’s, has associated with it a tractable (canonical) probability model usually derived by means of classical Bernoulli-Laplace reasoning in which elementary events are identified and presumed to have equal probability, thus making the probability of any other event simply its cardinality divided by the total number of elementary events. Although the analyst cannot rule out the possibility that participants may have different probability models in mind, in a well-designed implementation of an experiment the modeler can reasonably expect that the epistemic perspective of participants in the experiment is closely related to the canonical probability model. So the focus of the present paper is on well-conceived and well-designed experiments in which one can reasonably anticipate that the perspective of participants remains close in an appropriate sense to the modeler’s model of the experiment.

The main idea of the paper is that a robust challenge to a decision-theoretic model arising from an experiment remains a challenge in any probability model that is close to the canonical model. With this in mind, we demonstrate that many classic experiments such as the prototypical Allais and Ellsberg experiments and their derivatives (Machina (2009)), although well-conceived, do not pose convincing challenges to expected utility theory. We show, however, how they can be modified to overcome this

problem. We also explain how experimental design methods used in Ellsberg type experiments (Binmore et al. (2012)) can be understood as the robustification of the experiment within our framework.

The paper is organized as follows. Section 2 provides an articulation of the main concepts, results and insights in terms of Ellsberg’s single-urn thought experiment. The formal framework is set up in Section 3. Robust experiments are introduced in Section 4 and a number of examples of non-robust experiments are studied in Section 5.¹ Section 6 outlines different methods for robustifying an experiment. All proofs appear in Appendix A.

2. THE MAIN CONCEPTS, RESULTS AND INSIGHTS

Before the formal development of our framework that begins in Section 3, we provide in this section a summary of the main concepts, results and insights: introducing and illustrating each in the context of Ellsberg’s single-urn thought experiment.

In Ellsberg’s single-urn thought experiment, the reader is asked to “imagine an urn known to contain 30 red balls and 60 black and yellow balls, the latter in unknown proportion.” (Ellsberg, 1961, p. 653). A ball is to be drawn from the urn. On the basis of the color of the ball drawn, first consider a choice between a bet that pays \$100 if the ball drawn is red and nothing otherwise, denoted b_R , and a bet that pays \$100 if the ball drawn is black and nothing otherwise, denoted b_B . Next consider a choice between a bet that pays \$100 if the ball drawn is red or yellow and nothing if it is black, denoted b_{RY} , and a bet that pays \$100 if the ball drawn is black or yellow and nothing if it is red, denoted b_{BY} . Ellsberg argues that anyone exhibiting the preference pattern $b_R \succ b_B$ and $b_{BY} \succ b_{RY}$ is “simply not acting ‘as though’ they assigned numerical or even qualitative probabilities to the events in question.” (Ellsberg, 1961, p. 656) In particular, this means such a preference pattern is inconsistent with subjective expected utility theory.

Ellsberg’s reasoning rests on the assumption that the subject in such an experiment takes the state space to be the *sample* space $\{s_R, s_B, s_Y\}$, where s_c is the sample-state in which a ball of color c is drawn from the urn *independent* of which bet has been chosen by the subject in either problem. By identifying each of these three

¹An additional example from Machina (2009) is analyzed in Appendix B.

states with the corresponding vector of bet-consequences we obtain the following 4×3 consequence matrix:

$$C = \begin{array}{c} \\ b_R \\ b_B \\ b_{RY} \\ b_{BY} \end{array} \begin{array}{ccc} s_R & s_B & s_Y \\ \left[\begin{array}{ccc} 100 & 0 & 0 \\ 0 & 100 & 0 \\ 100 & 0 & 100 \\ 0 & 100 & 100 \end{array} \right] \end{array}.$$

The set of admissible preferences are ones that can be induced by a subjective expected utility maximizing decision-maker characterized by a pair (u, p) where

- (1) u is any (Bernoulli) utility function for which $u(0) < u(100)$.
- (2) p is any probability distribution over the sample space satisfying $p(s_R) = \frac{1}{3}$, $p(s_B) = q$ and $p(s_Y) = \frac{2}{3} - q$, for some number q , $\frac{1}{90} \leq q \leq \frac{59}{90}$.²

We refer to the set \mathcal{A} of all such pairs as the *admissible parameters* and the pair (C, \mathcal{A}) as the *canonical version* of the single-urn Ellsberg experiment *in belief form*.

For each pair $(u, p) \in \mathcal{A}$, we can map each bet into a lottery with a support comprising the two outcomes, $u(0)$ and $u(100)$. For example, the bet b_{RY} is mapped to the lottery that assigns probability $\frac{1}{3} + p(s_Y)$ to outcome $u(100)$ and assigns the complementary probability $p(s_B)$ to outcome $u(0)$. The final step is to translate, by the application of the *expected utility rule*, each lottery into a real number from which is generated the preferences over bets characterized by that pair.

Notice that for any $(u, p) \in \mathcal{A}$ the difference in the probability assigned to the outcome $u(100)$ by the lottery induced from bet b_R compared to that assigned by the lottery induced from bet b_B is $p(s_R) - p(s_B)$. But this can in turn be seen to be equal to the difference in the probability assigned to the outcome $u(100)$ by the lottery induced from bet b_{RY} compared to that assigned by the lottery induced from bet b_{BY} . That is, the lottery induced from bet b_R assigns more probability to the better outcome than does the lottery induced from bet b_B *if and only if* the lottery

²The restriction $p(s_R) = \frac{1}{3}$ accords with the information that 30 out of the 90 balls are red. The restriction to positive probability for the other two states accords with the information that the urn contains both black and yellow balls, albeit in unknown proportion. Alternatively, as Ellsberg writes, “imagine a sample of two drawn from the 60 black and yellow balls has resulted in one black and one yellow.” (Ellsberg, 1961, pp. 653-4)

induced from bet b_{RY} assigns more probability to the better outcome than does the lottery induced from bet b_{BY} . So we conclude any preference ordering \succsim over the four bets, with $b_R \succ b_B$ and $b_{BY} \succ b_{RY}$, is not admissible in this experiment.

Notice further, that such a preference pattern remains inadmissible even if we allow for perturbations of any admissible belief in the direction of *any* belief over the sample space $\{s_R, s_B, s_Y\}$.³ In this sense, we view the canonical version of the Ellsberg experiment as being *internally robust*.

Imagine now, however, the experimenter suspects that some participants have an alternative perception of the situation which corresponds to the experiment in belief form (C', \mathcal{A}') given by the consequence matrix

$$C' = \begin{array}{c} \\ b_R \\ b_B \\ b_{RY} \\ b_{BY} \end{array} \begin{array}{cccc} s_R & s_B & s_Y & s^* \\ \left[\begin{array}{cccc} 100 & 0 & 0 & 0 \\ 0 & 100 & 0 & 0 \\ 100 & 0 & 100 & 0 \\ 0 & 100 & 100 & 100 \end{array} \right] \end{array},$$

with \mathcal{A}' being the set of all (u, p') for which there is $(u, p) \in \mathcal{A}$ satisfying $p'(s) = p(s)$ for all $s \in \{s_R, s_B, s_Y\}$, thus making $p'(s^*) = 0$. That is, (C', \mathcal{A}') is very similar to the canonical version but with an additional state s^* that has zero probability in all the admissible parameters of the experiment. One possible interpretation for this state s^* is the participant conceives of, but places zero probability on, the possibility that the experimenter can “manipulate” the draw of a ball whose color is not red by substituting a yellow ball for a black one or a black ball for a yellow one, whenever such a substitution results in the bet paying out \$0 instead of \$100. Hence, the only bet that pays out \$100 in s^* is b_{BY} , the one that pays out \$100 no matter whether the color of a non-red ball is black *or* yellow.

The version (C', \mathcal{A}') admits the same admissible and inadmissible preferences as the canonical experiment. However, the set of possible perturbations in beliefs is richer in (C', \mathcal{A}') than was the case in the canonical version. In particular, notice that the preference ordering \succsim , in which $b_{BY} \succ b_{RY} \succ b_R \succ b_B$, although inadmissible in

³This includes inadmissible beliefs such as those for which $p(s_R) \neq \frac{1}{3}$.

both versions of the experiment, can be rationalized by a subjective expected utility maximizer characterized by a utility function satisfying $u(0) < u(100)$ and a belief

$$\left(\frac{1}{3}, \frac{1-\varepsilon}{3}, \frac{1-\varepsilon}{3}, \frac{2\varepsilon}{3} \right),$$

where ε can be *any* number in $(0, 1)$, no matter how small. That is, although inadmissible, once we allow for small perturbations of any admissible belief (such as $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0)$) in the direction of any other probability distribution over the sample-space $\{s_R, s_B, s_Y, s^*\}$ (such as $(\frac{1}{3}, 0, 0, \frac{2}{3})$) the preference ordering becomes what we refer to in the sequel as ε -admissible in (C', \mathcal{A}') .⁴

Although the canonical version of the Ellsberg experiment is clearly what Ellsberg had in mind, we contend that one cannot rule out, either from *a priori* reasoning or from any *ex post* experimental observation, that a participant did not have some alternative version in mind, for example, the version (C', \mathcal{A}') described above. Notice that both of these versions are probability models with the same set of bets and, for each bet, the same set of consequences and, as we have already noted, the sets of admissible preferences coincide. However, the sample-spaces of the two versions differ, which we have seen lead to differences in what is ε -admissible. Thus if we desire a notion of robustness that does not depend on which of these two versions a participant has in mind, then internal robustness will not suffice.

Even more troubling, the sample space of the non-canonical version (C', \mathcal{A}') was chosen somewhat arbitrarily. We introduced and motivated it by one story, but there are other possible stories one might tell each with its own distinct probability model. There are many more probability models with widely varying sample spaces that are also consistent with the underlying uncertainty described in this Ellsberg scenario. Any of these potentially can serve as the version in the mind of a participant.

Building on this insight we define in Section 4 an equivalence class of experiments in belief form which includes all these potential versions. Our formal notion of robustness will be one that is invariant across such an equivalence class of experiments.

⁴This is reminiscent of Kadane (1992) in which he proposes that the participants' 'healthy scepticism' of the experimenter and a suspicion that he might manipulate the design to their disadvantage could 'explain' both Allais and Ellsberg type phenomena without having to resort to a model of behavior that does not conform to expected utility.

The associated equivalence relation is based on the distributions over consequences induced from the bets by the beliefs thus making no explicit reference to the sample space. Loosely speaking, each equivalence class of experiments in belief form can be interpreted as embodying a notion of *statistical* equivalence.

Given the (statistical) equivalence class associated with an experiment in belief form (C, \mathcal{A}) , we shall say a preference ordering over bets is *weakly-admissible* if it is ε -admissible for some version from this class. This leads naturally to our notion of robustness both for observed violations of the theory within an experiment and for the entire experiment itself.

- A preference ordering is *robustly inadmissible* if it is not weakly-admissible.
- An experiment is *robust* if every inadmissible preference ordering is robustly inadmissible.

Although this notion of robustness requires a property to hold on an entire equivalence class of experiments, we show it can be characterized solely in terms of the admissible preferences over the bets for any given (statistically equivalent) version of the experiment. So in the case of the Ellsberg experiment, for example, we can characterize robustness purely in terms of its canonical version (C, \mathcal{A}) .

In particular, the main result of the paper is a characterization of robust inadmissibility and a corollary characterizing robust experiments. They are stated in terms of the following partial orderings over preference orderings. A preference ordering over bets is *finer* than another distinct preference ordering if any strict preference between a pair of bets in the latter implies the corresponding strict preference holds in the former. Correspondingly, we say one preference ordering over bets is *coarser* than another if the latter is finer than the former.

Our characterizations are based on the insight that any refinement of any admissible preference ordering is *weakly* admissible.

Characterization of Robust Inadmissibility (Theorem 1):

An inadmissible preference ordering is robustly inadmissible if and only if no coarser preference ordering is admissible.

Characterization of Robust Experiments (Corollary 1.2):

An experiment is robust if and only if every preference ordering that is finer than an admissible preference ordering is also admissible.

Applying these results to the Ellsberg experiment, we see that any inadmissible preference ordering \succsim exhibiting $b_R \succ b_B$ and $b_{BY} \succ b_{RY}$ is not robustly inadmissible since the coarser preference ordering \succsim' in which $b_R \sim' b_B$ and $b_{BY} \sim' b_{RY}$ is admissible as it can be induced by the admissible parameter pair (\hat{u}, \hat{p}) , where $\hat{u}(0) < \hat{u}(100)$ and $\hat{p}(s_B) = \frac{1}{3}$ ($= \hat{p}(s_R)$). Correspondingly, we see that the Ellsberg experiment is not robust since the *inadmissible* preference ordering \succsim is *finer* than the *admissible* preference ordering \succsim' .

The simple explanation of the characterizations is that non-robust anomalous-inadmissible preferences are highly susceptible to (misspecification) error in cases where true preferences are at indifferences. Even vanishingly small perturbations in beliefs can produce reported preferences that are inconsistent with the theory being tested. As we shall see, this susceptibility to small perturbations is not confined to the Ellsberg one-urn experiment. It arises with other classic experiments and reflects the way they were constructed: namely, as direct tests of axiom systems that imply preferences admit an expected utility representation.

Thus, the take-home message of Theorem 1 is that if a participant in an experiment reports a preference ordering that is inadmissible for that experiment, then the experimenter should check whether there is any admissible preference ordering that is *coarser*. If that turns out to be the case, then the experimenter cannot rule out the possibility that the participant is an expected utility maximizer, who has a statistically equivalent version ‘in mind’ for which the observed ‘anomalous’ preferences can be attributed to a (vanishingly) small perturbation of some admissible belief.

On the other hand, if every coarser preference ordering is also inadmissible, then this reported failure of admissibility is robust in the sense that it cannot be rationalized by a (vanishingly) small perturbation to *any* admissible belief within *any* of the equivalent versions. Thus, even if the experiment itself is not robust, this particular piece of data may be viewed as a *robust* rejection of the participant’s behavior conforming to the theory being tested. Section 5 features a number of examples of non-robust experiments.

In section 6 we present results showing how small adjustments to the payoffs in the consequence matrix of an experiment in belief form can make them robust. For the case of four bets the intuition underpinning these results can be illustrated in the

following modified version of the canonical form for the Ellsberg experiment $(\hat{C}, \hat{\mathcal{A}})$, where the consequence matrix is ‘slightly perturbed’ as follows:

$$\hat{C} = \begin{array}{c} \\ b_R \\ b_B \\ b_{RY} \\ b_{BY} \end{array} \begin{array}{ccc} s_R & s_B & s_Y \\ \left[\begin{array}{ccc} 99 & 0 & 0 \\ 0 & 100 & 0 \\ 100 & 0 & 100 \\ 0 & 100 & 100 \end{array} \right] \end{array},$$

and the set of admissible parameters $\hat{\mathcal{A}}$ is correspondingly expanded by allowing for any utility function u , satisfying $u(0) < u(99) < u(100)$. Notice for any admissible preference ordering \succsim , we have: (i) $b_{RY} \succ b_R$; (ii) $b_{BY} \succ b_B$; (iii) $b_{BY} \succ b_R$; and, (iv) $b_R \succsim b_B \Rightarrow b_{RY} \succ b_{BY}$. Thus for the inadmissible preference ordering $\hat{\succsim}$ in which $b_{BY} \hat{\succ} b_{RY} \hat{\succ} b_R \hat{\succ} b_B$, *no* coarsening is admissible. Hence, by Theorem 1 this preference ordering is robustly inadmissible. Furthermore, for any admissible preference ordering, it is straightforward to check that every finer preference ordering is also admissible, so by Corollary 1.2 the experiment is robust. Essentially, this robustified Ellsberg experiment has become a ‘one-sided’ test for ambiguity aversion, since the preference ordering \succsim' in which $b_{RY} \succ' b_{BY} \succ' b_B \succ' b_R$ is *now* admissible.⁵

3. AN EXPERIMENT IN BELIEF FORM

An experiment is designed to test whether decision makers in a certain population behave in a way that is consistent with a given decision-making theory involving subjective expected utility.

3.1. The experiment in belief form. An experiment comprises a finite number B of *bets* (sometimes referred to as *actions*, *choices* or *prospects*) and a finite number S of *sample states* that determine the outcomes of the bets. We abuse notation and write $B = \{b_1, b_2, \dots, b_B\}$ for the set of bets and $S = \{s_1, s_2, \dots, s_S\}$ for the *sample space*. We use b for elements of B to denote a bet and s for elements of S to denote a state.

⁵Notice \succsim' can be induced by any $(u', p') \in \hat{\mathcal{A}}$ for which $\frac{u'(99) - u'(0)}{u'(100) - u'(0)} < 3p'(s_B) < 1$.

A *consequence matrix* is a function from $S \times B$ to \mathbb{R} associating with each pair bs a number c_{bs} and thus can be represented in matrix form:

$$C = \begin{matrix} & s_1 & s_2 & \cdots & s_S \\ \begin{matrix} b_1 \\ b_2 \\ \vdots \\ b_B \end{matrix} & \begin{bmatrix} c_{b_1 s_1} & c_{s_1 b_2} & \cdots & c_{b_1 s_S} \\ c_{b_2 s_1} & c_{b_2 s_2} & \cdots & c_{b_2 s_S} \\ \vdots & \vdots & \ddots & \vdots \\ c_{b_B s_1} & c_{b_B s_2} & \cdots & c_{b_B s_S} \end{bmatrix} \end{matrix}$$

Given a bet b and sample-state s let

$$C_b = \begin{bmatrix} c_{bs_1} & c_{bs_2} & \cdots & c_{bs_S} \end{bmatrix} \quad \text{and} \quad C_s = \begin{bmatrix} c_{b_1 s} \\ c_{b_2 s} \\ \vdots \\ c_{b_B s} \end{bmatrix}$$

denote the bet b row of C and the sample-state s column of C , respectively. For each bet b let \underline{C}_b be the set $\{c_{bs} : s \in S\}$, which is the set of *underlying consequences of the bet b* . Likewise, let \underline{C} be the set $\{c_{bs} : b \in B, s \in S\}$, of *underlying consequences of the experiment*.

A (*Bernoulli*) *utility index* is a strictly increasing function $u : \underline{C} \rightarrow \mathbb{R}$. Each utility index gives rise to the $S \times B$ utility matrix:

$$\begin{bmatrix} u(c_{b_1 s_1}) & u(c_{b_1 s_2}) & \cdots & u(c_{b_1 s_S}) \\ u(c_{b_2 s_1}) & u(c_{b_2 s_2}) & \cdots & u(c_{b_2 s_S}) \\ \vdots & \vdots & \ddots & \vdots \\ u(c_{b_B s_1}) & u(c_{b_B s_2}) & \cdots & u(c_{b_B s_S}) \end{bmatrix}.$$

We will often identify a Bernoulli utility index with its utility matrix, and denote both by u . Of course, the consequence matrix C is the utility matrix for the utility index $i(x) = x$, which we call the *identity utility index*.

A (*subjective*) *belief p over S* is a probability distribution (density function) over the sample space S . An ordered pair (u, p) , where u is a utility index and p is a belief, is called a *parameter*. Parameters can be identified with lotteries in a standard way. A (*simple*) *lottery* is a probability density function $\ell : \mathbb{R} \rightarrow [0, 1]$ whose distribution

has finite support: we write $\underline{\ell}$ for the support $\{x: \ell(x) > 0\}$ of ℓ . Each parameter (u, p) induces from each bet $b \in B$ the lottery over outcomes ℓ_b^{up} defined by

$$\ell_b^{up}(x) = \sum_{s: x=u(c_{bs})} p_s.$$

For each parameter (u, p) we write $\mathbb{E}(\ell_b^{up})$ for the expectation of the lottery ℓ_b^{up} induced from bet b . That is,

$$\mathbb{E}(\ell_b^{up}) = \sum_{x \in \underline{\ell}_b^{up}} x \ell_b^{up}(x) = \sum_{s=1}^S u(c_{bs}) p_s,$$

and so can be interpreted as the *subjective expected utility* of bet b . We say that a parameter pair (u, p) represents a binary relation \succsim over the set of bets B whenever $b \succsim b'$ if and only if $\mathbb{E}(\ell_b^{up}) \geq \mathbb{E}(\ell_{b'}^{up})$. Notice that if (u, p) represents a binary relation \succsim then that relation must be a preference ordering, that is, complete and transitive.

A bet b is said to be *deterministic* for parameter (u, p) if the induced lottery ℓ_b^{up} is degenerate, that is, $|\underline{\ell}_b^{up}| = 1$.

A parameter (u, p) is said to be *non-degenerate* if for all distinct bets b and b' , if b is deterministic under (u, p) then $\ell_b^{up} \neq \ell_{b'}^{up}$.⁶ This non-degeneracy condition states that all induced degenerate lotteries are distinct. We are ready to define an experiment.

Definition 1. An *experiment (in belief form)* is a pair (C, \mathcal{A}) where C is a $B \times S$ consequence matrix and \mathcal{A} is a set of non-degenerate parameters (u, p) (on C). The parameters in \mathcal{A} are called *admissible* parameters of the experiment.

A preference ordering \succsim on B is called *admissible in the experiment* (C, \mathcal{A}) if it is represented by an admissible parameter $(u, p) \in \mathcal{A}$. We refer to any binary relation \succsim on B that is not admissible in the experiment as *inadmissible in the experiment*.

As an example, the canonical version (C, \mathcal{A}) of the Ellsberg experiment from Section 2 is indeed an experiment, as is the alternative version (C', \mathcal{A}') .

The experimenter designs an experiment in belief form with the aim of testing expected utility theory, possibly in combination with some other global restriction, such as constant absolute risk aversion. The belief form allows her to choose a

⁶Since utility indexes are strictly increasing, notice that if (u, p) is non-degenerate, then so is (u', p) for any utility index u' .

consequence matrix C and an accompanying set of parameters \mathcal{A} that is tailored to express her challenge to the theory. The subjects participate in the designed experiment and their preferences over the set of bets are elicited. If the observed preferences of a subject turn out to be inadmissible then the analyst concludes that the decision-maker in question has preferences that do not conform to expected utility theory (as restricted by \mathcal{A}), posing a challenge to the theory.

3.2. Lottery based experiments. In special experiments, like the Allais common consequence and common ratio experiments, the design is presented effectively in terms of lotteries with objective probabilities. The experimenter then elicits preferences over these lotteries and checks for instance whether they can be represented by the von Neumann-Morgenstern utility function. These can be treated as a subclass of experiments in belief form since we shall see in Proposition 1, they always admit at least one rendition in belief form, with a state space and a single admissible belief that induces those lotteries.

Suppose that we are given a finite set of lotteries $\ell_1, \ell_2, \dots, \ell_B$. We write $K = \{1, 2, \dots, K\}$ as an index set for the set of consequences given by the set $\cup_{b \in B} \underline{\ell}_b$, that is, the union of the supports of the B lotteries. We also write $B = \{1, 2, \dots, b\}$ as an index set for the lotteries. The *lottery matrix* L induced by the lotteries is the $B \times K$ matrix

$$L = \begin{array}{c} \ell_1 \\ \ell_2 \\ \vdots \\ \ell_B \end{array} \begin{bmatrix} c_1 & c_2 & \cdots & c_K \\ \ell_{1c_1} & \ell_{1c_2} & \cdots & \ell_{1c_K} \\ \ell_{2c_1} & \ell_{2c_2} & \cdots & \ell_{2c_K} \\ \vdots & \vdots & \ddots & \vdots \\ \ell_{Bc_1} & \ell_{Bc_2} & \cdots & \ell_{Bc_K} \end{bmatrix},$$

where ℓ_{bc} is the probability assigned by lottery ℓ_b to consequence c , with $\ell_{bc} > 0$ if and only if $c \in \underline{\ell}_b$. Notice that a lottery matrix has no column of all zeros.

Let Λ denote the set of simple lotteries. We define lottery based experiments.

Definition 2. A *lottery based experiment* is a pair (L, \mathcal{U}) , where L is the lottery matrix induced from a lottery profile $(\ell_1, \ell_2, \dots, \ell_B) \in \Lambda^B$ such that

- (1) For any distinct bets b and b' , if $|\underline{\ell}_b| = 1$ then $\ell_b \neq \ell_{b'}$.

(2) \mathcal{U} is a set of (Bernoulli) utility indexes over K .

The lottery based experiment (L, \mathcal{U}) is *induced* by the experiment in belief form (C, \mathcal{A}) with B bets if $\mathcal{A} = \{(p, u) : u \in \mathcal{U}\}$ for some belief p on S and

$$\ell_b = \ell_b^{ip},$$

for the identity utility index i and all $b \in B$.

The next proposition states that for any lottery based experiment we can “reverse engineer” an experiment in belief form that induces it.

Proposition 1. Every lottery based experiment (L, \mathcal{U}) is induced by an experiment in belief form (C, \mathcal{A}) .

We illustrate this proposition for the Allais common consequence experiment. Moreover, we demonstrate that a lottery based experiment can be induced by multiple experiments in belief form with distinct sample spaces.

Allais Common Consequence Experiment. In the common consequence experiment (Allais, 1953, p. 527) the reader is asked to consider the following four lotteries $\{\ell_1, \ell_2, \ell_3, \ell_4\}$ over monetary consequences specified (in millions) as:

$$\ell_1(c) = \begin{cases} 1 & \text{if } c = \$1 \\ 0 & \text{if } c \neq \$1 \end{cases} \quad \ell_2(c) = \begin{cases} \frac{10}{100} & \text{if } c = \$5 \\ \frac{89}{100} & \text{if } c = \$1 \\ \frac{1}{100} & \text{if } c = \$0 \\ 0 & \text{if } c \notin \{\$0, \$1, \$5\} \end{cases}$$

$$\ell_3(c) = \begin{cases} \frac{11}{100} & \text{if } c = \$1 \\ \frac{89}{100} & \text{if } c = \$0 \\ 0 & \text{if } c \notin \{\$0, \$1\} \end{cases} \quad \ell_4(c) = \begin{cases} \frac{10}{100} & \text{if } c = \$5 \\ \frac{90}{100} & \text{if } c = \$0 \\ 0 & \text{if } c \notin \{\$0, \$5\} \end{cases}$$

Notice that $\underline{\ell}_1 \cup \underline{\ell}_2 \cup \underline{\ell}_3 \cup \underline{\ell}_4 = \{0, 1, 5\}$, thus the lottery based experiment representing this situation is the pair (L, \mathcal{U}) , where L is the 4×3 lottery matrix

$$L = \begin{bmatrix} 0 & 1 & 0 \\ \frac{1}{100} & \frac{89}{100} & \frac{1}{10} \\ \frac{89}{100} & \frac{11}{100} & 0 \\ \frac{9}{10} & 0 & \frac{1}{10} \end{bmatrix}$$

and \mathcal{U} is the set of all utility functions satisfying $u(0) < u(1) < u(5)$.

To construct one experiment in belief form that generates (L, \mathcal{U}) , take a sample space comprising just three elements $S = \{s_1, s_2, s_3\}$ with a corresponding consequence matrix

$$C = \begin{bmatrix} 1 & 1 & 1 \\ 5 & 0 & 1 \\ 1 & 1 & 0 \\ 5 & 0 & 0 \end{bmatrix}.$$

Set the unique admissible belief to be

$$p = \begin{bmatrix} p_1 \\ p_2 \\ p_3 \end{bmatrix} = \begin{bmatrix} \frac{1}{10} \\ \frac{1}{100} \\ \frac{89}{100} \end{bmatrix}.$$

The set of admissible parameters \mathcal{A} is the set of pairs (u, p) , where u is a utility function satisfying $u(0) < u(1) < u(5)$.⁷

Although the experiment in belief form above induces a lottery based experiment that corresponds to the Allais common consequence experiment, we stress that it is not the only experiment in belief form that does this. Furthermore, the sample spaces of alternative experiments in belief form will generally involve different correlation

⁷This formulation is closest to the way the example is presented in Allais (1953). It represents any experiment in which the design is essentially equivalent to the following. The subject is asked to consider the draw of a ball from an urn containing 100 balls numbered from 1 to 100, with (i) state s_1 corresponding to the event of the draw of a ball with any number from 1 to 10, (ii) state s_2 corresponding to the draw of the ball numbered 11, and (iii) state s_3 corresponding to the event of the draw of a ball with any number from 12 to 100.

structures for the lotteries. The version above is one in which the correlation is greatest, allowing us to specify a sample space of minimal cardinality of three.

Another possible version has the four lotteries distributed independently as formulated by Loomes and Sugden (1982). In this case, the sample space S' needs at least twelve elements. An example of a consequence matrix with twelve states is:

$$C' = \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 5 & 5 & 5 & 5 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 \\ 5 & 5 & 0 & 0 & 5 & 5 & 0 & 0 & 5 & 5 & 0 & 0 \end{bmatrix}.$$

The single admissible belief p' is the 12-dimensional column vector, in which

$$p'_s = \ell_1(c_{1s})\ell_2(c_{2s})\ell_3(c_{3s})\ell_4(c_{4s}),$$

for all the twelve states $s \in S'$. □

4. ROBUST EXPERIMENTS

4.1. Internally robust experiments. For the experimental results to constitute a robust challenge to the theory being tested, the set of *inadmissible* preference orderings in the experiment should be unaffected by (vanishingly) small perturbations in the beliefs in the experiment (C, \mathcal{A}) . To capture robustness, we introduce a weaker notion of admissibility that expands the set of admissible preference orderings and makes it more difficult to reject a theory.

Definition 3. A preference ordering \succsim over bets B is ε -admissible in an experiment in belief form (C, \mathcal{A}) if there is a pair $(u, p) \in \mathcal{A}$ and a belief \hat{p} over S such that

- (1) $\underline{\ell}_b^{u\hat{p}} \subseteq \underline{\ell}_b^{up}$ for all bets b , and
- (2) the preference \succsim is represented by the parameter $(u, ((1 - \varepsilon)p + \varepsilon\hat{p}))$ for all $\varepsilon \in (0, 1)$.

An experiment in belief form is *internally robust* if every ε -admissible preference ordering is admissible.

Our notion of robustness is based on the idea that a subject may have a slightly different experiment in mind, which can be captured by a slight perturbation in

beliefs outside the admissible set. We view $(1 - \varepsilon)p + \varepsilon\hat{p}$ as a perturbation of the admissible belief p in the direction of the belief \hat{p} . Notice that we allow for perturbations toward any belief over S that does not add positive probability to consequences that have zero probability of realising; in particular, to those outside the admissible set of beliefs. In an internally robust experiment an inadmissible preference ordering challenges theory more convincingly than in an experiment that is not internally robust, where inadmissible preferences may in fact be ε -admissible. Internal robustness ensures that for vanishingly small differences between the actual experiment, and the one in the subject's mind, the set of admissible preferences are the same.

Returning to the illustrative Ellsberg example given in Section 2, we can now verify that the canonical version (C, \mathcal{A}) is internally robust while the alternative version (C', \mathcal{A}') is not. So, if the analyst is convinced that the participants perceive the experimental design according to (C, \mathcal{A}) and its associated state space S , then the analyst will view the observation of any inadmissible preference pattern in the experiment (C, \mathcal{A}) as constituting a violation of the theory. With the second experiment (C', \mathcal{A}') , the violation of subjective expected utility posed by the preference pattern $b_{BY} \succ b_{RY}$ and $b_R \succ b_B$ is not as robust. The apparent violation may arise from a small perturbation of an admissible belief of a subjective expected utility maximizer.

4.2. Equivalence classes of experiments. In this section we first define an equivalence class of experiments in belief form which includes both the canonical version of the Ellsberg experiment and the alternative version in the illustrative example in Section 2.

In an experiment in belief form, (C, \mathcal{A}) , each admissible pair (u, p) induces a *lottery profile* $(\ell_1^{up}, \ell_2^{up}, \dots, \ell_B^{up}) \in \Lambda^B$. We use this fact to define equivalent versions of an experiment in belief form.

Definition 4. Fix an experiment in belief form (C, \mathcal{A}) . Its *reduction*, denoted $L_{(C, \mathcal{A})}$, is the set

$$L_{(C, \mathcal{A})} = \{(\ell_1^{up}, \ell_2^{up}, \dots, \ell_B^{up}) : (u, p) \in \mathcal{A}\}.$$

Two experiments are *versions of each other* if they have the same reduction.

Notice that any two experiments that are versions of each other have the same set of bets and the same set of outcomes for each of the corresponding bets.

Returning to the Ellsberg single-urn experiment, we see that the canonical and alternative forms described in Section 2 are versions of each other in the sense of Definition 4. Furthermore, the set of possible lottery profiles are those with support $\{\gamma, \delta\}$, $\gamma < \delta$ and probabilities of the form

$$\begin{array}{c} \gamma \quad \delta \\ \ell_{b_R} \left[\begin{array}{cc} \frac{2}{3} & \frac{1}{3} \\ \frac{90-n}{90} & \frac{n}{90} \\ \frac{n}{90} & \frac{90-n}{90} \\ \frac{1}{3} & \frac{2}{3} \end{array} \right], \\ \ell_{b_B} \\ \ell_{b_{RY}} \\ \ell_{b_{BY}} \end{array}$$

where $1 \leq n \leq 59$ and is naturally interpreted as the expected number of black balls.

These experiments would be regarded as equivalent by any decision-maker who is *probabilistically sophisticated* in the sense of Machina and Schmeidler (1992). Recall, a probabilistically sophisticated decision-maker has state-independent preferences over the consequences and only cares about a state in terms of its associated consequence and its likelihood.

By virtue of having the same reduction, it follows that any two versions of an experiment have the same set of admissible preferences. However, as we have already demonstrated in the context of the Ellsberg experiment, what can be ‘rationalized’ by small perturbations of the admissible beliefs in different versions will in general differ. As we noted above, ideally a notion of robustness should be a condition that holds for *all* versions of an experiment. This idea motivates the following definition.

Definition 5. A preference ordering over bets is *weakly-admissible* in an experiment if it is ε -admissible for some version of the experiment. A preference ordering is *robustly inadmissible* if it is not weakly-admissible. An experiment is *robust* if every inadmissible preference ordering is robustly inadmissible.

Underpinning our robustness notion is the idea that the analyst who designs the experiment (C, \mathcal{A}) does not know which version of the experiment is in the mind of the participant. So, for a *particular* observed violation of expected utility to be deemed robust, we require that no matter which version of the experiment the participant

has in mind, the violation cannot be attributed to a (vanishingly) small perturbation of an admissible belief within that version. Correspondingly, the experiment itself is deemed robust, if a perturbation of a belief in *any* version of the experiment does not affect its set of admissible preferences.

We state the following characterization of robust experiments.

Proposition 2. The following statements are equivalent for an experiment in belief form (C, \mathcal{A}) :

- (1) It is robust.
- (2) Every weakly-admissible preference ordering is admissible.
- (3) Every version of the experiment is internally robust.

Although robustness requires a property to hold on an entire equivalence class of experiments, we characterize it entirely in terms of the admissible preferences over the bets for any given version of the experiment. In particular, we characterize robustness purely in terms of the analyst's version of the experiment (C, \mathcal{A}) .

For any pair of distinct preference orderings, \succsim and \succsim' , we say that \succsim' is *finer* than \succsim on B , if for all $b, b' \in B$ we have $b \succ b'$ implies $b \succ' b'$.⁸ Correspondingly, we say that \succsim' on B is *coarser* than \succsim on B , whenever \succsim is finer than \succsim' .

We characterize a robustly inadmissible preference ordering in terms of inadmissibility of its coarsenings.

Theorem 1. An inadmissible preference ordering is robustly inadmissible if and only if no coarser preference ordering is admissible.

This theorem can be restated in terms of ε -admissibility.

Corollary 1.1. An inadmissible preference ordering is weakly admissible if and only if it is finer than an admissible one.

This is the form of Theorem 1 that we shall prove. We give two comments about the result: one about the only-if part, and one about the if-part. The only-if part of Corollary 1.1 states that if an inadmissible ordering is weakly admissible, then it must be finer than some admissible preference ordering. The proof of this part relies

⁸That is, \succsim' is finer than \succsim , if $\succ \subset \succ'$.

essentially on the continuity of the expected utility representation with respect to beliefs. Since many non-expected utility representations satisfy such continuity, this result may potentially be extended to these cases.

The if-part of Corollary 1.1 states that *every* inadmissible preference ordering that is finer than an admissible one is *also* weakly admissible. This comes from our strong notion of robustness which considers all versions and all perturbations of beliefs. If, on the other hand, the analyst had some additional information about the possible versions or perturbations, then some of the refinements of admissible preferences could be excluded. We do not explore such additional restrictions here and presume instead that the analyst considers all versions as possible.

Using Theorem 1, we can immediately characterize robust experiments in terms of refinements of all admissible preferences.

Corollary 1.2. An experiment in belief form is robust if and only if every preference ordering that is finer than an admissible preference ordering is also admissible.

4.3. Robustness for lottery-based experiments. Using the fact that lottery based experiments can be induced by experiments in belief form, we extend the notions of admissibility and robustness to lottery based experiments in the natural way. That is, a preference ordering in a lottery based experiment is admissible (respectively, [robustly] inadmissible) if it is admissible (respectively, [robustly] inadmissible) in any experiment in belief form that induces it. Similarly, a lottery based experiment is robust if any experiment in belief form that induces it is robust.

However, as was demonstrated in the Allais common consequence experiment, a lottery based experiment can be induced by multiple distinct experiments in belief form. Hence it is not possible to define unambiguously the corresponding notion of ε -admissible preferences over lotteries or internal robustness of a lottery based experiment.

5. EXAMPLES OF NON-ROBUST EXPERIMENTS

We present four examples. The first is a one-urn Ellsberg-style experiment based on an example that appears in Eichberger et al. (2007, p. 892). Although the experiment itself is not robust, we show that there is a robustly inadmissible preference

pattern that accords with our intuition of how an ambiguity averse decision-maker might choose. The second is the Allais common-consequence experiment which provides an example of a non-robust lottery-based experiment. The third is a two-urn Ellsberg-style experiment which corresponds to the ‘‘Reflection Example’’ from Machina (2009). The fourth is a non-robust test of CARA.

Example 1 (A Robustly Inadmissible Preference). Consider an urn that contains 200 balls numbered 1 to 200. The balls numbered 1 to 66 are red, the balls numbered 67 to $200 - 2n$ are black and the remainder (that is, those numbered from $[201 - 2n]$ to 200) are yellow. The only information a participant has about n is that it is an integer and that $1 \leq n \leq 66$. Let O (respectively, E) be the event that the ball drawn from the urn has an odd (respectively, even) number on it. Let R (respectively, B , Y) be the event that color of the ball drawn is red (respectively, black, yellow). Let OR be the event that the ball drawn from the urn has an odd number and its color is red, and so on. Notice that the number of balls that are black with an odd number on them or yellow with an even number on them is 67 no matter what value n takes. Similarly, the number of balls that are black with an even number on them or yellow with an odd number on them is also 67 no matter what value n takes. We take the sample space to be $S = \{OR, OB, OY, ER, EB, EY\}$ and the set of bets to be $B = \{b_1, b_2, b_3, b_4, b_5, b_6\}$. The payoffs are given in the following consequence matrix

$$C^{\text{RI}} = \begin{matrix} & OR & OB & OY & ER & EB & EY \\ \begin{matrix} b_1 \\ b_2 \\ b_3 \\ b_4 \\ b_5 \\ b_6 \end{matrix} & \left[\begin{array}{cccccc} \$100 & \$0 & \$0 & \$100 & \$0 & \$0 \\ \$0 & \$100 & \$0 & \$0 & \$0 & \$100 \\ \$100 & \$0 & \$0 & \$0 & \$0 & \$0 \\ \$0 & \$100 & \$0 & \$0 & \$0 & \$0 \\ \$0 & \$0 & \$0 & \$100 & \$0 & \$0 \\ \$0 & \$0 & \$0 & \$0 & \$0 & \$100 \end{array} \right] \end{matrix}.$$

The bet b_1 is a standard ‘unambiguous’ bet that the color of the ball drawn is red. The bet b_2 can be viewed as a way of implementing the suggestion by Raiffa (1961) to avoid the ambiguity associated with a bet on black or a bet on yellow by randomly choosing which of these two colors to bet on. Here we are using the property of

whether the number on the ball drawn is odd or is even ‘to randomize’ between the choices of black or yellow. A choice between bets b_3 and b_4 corresponds to a choice between betting on red versus betting on black *conditional* on the number of the ball drawn is odd. Similarly, a choice between bets b_5 and b_6 is a choice between betting on red versus betting on yellow, conditional on the number of the ball drawn is even.

The set of admissible parameters \mathcal{A}^{RI} consists of pairs (u, p) where u is any Bernoulli utility function with $u(0) < u(100)$ and p is a probability defined on S , satisfying

$$p(OR) = p(ER) = \frac{33}{200} \quad p(OB) = p(EB) = q \quad p(OY) = p(EY) = \frac{67}{200} - q,$$

where $\frac{1}{200} \leq q \leq \frac{66}{200}$.

We note that in the experiment $(C^{\text{RI}}, \mathcal{A}^{\text{RI}})$, the preference pattern $b_2 \succ b_1$, $b_3 \succ b_4$ and $b_5 \succ b_6$ is robustly inadmissible. To see this, notice that for any admissible preference ordering we must have $b_2 \succ b_1$, and $b_3 \succ b_4 \Rightarrow b_6 \succ b_5$. This follows since for any $(u, p) \in \mathcal{A}$,

$$\begin{aligned} p(OB) + p(EY) &> p(OR) + p(ER) \\ \max\{p(OB), p(EY)\} &> p(OR) = p(ER). \end{aligned}$$

Hence, any coarsening of $b_2 \succ b_1$ or $b_3 \succ b_4$ or $b_5 \succ b_6$ is inadmissible, thus by Theorem 1 the inadmissible preference pattern is robustly inadmissible. But this pattern accords with what we expect from someone who exhibits ambiguity aversion, since b_1, b_2, b_3, b_5 are all unambiguous (and b_2 first-order stochastically dominates b_1) while b_4 and b_6 are bets for which there is ambiguity about the probability of winning. \square

As we shall now see the Allais lottery based experiment is not robust. Furthermore, as a lottery-based experiment there is no ‘‘natural’’ experiment in belief form that induces it, one might argue it is even ‘‘less robust’’ than the Ellsberg experiment for which the canonical version was shown to be internally robust.

Example 2 (Allais Common Consequence Experiment). Recall the two experiments in belief form (C, \mathcal{A}) and (C', \mathcal{A}') from Section 3.2 that were both shown to induce the lottery based Allais common consequence experiment (L, \mathcal{U}) .

The preference ordering \succsim satisfying

$$l_1 \succ l_2 \succ l_4 \succ l_3$$

is inadmissible in both (C, \mathcal{A}) and (C', \mathcal{A}') . Thus it is inadmissible in the lottery based experiment (L, \mathcal{U}) . Consider, however, the belief p^* on S' , in which $p^*(s^*) = 1$, where s^* is the state in S' that has associated with it the vector of bet consequences

$$C'_{s^*} = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 5 \end{bmatrix}.$$

By straightforward calculation it follows that for a utility function u' in which $u'(1) = \frac{10}{11}u'(5) + \frac{1}{11}u'(0)$ and a belief $\hat{p} = \frac{1}{2}p' + \frac{1}{2}p^*$ on S' , we have

$$\mathbb{E} \left(\ell_1^{u'((1-\varepsilon)p' + \varepsilon\hat{p})} \right) > \mathbb{E} \left(\ell_2^{u'((1-\varepsilon)p' + \varepsilon\hat{p})} \right) > \mathbb{E} \left(\ell_4^{u'((1-\varepsilon)p' + \varepsilon\hat{p})} \right) > \mathbb{E} \left(\ell_3^{u'((1-\varepsilon)p' + \varepsilon\hat{p})} \right),$$

for any $\varepsilon \in (0, 1)$. Hence, \succsim is ε -admissible in (C', \mathcal{A}') which means it is not robustly inadmissible in (C, \mathcal{A}) and hence not robustly inadmissible in (L, \mathcal{U}) either.

Furthermore, the coarser preference ordering \succsim' satisfying

$$l_1 \sim' l_2 \succ' l_4 \sim' l_3$$

is admissible in (C, \mathcal{A}) as it is generated by the admissible pair (u', p) . So applying Corollary 1.2 we conclude that (C, \mathcal{A}) is not robust and thus neither is the Allais common consequence lottery based experiment (L, \mathcal{U}) . \square

We now turn to the ‘‘Reflection Example’’ from Machina (2009). It was designed as an Ellsberg-style experiment to generate choice paradoxes that could not be explained by any member of the generalization of expected utility known as Choquet Expected Utility (CEU). So we shall begin by undertaking the appropriate modifications of our framework to enable it to accommodate this larger family of preferences. We then show that, although the experiment is designed to elicit preference patterns that are inadmissible for this larger class of preferences, such inadmissible patterns are not robustly inadmissible even for the smaller class of subjective expected utility maximizers.

Example 3 (The Reflection Example from Machina (2009)). The subject is presented with two urns each containing 100 balls that are either Red or Black. Urn 1 is known to contain 50 balls of each color. The proportion of red balls in urn 2 is unknown. There are four bets $b_1, b_2, b_3,$ and b_4 .

We take the sample space to be $S = \{RR, RB, BR, BB\}$, where $\gamma_1\gamma_2$ is the state in which a ball of color γ_i is drawn from urn i . The consequence matrix is

$$C^R = \begin{array}{c} \\ b_1 \\ b_2 \\ b_3 \\ b_4 \end{array} \begin{array}{cccc} RR & RB & BR & BB \\ \left[\begin{array}{cccc} \$4000 & \$8000 & \$4000 & \$0 \\ \$4000 & \$4000 & \$8000 & \$0 \\ \$0 & \$8000 & \$4000 & \$4000 \\ \$0 & \$4000 & \$8000 & \$4000 \end{array} \right] \end{array}.$$

Exploiting the symmetry between bets b_1 and b_4 , and similarly, between bets b_2 and b_3 , Machina argues that we might expect a decision-maker to exhibit the preferences $b_1 \sim b_4$ and $b_2 \sim b_3$. When comparing b_1 to b_2 he observes that although the events in which they yield the best (respectively, the worst) outcome are ‘similarly’ ambiguous, the event in which the middle outcome occurs is unambiguous for b_2 but not for b_1 . Hence, one might argue that an ambiguity-averse decision-maker would strictly prefer b_2 to b_1 .

The class of preferences Machina has in mind is Choquet expected utility. Preferences in this class are characterized by a Bernoulli utility function u and a subjective belief which is now a capacity ν (or ‘non-additive probability’) defined over the set of subsets of S , that is normalized ($\nu(\emptyset) = 0$ and $\nu(S) = 1$) and set-monotonic ($E \subset F \subseteq S \Rightarrow \nu(E) \leq \nu(F)$).⁹ Fixing a capacity ν , its conjugate, denoted $\bar{\nu}$, is the capacity given by $\bar{\nu}(E) = 1 - \nu(E^c)$.

Each pair (u, ν) induces from each bet $b \in B$ the lottery $\ell_b^{u\nu}$ defined by

$$\ell_b^{u\nu}(x) = \nu(\{s : u(c_{sb}) \geq x\}) - \nu(\{s : u(c_{sb}) > x\}).$$

Correspondingly, we say that a pair (u, ν) represents a preference ordering \succsim over the set of bets B whenever $b \succsim b'$ if and only if $\mathbb{E}(\ell_b^{u\nu}) \geq \mathbb{E}(\ell_{b'}^{u\nu})$.

⁹The capacity ν is a probability if, in addition to being normalized and set-monotonic, it is additive, that is, $\nu(E) + \nu(F) = \nu(E \cup F) + \nu(E \cap F)$.

We take \mathcal{A}^R to be the set of pairs (u, ν) , where u is any Bernoulli utility function with $u(0) < u(4000) < u(8000)$ and ν is a capacity that along with its conjugate $\bar{\nu}$ satisfy the following ‘natural’ symmetry conditions:

$$\begin{aligned}\nu(RR) &= \nu(RB) = \nu(BR) = \nu(BB) > 0 \\ \bar{\nu}(RR) &= \bar{\nu}(RB) = \bar{\nu}(BR) = \bar{\nu}(BB) > 0\end{aligned}$$

We say a preference ordering \succsim is (CEU-)admissible in the experiment (C^R, \mathcal{A}^R) if it can be represented by some $(u, \nu) \in \mathcal{A}^R$. However, notice that for any $(u, \nu) \in \mathcal{A}^R$, we have:

$$\begin{aligned}\mathbb{E}(\ell_{b_1}^{u\nu}) &= \ell_{b_1}^{u\nu}(8000)u(8000) + \ell_{b_1}^{u\nu}(4000)u(4000) + \ell_{b_1}^{u\nu}(0)u(0) \\ &= (\nu(RB) - \nu(\emptyset))u(8000) + (\nu(\{RR, RB, BR\}) - \nu(RB))u(4000) \\ &\quad + (\nu(S) - \nu(\{RR, RB, BR\}))u(0) \\ &= \nu(RB)u(8000) + (1 - \nu(RB) - (1 - \nu(\{RR, RB, BR\})))u(4000) \\ &\quad + (1 - \nu(\{RR, RB, BR\}))u(0) \\ &= \nu(RB)u(8000) + (1 - \nu(RB) - \bar{\nu}(BB))u(4000) + \bar{\nu}(BB)u(0).\end{aligned}$$

Similarly, the (Choquet) expected utilities of the other three bets are given by, respectively,

$$\begin{aligned}\mathbb{E}(\ell_{b_2}^{u\nu}) &= \nu(BR)u(8000) + (1 - \nu(BR) - \bar{\nu}(BB))u(4000) + \bar{\nu}(BB)u(0), \\ \mathbb{E}(\ell_{b_3}^{u\nu}) &= \nu(RB)u(8000) + (1 - \nu(RB) - \bar{\nu}(RR))u(4000) + \bar{\nu}(RR)u(0), \\ \mathbb{E}(\ell_{b_4}^{u\nu}) &= \nu(BR)u(8000) + (1 - \nu(BR) - \bar{\nu}(RR))u(4000) + \bar{\nu}(RR)u(0).\end{aligned}$$

Given the above equality constraints on any admissible capacity, it follows that the (Choquet) expected utilities of all four bets must be equal. Hence the only preference ordering that is admissible is the trivial one in which $b \sim^T b'$, for all b, b' in B . Thus the preference ordering in which b_2 is strictly preferred to b_1 as suggested by Machina, is inadmissible.

However, the trivial preference relation is also admissible for the probability $p = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$. Hence it follows from Theorem 1 that any inadmissible preference relation, including the one suggested by Machina, is not robustly inadmissible even restricting preferences to the smaller class of subjective expected utility maximizers.

In Appendix B we analyze the “50:51” experiment, the other main thought experiment presented in Machina (2009). We show the preference pattern Machina argues as being intuitively plausible for an ambiguity averse individual is not only inadmissible for CEU maximizers but it is also robustly inadmissible for SEU maximizers. However, since a coarsening of these preferences is admissible for some CEU maximizer, and we conjecture a result analogous to Theorem 1 can be established for the family of CEU preferences, this suggests this preference pattern would not be robustly “CEU-inadmissible.” \square

Example 4 (A Non-robust Test of CARA). The subjects are presumed to be expected utility maximizers that are either CARA or DARA. For this experiment we seek to challenge the hypothesis that the subjects are all CARA. We offer the subjects four bets involving the toss of a fair coin. So the canonical state space is $S = \{H, T\}$ and the payoffs of the bets are given by the following consequence matrix:

$$C = \begin{array}{c} \\ b_1 \\ b_2 \\ b_3 \\ b_4 \end{array} \begin{array}{cc} H & T \\ \left[\begin{array}{cc} 5 & 5 \\ 2 & 10 \\ 55 & 55 \\ 52 & 60 \end{array} \right] \end{array}.$$

Here the set of parameters \mathcal{A} are given by pairs (α, p^*) , where $\alpha > 0$ is the coefficient of absolute risk aversion for the Bernoulli utility function of the form $u(x) = -e^{-\alpha x}$ and $p^*(H) = \frac{1}{2}$.

Notice that the preference ordering $b_4 \succ b_3 \succ b_1 \succ b_2$ is inadmissible for the CARA model. However, for the admissible pair (α^*, p^*) in which

$$-e^{-5\alpha^*} = -\frac{1}{2}e^{-2\alpha^*} - \frac{1}{2}e^{-10\alpha^*}$$

we have $b_4 \sim^* b_3 \succ^* b_1 \sim^* b_2$, thus making $b_4 \succ b_3 \succ b_1 \succ b_2$ weakly admissible and hence not robustly inadmissible. \square

6. ROBUSTIFYING AN EXPERIMENT

The idea of robustification begins with a given non-robust thought experiment (C, \mathcal{A}) and a specific preference ordering \succsim that is inadmissible. The experimenter has a well-reasoned intuition that participants could display these preferences. The first task of robustification is to modify the experiment so that the hypothesized preference ordering is robustly inadmissible. A more ambitious objective is to do so in a way such that the modified experiment is robust while at the same time the hypothesized preference ordering remains inadmissible.

We begin by showing how the first task may be accomplished for a class of experiments, including the classic Ellsberg and Allais, involving two pairs of bets. Next, we discuss a participant based approach in which experiments are modified for each participant to guarantee that, if a participant's elicited answers imply preferences that are inadmissible, then those preferences are robustly inadmissible in that personalized experiment.

6.1. Robustification of the Four Bet Experiment. Important characteristics of preferences may be expressed in terms of consistency requirements. Examples include the sure-thing principle in subjective expected utility, the independence axiom in expected utility, and CARA in risk theory. Experiments designed to test these consistency requirements directly will often give rise to inadmissible preferences that are not robustly so. This may be explained as follows.

Suppose that in an experiment (C, \mathcal{A}) involving four bets, the following consistency property holds,

$$b_1 \succsim b_2 \iff b_3 \succsim b_4, \text{ for all } (u, p) \in \mathcal{A}.$$

Furthermore suppose that the experimenter has an intuition that participants could display the inadmissible preference pattern $b_1 \succ b_2$ and $b_4 \succ b_3$ as in the Ellsberg one-urn example from Section 2, and as in Examples 2 and 4. Since all these admit a pair $(u, p) \in \mathcal{A}$ that induces indifference between b_1 and b_2 (and hence, by the consistency property, between b_3 and b_4 , as well), this pattern is not robustly inadmissible.

Since any u is an order-preserving transformation of the consequences, we can exploit the monotonicity of the induced preferences to robustify experiments of this kind. Consider a perturbed consequence matrix $\hat{C} \neq C$, in which $C_{b_1} \gg \hat{C}_{b_1}$ and $\hat{C}_{\hat{b}_i} = C_{b_i}$ for $i = 2, 3, 4$.¹⁰ Notice that for any $(u, p) \in \mathcal{A}$, $\mathbb{E}(\ell_{b_1}^{up}) > \mathbb{E}(\ell_{\hat{b}_1}^{up})$. Hence monotonicity and the consistency property together imply

$$\hat{b}_1 \succsim \hat{b}_2 \implies b_1 \succ b_2 \iff b_3 \succ b_4 \iff \hat{b}_3 \succ \hat{b}_4, \text{ for all } (u, p) \in \mathcal{A}.$$

Thus for the inadmissible preference ordering $\hat{b}_1 \succ \hat{b}_2$ and $\hat{b}_4 \succ \hat{b}_3$, no coarsening is admissible. So by Theorem 1 the ordering is robustly inadmissible.

Returning to Example 4, consider the perturbed consequence matrix

$$\hat{C} = \begin{array}{c} H \quad T \\ \hat{b}_1 \left[\begin{array}{cc} 4 & 4 \end{array} \right] \\ \hat{b}_2 \left[\begin{array}{cc} 2 & 10 \end{array} \right] \\ \hat{b}_3 \left[\begin{array}{cc} 55 & 55 \end{array} \right] \\ \hat{b}_4 \left[\begin{array}{cc} 52 & 60 \end{array} \right] \end{array},$$

obtained by reducing the consequence of bet b_1 in each state from 5 to 4. The choice pattern $\hat{b}_1 \succ \hat{b}_2$ and $\hat{b}_4 \succ \hat{b}_3$ is now robustly inadmissible for CARA preferences.

Notice that there is a price to pay. In the original experiment, the pattern $b_2 \succ b_1$ and $b_3 \succ b_4$ was also inadmissible. In the new experiment the preference ordering $\hat{b}_3 \succ \hat{b}_4 \succ \hat{b}_2 \succ \hat{b}_1$ is admissible. That is, the new experiment is a ‘one-sided’ instead of a ‘two-sided’ test. The acceptability of this trade-off depends on the strength of the intuition that the preference pattern $b_1 \succ b_2$ and $b_4 \succ b_3$ is more plausible than the reverse. In the case of Example 4, CARA is likely to be rejected in favor of DARA but not in favor of IARA.

6.2. Participant based robustness. Every expected utility maximizing subject in an experiment in belief form (C, \mathcal{A}) can be taken to have a single subjective belief p . If p is an admissible belief, that is, $(u, p) \in \mathcal{A}$ for some utility function u , then that subject can be viewed as participating in a ‘personalized’ lottery based experiment. So for an experiment (C, \mathcal{A}) and each admissible belief p , let $\mathcal{A}^p = \{(u, p) : (u, p) \in$

¹⁰For any two vectors $x, y \in \mathbb{R}^n$ we write $x \gg y$ for the component-wise strict inequality, that is, $x \gg y \iff x_i > y_i$ for all i .

$\mathcal{A}\}$. The pair (C, \mathcal{A}^p) is an experiment in belief form that induces a lottery based experiment.

The next result states that an experiment in belief form is robust if and only if each admissible personalized lottery based experiment is robust.

Proposition 3. An experiment (C, \mathcal{A}) is robust if and only if the ‘personalized’ experiment (C, \mathcal{A}^p) is robust for every admissible belief p .

So far we have considered an experiment (C, \mathcal{A}) for which the experimenter selects a sample of participants from the population that is being studied. The robustness of this experiment guarantees that any inadmissible elicited preference ordering from any participant is robustly inadmissible. Consider an individual participant i . If the experimenter elicits a preference ordering \succsim^i over the bets, and it is inadmissible but weakly admissible in the experiment (C, \mathcal{A}) , then the experimenter may not have enough information about the preferences of the particular participant to rule out the possibility that \succsim^i is arising from a small perturbation in the expected utility maximizer’s beliefs.

To see this, we assume by way of contradiction that participant i is a subjective expected utility maximizer satisfying the theory being tested. This participant i has a belief p^i over S and a utility function u^i , which represent her preferences but which are not known. From the perspective of the experimenter, i participates in the trivial experiment (C, \mathcal{A}^i) where $\mathcal{A}^i = \{(u^i, p^i)\}$. Because she has just one belief, effectively this participant faces the lottery-based experiment $L_{(C, \mathcal{A}^i)}$.

The elicited preferences \succsim^i are inadmissible in $L_{(C, \mathcal{A}^i)}$ because they are inadmissible in (C, \mathcal{A}) . Importantly, although \succsim^i may be weakly admissible in (C, \mathcal{A}) , if the experimenter can establish that $L_{(C, \mathcal{A}^i)}$ is robust, then the experimenter is able to conclude that preferences \succsim^i are not due to a small perturbation in the subjective beliefs of the particular participant i .

The experimenter, however, does not know the parameter (u^i, p^i) of the (C, \mathcal{A}^i) and cannot determine these simply by the information given by the preferences \succsim^i . So the experimenter will want to gather additional information from the participant and econometrically estimate (u^i, p^i) from these. In an Ellsberg experiment where

there are only two consequence, one need only estimate the participant i 's beliefs p^i setting the range of u^i as $\{0,1\}$, then checking for the robustness of (C, \mathcal{A}^i) .

For the special case of an Ellsberg experiment Binmore et al. (2012) employed the following procedure for each participant i :

- (1) Estimate the implied beliefs p^i of each participant by iteratively adding and subtracting balls in the experiment, eliciting the participant i 's preferences, and stopping the iteration when i switches her strict preferences over two bets. Coupled with the contrapositive assumption of expected utility maximization, this switching (taken as representing indifferences) gives an estimate \hat{p}^i of p^i .
- (2) Use the information given by \hat{p}^i to robustify the experiment $L_{(C, (u, \hat{p}^i))}$ by means of a perturbation of the induced lotteries.
- (3) Elicit preferences in the robustified experiment.

We note that the Binmore et al. (2012) procedure has a modification in which the elicited preferences \succsim^i are given for the same experiment for all participants. For instance, first elicit the participant's preferences \succsim^i for all i . Then follow the Binmore et al. (2012) procedure and estimate \hat{p}^i . If \succsim^i is robustly inadmissible in $L_{(C, (u, \hat{p}^i))}$, then we can take it as a robust rejection that i is an expected utility maximizer.¹¹

7. CONCLUDING COMMENTS

Experiments like those proposed by Allais and Ellsberg have been influential in the development of alternatives to, and generalizations of, expected utility theory. However, it has often been suggested that the apparently anomalous results of these experiments may result from sensitivity to small errors in decisions or deviations between the perceptions of the subject and those assumed by the experimenter.

The central task of this paper has been to formulate a rigorous definition of robustness for experiments and for observed choices that are inadmissible for a class of preferences under consideration. Most commonly, this is the class of expected

¹¹Although primarily concerned with testing properties of dynamic preferences such as consequentialism and dynamic consistency, Dominiak et al. (2012) note that their design can also be seen as a more robust test of expected utility. They found a non-negligible fraction of their subjects who were classified as conforming to subjective expected utility theory in the static Ellsberg experiment exhibited behavior after the arrival of new information that was not consistent with standard (i.e. Bayesian) updating.

utility preferences, but the method is equally applicable to tests of such hypotheses as constant absolute risk aversion. The core result is that if inadmissible preferences can be made admissible by coarsening (that is, by one or more conversions of strict preference to indifference), then the inadmissibility is not robust.

APPENDIX A. PROOFS

Let B be the finite set of bets and $K \subseteq \mathbb{R}$ be a finite set of distinct consequences. Let L be any $B \times K$ (Markov or lottery) matrix: that is, each row L_b of L is non-negative and $\sum_{k=1}^K \ell_{bk} = 1$. With slight abuse of notation, let ℓ_b denote the lottery corresponding to each row L_b . That is, $\ell_b(c) = \ell_{bc}$ if $c \in K$ and $\ell_b(c) = 0$, if $c \notin K$.

Consider the set of states $S = K^B$. Fix the $B \times S$ consequence matrix C in which $c_{bs} = s_b \in K$ for all bs . Let p^L be the belief on S given by

$$p_s^L = \ell_1(c_{1s})\ell_2(c_{2s}) \dots \ell_B(c_{Bs}).$$

We shall use the following lemma.

Lemma 1. For any L and $c \in K$ we have $\ell_b^{ip^L}(c) = \ell_{bc}$, where i is the identity Bernoulli utility function.

Proof. Notice that for each $b \in B$ and each $c \in K$

$$\begin{aligned} \ell_b^{ip^L}(c) &= \sum_{s: c_{bs}=c} p_s^L \\ &= \ell_b(c) \left(\sum_{s: c_{bs}=c} \ell_1(c_{1s})\ell_2(c_{2s}) \dots \ell_{b-1}(c_{(b-1)s}) \ell_{b+1}(c_{(b+1)s}) \dots \ell_B(c_{Bs}) \right) \\ &= \ell_b(c) \prod_{b' \neq b} \left(\sum_{k \in K} \ell_{b'}(k) \right) \\ &= \ell_b(c) = \ell_{bc}, \end{aligned}$$

as required. □

Proof of Proposition 1. The proof is an immediate consequence of Lemma 1 where the assumption that parameters are non-degenerate in an experiment in belief form corresponds to condition (1) in the definition of a lottery based experiment. □

Proof of Proposition 2. The proposition simply translates definitions and the proof is an immediate consequence of these. \square

Proof of Corollary 1.1. Suppose \succsim is an inadmissible preference ordering in an experiment in belief form (C, \mathcal{A}) .

If-part : Suppose that \succsim is weakly admissible. Then, for some version (C', \mathcal{A}') , there is an admissible parameter $(u, p) \in \mathcal{A}'$ and a belief \hat{p} such that \succsim is represented by the inadmissible parameter $(u, ((1 - \varepsilon)p + \varepsilon\hat{p}))$ for all $\varepsilon \in (0, 1)$.

Let \succsim' be the preference ordering represented by the admissible parameter (u, p) in \mathcal{A}' . We will show that \succsim is finer than the admissible preference \succsim' . It suffices to show that $b \succ' b'$ implies $b \succ b'$. Since $b \succ' b'$, we know that $\mathbb{E}(\ell_b^{up}) > \mathbb{E}(\ell_{b'}^{up})$. By the linearity of expected utility, $\mathbb{E}(\ell_b^{u((1-\varepsilon)p+\varepsilon\hat{p})}) = (1 - \varepsilon)\mathbb{E}(\ell_b^{up}) + \varepsilon\mathbb{E}(\ell_b^{u\hat{p}})$. So when ε is close to zero, $\mathbb{E}(\ell_b^{u((1-\varepsilon)p+\varepsilon\hat{p})}) > \mathbb{E}(\ell_{b'}^{u((1-\varepsilon)p+\varepsilon\hat{p})})$. Thus, $b \succ b'$.

Only-if-part : For the converse, suppose that \succsim is finer than an admissible preference ordering. We will construct a version (C', \mathcal{A}') of (C, \mathcal{A}) for which \succsim is ε -admissible.

Let $K \subseteq \mathbb{R}$ be any finite set that includes all the elements of C . Let $S' = K^B$ and C' be the associated consequence matrix as in Lemma 1. For each belief p of S let q^p be the belief on S' satisfying $\ell_b^{ip} = \ell_b^{iq^p}$ for all b as constructed in Lemma 1. Let $\mathcal{A}' = \{(u, q^p) : (u, p) \in \mathcal{A}\}$. Clearly, (C', \mathcal{A}') is a version of (C, \mathcal{A}) . Further, a preference ordering on bets is weakly admissible in (C, \mathcal{A}) if and only if it is ε -admissible in (C', \mathcal{A}') .

We now fix our attention on the experiment (C', \mathcal{A}') . Since \succsim is finer than an admissible preference ordering \succsim' , the relation \succsim can be obtained from \succsim' by an iterative application of steps (1) and (2), below.

Let $\hat{\succsim}$ be a weakly admissible preference ordering. Pick an arbitrary bet b , and let $\hat{\succsim}^{+b}$ and $\hat{\succsim}^{-b}$ be the preference orderings defined as follows

- (1) For $\tilde{b}, b' \neq b$, we have $\tilde{b} \hat{\succ}^{+b} b'$ ($\tilde{b} \hat{\succ}^{-b} b'$) if and only if $\tilde{b} \hat{\succ} b'$.
- (2) For $b' \neq b$, we have
 - (a) $b' \hat{\succ} b$ implies $b' \hat{\succ}^{+b} b$ ($b' \hat{\succ}^{-b} b$).
 - (b) $b \hat{\succ} b'$ implies $b \hat{\succ}^{+b} b'$ ($b \hat{\succ}^{-b} b'$).
 - (c) $b' \hat{\sim} b$ implies $b \hat{\succ}^{+b} b'$ and $b' \hat{\succ}^{-b} b$.

It suffices to show that both \succsim^{+b} and \succsim^{-b} are weakly admissible.

We consider two cases: $\ell_b^{u^*p}$ is not degenerate and $\ell_b^{u^*p}$ is degenerate. Consider the first case. By the assumption of non-degenerate parameters, we have $c_{bs'}, c_{bs''} \in \underline{\ell}_b^{ip}$ and $u(c_{bs'}) > \mathbb{E}(\ell_b^{up}) > u(c_{bs''})$. Now consider the lottery profile ℓ^δ defined by

$$\ell_{b'}^\delta(c) = \begin{cases} \ell_b^{ip}(c_{bs'}) + \delta & \text{if } b' = b \text{ and } c = c_{bs'} , \\ \ell_b^{ip}(c_{bs''}) - \delta & \text{if } b' = b \text{ and } c = c_{bs''} , \\ \ell_{b'}^{ip}(c_{b's}) & \text{otherwise .} \end{cases}$$

Notice that for δ (positive or negative) close to zero ℓ^δ indeed is a lottery profile. Let p^δ be the belief over S' consistent with ℓ^δ given in Lemma 1. Clearly, (u^*, p^δ) , for $\delta > 0$ close enough to zero, generates \succsim^{+b} and the preference \succsim^{-b} is generated by (u^*, p^δ) for $\delta < 0$ close enough to zero, which implies that \succsim^{+b} and \succsim^{-b} are ε -admissible in the experiment (C', \mathcal{A}') , thus weakly admissible in (C, \mathcal{A}) .

We can apply a similar argument in the second case where $\ell_b^{u^*p}$ is degenerate, keeping in mind that by the assumption of non-degenerate parameters we have $\ell_b^{u^*p} \neq \ell_{b'}^{u^*p}$ for all $b' \neq b$. \square

Proof of Theorem 1. Theorem 1 is an equivalent form of Corollary 1.1 obtained by taking contrapositives. \square

Proof of Corollary 1.2. This proof follows from Theorem 1 and the definition of robust experiment. \square

Proof of Proposition 3. This follows from the definitions of robust experiment and personalized experiment. \square

APPENDIX B. THE 50:51 EXAMPLE FROM MACHINA (2009)

Recall the preference pattern in Example 3 that Machina hypothesized was intuitively plausible for an ambiguity averse individual. Although such a pattern was not admissible for any member of the larger family of Choquet expected utility (CEU) maximizers, we showed it was weakly-admissible and hence not robustly inadmissible for subjective expected utility (SEU) maximizers.

In this appendix we analyze the other main thought experiment presented in Machina (2009), which he dubbed the “50:51” example. We show the preference pattern he argues as being intuitively plausible for an ambiguity averse individual is not only inadmissible for CEU maximizers but it is also robustly inadmissible for SEU maximizers. However, since a coarsening of these preferences is admissible for some CEU maximizer we conjecture they would not be “robustly CEU-inadmissible.”

The subject is presented with a single urn containing 101 balls. Fifty balls are marked with either 1 or 2, the other fifty-one balls are marked with either 3 or 4. Each ball is equally likely to be drawn. There are four bets b_1, b_2, b_3, b_4 .

We take the sample space to be $S = \{1, 2, 3, 4\}$, where s is the event in which a ball marked with an s is drawn from the urn. The consequence matrix is

$$C^M = \begin{array}{c} \\ b_1 \\ b_2 \\ b_3 \\ b_4 \end{array} \begin{array}{cccc} 1 & 2 & 3 & 4 \\ \left[\begin{array}{cccc} \$8000 & \$8000 & \$4000 & \$4000 \\ \$8000 & \$4000 & \$8000 & \$4000 \\ \$12000 & \$8000 & \$4000 & \$0 \\ \$12000 & \$4000 & \$8000 & \$0 \end{array} \right]. \end{array}$$

An individual who prefers bets with known odds might well express the preference pattern $b_1 \succ b_2$ and $b_4 \succ b_3$.

As was the case in Example 3, the class of preferences Machina has in mind is Choquet expected utility. We take \mathcal{A}^M to be the set of pairs (u, ν) , where u is any Bernoulli utility function with $u(0) < u(4000) < u(8000) < u(12000)$ and ν is a capacity that along with its conjugate $\bar{\nu}$ satisfy the following ‘natural’ symmetry conditions along with inequalities reflecting the known aspects of the urn’s composition:

$$\begin{aligned} 0 < \nu(1) = \nu(2) &\leq \nu(3) = \nu(4), & \nu(\{1, 2\}) < \nu(\{3, 4\}), \\ 0 < \bar{\nu}(1) = \bar{\nu}(2) &\leq \bar{\nu}(3) = \bar{\nu}(4). \end{aligned}$$

We say a preference ordering \succsim is (CEU-)admissible in the experiment (C^M, \mathcal{A}^M) if it can be represented by some $(u, \nu) \in \mathcal{A}^M$. However, by calculations similar to the

ones detailed in Example 3, it follows that for any $(u, \nu) \in \mathcal{A}^M$:

$$\mathbb{E}(\ell_{b_1}^{u\nu}) \geq \mathbb{E}(\ell_{b_2}^{u\nu}) \iff \nu(\{1, 2\}) \geq \nu(\{1, 3\}) \iff \mathbb{E}(\ell_{b_3}^{u\nu}) \geq \mathbb{E}(\ell_{b_4}^{u\nu}).$$

Hence the preference pattern $b_1 \succ b_2$ and $b_4 \succ b_3$ is (CEU-)inadmissible. Moreover, it is robustly (SEU-)inadmissible since the only probability that satisfies the above equality and inequality restrictions is $p = (\frac{25}{101}, \frac{25}{101}, \frac{50}{202}, \frac{50}{202})$. Thus the only preference orderings that are (SEU-)admissible must have $b_2 \succ_{\text{SEU}} b_1$ and $b_4 \succ_{\text{SEU}} b_3$.

However, notice that for any $(u, \nu) \in \mathcal{A}^M$ with $\nu(\{1, 2\}) = \nu(\{1, 3\})$, the induced preference ordering \succsim' has $b_1 \sim' b_2$ and $b_3 \sim' b_4$. Since we conjecture a result analogous to Theorem 1 can be established for the larger family of Choquet expected utility preferences, this suggests the preference pattern $b_1 \succ b_2$ and $b_4 \succ b_3$ would not be robustly (CEU-)inadmissible. \square

REFERENCES

- Al-Najjar, N. I. and J. Weinstein (2009). The ambiguity aversion literature: a critical assessment. *Economics and Philosophy* 25(3), 249–284.
- Allais, M. (1953). Le comportement de l’homme rationnel devant le risque: critique des postulats et axiomes de l’école Américaine. *Econometrica* 21(4), 503–546.
- Binmore, K., L. Stewart, and A. Voorhoeve (2012). How much ambiguity aversion? Finding indifference between Ellsberg’s risky and ambiguous bets. *Journal of Risk and Uncertainty* 45, 215–238.
- Dominiak, A., P. Duersch, and J.-P. Lefort (2012). A dynamic Ellsberg urn experiment. *Games and Economic Behavior* 75(2), 625–638.
- Eichberger, J., S. Grant, and D. Kelsey (2007). Updating Choquet beliefs. *Journal of Mathematical Economics* 43, 888–899.
- Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Economics* 75(4), 643–669.
- Fishburn, P. and P. Wakker (1995). The invention of the independence condition for preferences. *Management Science* 41(7), 1130–1144.
- Grabiszewski, K. (2014). On the rejectability of the subjective expected utility theory. Mimeo.

- Harrison, G. W. (1994). Expected utility theory and the experimentalists. *Empirical Economics* 19(2), 223–253.
- Hey, J. (1995). Experimental investigations of errors in decision making under risk. *European Economic Review* 39, 633–640.
- Hey, J. and C. Orme (1994). Investigating generalizations of expected utility theory using experimental data. *Econometrica* 62, 1291–1326.
- Kadane, J. (1992). Healthy scepticism as an expected-utility explanation of the phenomena of Allais and Ellsberg. *Theory and Decision* 32, 57–64.
- Kahneman, D. and A. Tversky (1979). Prospect theory: an analysis of decision under risk. *Econometrica* 47(2), 263–291.
- Karni, E. and D. Schmeidler (1991). Utility theory with uncertainty. In W. Hildenbrand and H. Sonnenschein (Eds.), *Handbook of Mathematical Economics*, Volume IV, Chapter 33, pp. 1763–1831. Elsevier.
- Loomes, G. and R. Sugden (1982). Regret theory: an alternative theory of rational choice under uncertainty. *The Economic Journal* 92(368), 805–824.
- MacCrimmon, K. and S. Larsson (1979). Utility theory: axioms versus paradoxes. In M. Allais and O. Hagen (Eds.), *Expected Utility and the Allais Paradox*, Chapter 14, pp. 333–409. Reidel.
- Machina, M. (2009). Risk, ambiguity, and the rank-dependence axioms. *American Economic Review* 99(1), 385–392.
- Machina, M. and D. Schmeidler (1992). A more robust definition of subjective probability. *Econometrica* 60(4), 745–780.
- Markowitz, H. (1952). Portfolio selection. *Journal of Finance* 7(1), 77–91.
- Neugebauer, T. and U. Schmidt (2007). Testing expected Utility in the presence of errors. *Economic Journal* 117, 470–485.
- Quiggin, J. (1993). *Generalized Expected Utility Theory: The Rank Dependent Model*. Boston: Kluwer Academic Publishers.
- Raiffa, H. (1961). Risk, ambiguity, and the Savage axioms: comment. *Quarterly Journal of Economics* 75, 690–694.
- Savage, L. J. (1954). *The Foundations of Statistics*. John Wiley & Sons.
- Tversky, A. and D. Kahneman (1974). Judgment under uncertainty: heuristics and biases. *Science* 185, 1124–1131.

von Neumann, J. and O. Morgenstern (1944). *Theory of Games and Economic Behavior*. Princeton University Press.