Economics and Philosophy, **25** (2009) 303–334
doi:10.1017/S0266267109990265

COPING RATIONALLY WITH AMBIGUITY: ROBUSTNESS VERSUS AMBIGUITY-AVERSION

KLAUS NEHRING
University of California, Davis

Al-Najjar and Weinstein (2009) argue that the extant literature on ambiguity aversion is not successful in accounting for Ellsberg choices as rational responses to ambiguity. We concur, and propose that rational choice under ambiguity aims at robustness rather than avoidance of ambiguity. A central argument explains why robust choice is intrinsically context-dependent and legitimately violates standard choice consistency conditions. If choice consistency is forced, however, ambiguity-aversion emerges as a semi-rational response to ambiguity.

1. INTRODUCTION

The growth of the literature on decision making under ambiguity over the last two decades since the publication of Schmeidler's (1989, first version: 1982) seminal paper on "Subjective Probability and Expected Utility without Additivity" must seem remarkable to many outsiders, as it is to someone who has been involved in this literature since its earliest days. Is this yet another bubble, or does this growth rely on strong fundamentals? The critical assessment offered by Al-Najjar and Weinstein in this issue (henceforth ANW) appears to lean to the former view. Whether or not one ultimately sides with their conclusions, their contribution is to be welcomed by the actual and potential producers of that literature for its timely challenge to take a step back and review how sound its intellectual outlook and bearings really are.

ANW focus on the mainstream of the ambiguity literature which they characterize in terms of the view that Ellsberg choices are rational responses to ambiguity, to be explained by relaxing Savage's sure thing principle and adding an ambiguity-aversion postulate. To keep in mind the distinction between ambiguity and ambiguity-aversion, we will refer to this mainstream as the "taste-based approach to ambiguity" – ANW put their finger on a key ambiguity of the taste-based approach – is it meant to provide primarily descriptive or primarily normative models? While the centrality of a particular behavioural phenomenon (namely that exhibited by Ellsberg choices) indicates a primacy of the descriptive interpretation, ANW are right to insist that this interpretation has limited plausibility. If Ellsberg choices were really "just another behavioral anomaly", what would then justify maintaining strong rationality assumptions everywhere else in the model, and how could one reasonably motivate giving ambiguity-averse behaviour centre stage in applied economic models such as in the influential macroeconomic work of Hansen and Sargent (2001a, 2001b)?

Yet a normative interpretation of the taste-based approach has its own difficulties as well. ANW emphasize the difficulties in providing a convincing account of rational dynamic choice under taste-based ambiguity. Indeed, we would concur that these difficulties can be disconcerting and reveal substantial gaps in our understanding from both the descriptive and normative perspectives. Yet it is not clear whether their arguments establish a fatal blow to a view of ambiguity-averse behaviour as rational. One might, it would seem, equally argue the contrapositive: that if and since ambiguity-aversion is rational, difficulties in coping with ambiguity dynamically (both for the decision-maker and for the modeller) may be for real, and may require getting used to. And, surely, the usual understanding of beliefs, updating etc. by economists and game-theorists can hardly be the ultimate arbiter in this matter, since that understanding is so deeply steeped in probabilistic reasoning, and since it cannot be assumed a priori that introducing ambiguity will leave all our basic preconceptions untouched. To the contrary, ambiguity may well turn out to be a disruptive and unsettling intellectual innovation that may prove hard to fully digest. Indeed, the present paper argues that successfully coping with ambiguity requires the abandonment of some deeply held views of what rational decision-making must look like.

While we agree with ANW that the taste-based approach to ambiguity has not been successful in providing a convincing normative account, we would locate its failure elsewhere, namely in the failure to provide an account of Ellsberg choices *as rational responses to ambiguity*, the phrase "responses to ambiguity" being key. It is this *epistemic charge*, we would submit, to which the study of Ellsberg choices owes its allure and claim to sustained attention, and which raises it beyond the level of "just another behavioral anomaly". In a time of black swans, even educated common sense may be inclined to believe that there is something economically real to an "epistemically charged" notion of ambiguity.

The core of the ambiguity-aversion literature described by ANW consists of various preference models that weaken the sure thing principle in order to accommodate Ellsberg-like behaviour and obtain a generalized representation of preferences. This modelling approach suffers from a self-imposed two-fold handicap in providing an account of Ellsberg choices as rational responses to ambiguity.

First, rationality is present at best in a negative form, namely in a rejection of the normative force of a particular axiom. Admitting Ellsberg choices as normatively acceptable gives a friendly pat on the shoulder of an Ellsbergian decision maker, but otherwise does little if any detectable *work* in the modelling itself. In particular, these models provide little if any basis to distinguish between rationally motivated departures from probabilistic sophistication and arbitrary or "subrational" departures. For example, rather than representing a response to some underlying ambiguity, Ellsberg choices might simply reflect differential attitudes towards risk across sources of uncertainty or implicit failures of using subjective probabilities coherently. Indeed, the recent literature offers a number of models along these lines (see e.g. Chew and Sagi 2008, Ergin and Gul 2009, Tversky and Wakker 1995), dispensing with any epistemic *hinterwelt*.[1]

Second, the proper function and "official" purpose of a representation theorem is a mathematically compact and convenient redescription of an axiomatically circumscribed class of behaviours, not an explanation of such behaviour. The representation is often chosen to be "suggestive" of an underlying motivation or psychological story, but the cognitive content and value of such suggestions is in doubt, for it is evidently impossible to go from purely behavioural data to something beyond behaviour by the means of formal inference only.[2] The difficulty of providing substance to the suggestions proffered by a representation is discussed further in section 5 in the context of Gilboa and Schmeidler's (1989) "Minimum Expected Utility" (MEU) model.

As a result of these two features of the taste-based approach, there is a tendency for ambiguity to disappear behind ambiguity aversion itself, or even to collapse into it. This tendency is mirrored in an ambiguity and ambivalence of the ANW critique itself, as ANW appear to doubt the soundness not just of the ambiguity aversion literature as a set of preference models, but of the very notion of ambiguity as

1 "Hinterwelt" is Nietzsche's term for a "world behind the world" in *Thus Spoke Zarathustra*.

2 However, see Gul and Pesendorfer's "The Case for Mindless Economics" (2008) for a spirited defence of "suggestiveness" as an integral part of the contribution of representation theorems.

non-probabilistic uncertainty. This scepticism surfaces in a sentence such as "being ambiguity averse would then amount to being cautious about things that 'do not exist'"; similarly, ANW later speak of the "understandable . . . temptation to interpret experimental anomalies as an expression of agents' being unsure about the right prior". Putting it simply and starkly, the purpose of the present paper is to salvage ambiguity from its disappearance in the jaws of ambiguity aversion.

Specifically, we shall argue that it is possible to provide a normative account of decision making under ambiguity on which, first, ambiguity is epistemically distinct from subjective probability and reflects an agent's uncertainty about the right prior, and, second, ambiguity thus conceived mandates (and not merely permits) departures from Subjective Expected Utility.

As the point of departure, we follow Bewley (1986, 2002), Walley (1991) and others in assuming that ambiguity is expressed most directly as a failure of preference completeness rather than the Sure-Thing Principle respectively the Independence axiom. Viewed as a model onto itself, the resulting "Partial Expected Utility" (PEU) model suffers from none of the main critiques of ANW: it provides a clear epistemic notion of ambiguity, namely as a set of reasonable priors, and this is accompanied by a canonical notion of updating of preferences and beliefs (namely point-by-point updating of the set of priors) that escapes the dynamic inconsistency problems that arise with ambiguity aversion.[3]

Despite these patently attractive features, the PEU model has remained on the sidelines,[4] presumably in large part due to its descriptive incompleteness. While descriptive incompleteness cannot be a decisive concern from a normative point of view, that point of view raises its own conundrum: why is there ever an *advantage* to asserting incomplete rather than complete preferences? In particular, when it comes to making an actual choice, why not make this choice based on completing the given PEU? At worst, such completion may be arbitrary, but what would be wrong with that? To address this normative challenge, an account of rational choice given incomplete preference is needed that rules out at least some ways of completing the PEU as suboptimal and hence impermissible.

We submit that such an account is indeed possible as an account of rational *robust* choice. To put it in a slogan: *robustness (not ambiguity aversion) is the rational response to ambiguity*. As a matter of fact, a specific model of robust choice has been proposed by the author in the past; see Nehring (1991, 1992, 2000). It seems safe to say that this "Simultaneous Expected Utility" (SIMEU) model has failed to leave an indelible imprint

3 In line with this, ANW explicitly abstain from commenting on this model.
4 See Rigotti and Shannon (2005) for an exception.

on the extant literature. Rather than to discuss the particular features of that model or advertise its merits, the main purpose of the present paper is to distinguish and contrast the concepts of robustness and ambiguity aversion, and to prepare the ground for the reception of any theory of robust choice in the first place.

The need for such ground-clearing arises from at least two significant intellectual obstacles that appear to stand in the way from seeing such a theory straight into the eye. First, how is it possible at all to meaningfully select among non-comparable alternatives? Doesn't non-comparability *mean* that no further selection is possible?[5] Second, robust choice turns out to be *essentially context-dependent*, violating any of the classical choice conditions such as WARP or IIA. Since such conditions are frequently viewed as one of *the* hallmarks of rationality, it may be hard to grasp the very possibility of a conception of rational robust choice as envisaged here.

To address both of these issues, we will argue that robustness is an attribute not of an act viewed as a performance (an uncertain state of affairs brought about by the decision maker), but an attribute of the endorsement of the choice of an act from a particular set of alternatives; robustness is not "goodness" under another name. Once this difference in category is understood, the possibility of selecting among non-comparables as well as the essential context-dependence of such a selection emerge transparently.

In the final substantive section, we obtain ambiguity aversion as a cousin of robustness which results from trading in the Independence axiom (for choice) for Choice Consistency. This substitution leads to a "belief-based" account of ambiguity aversion as a semi-rational response to ambiguity. Specifically, it yields an "exact" lexicographic version of the MEU model in which the set of decision weights is equal to the set of reasonable priors.

The plan for the remainder of the paper is as follows. In section 2, we introduce the Basic Argument from Robustness in the context of dialogue concerning a version of Pascal's wager with multiple rivalrous gods. The general framework is introduced in section 3. Section 4 discusses the basic axiomatic desiderata for a theory of rational robust choice. Its central argument explains why robust choice is essentially context-dependent and legitimately violates standard choice consistency conditions. Section 5 describes an account of belief-based ambiguity aversion, and sketches a model with incomplete preferences that provides room for both ambiguity aversion and robustness. Section 6 concludes.

[5] A view of this kind is suggested for example by the literature on rationalizable choice-behaviour; in that literature, incomplete preferences serve to rationalize set-valued choice *qua* choice of their maximal elements; see for example Moulin (1985) for a concise summary.

2. PASCAL'S WAGER WITH IMPRECISE PROBABILITIES

Blaise faces a serious decision problem: should I lead a religious life or a secular one? The stakes are high: if God truly exists, I am promised to enter paradise as a faithful, and otherwise will be damned. If God does not exist, the stakes will be lower but still substantial: I will live in untruth, and I will miss out on some worldly pleasures. So, my payoff matrix in utiles is as follows:[6]

	God	No God
Secular Life	–10	0
Religious Life	+9	–1

It is hard to decide what to do, though, since the evidence is murky. While it seems quite possible that God exists – so, for example, it would be not unreasonable to assign a 10% probability to this event – there is little hard evidence to support this if any; so a very low probability, even a 0% probability, seems defensible too. So, depending on which reasonable probability I base my decision on, the religious or the secular life may come out ahead. But I do not think it would be sound to just follow my nose and pick – more or less arbitrarily – one particular probability number; rather, I should honestly acknowledge the lack of reliable evidence and suspend judgment among reasonable estimates. But where does this leave me? What, in the end, should I do?

After pondering this conundrum for a while, Blaise turns to his friend Denis for advice. Denis listens, and, taking into account Blaise's stern and rigorous temperament, offers the following suggestion: "My dear friend, Blaise, it is understandable that you cannot make up your mind. And you are right: if you are sufficiently ignorant and honest at the same time, expected utility comparisons are often of limited help. We all have this problem. But why don't you loosen up and lower your standards a little bit: after all, I don't think that, if you consult your heart, you are really completely indifferent or completely agnostic between the two options."

Blaise is relieved to hear this advice from a friend whose sincerity and wisdom he had always admired. "Yes, Denis, I do admit to being strongly attracted to the greater hope and the greater security that the religious life offers, notwithstanding the somewhat tenuous nature of the evidence for His existence. So I know what do, and will choose a life devoted to Him."

Multiple Rivalrous Gods One of the classic arguments against Pascal's Wager was put forward by Denis Diderot, who observed that "an Imam

[6] Assuming finite utility, this is a tamed version of the original argument in Pascal's *Penséés* 233. According to Hacking (1975), Pascal's wager is "the first well-understood contribution to decision theory". See Hajek (2008) for a concise review of the extensive philosophical literature.

could reason just as well". Thus, let us now retell this story in a multi-cultural, more precisely, in a multi-monotheistic setting, in which there are many (jealous) gods who make similar grand promises but demand personal allegiance.

In this version of the story, Blaise's payoff matrix has the following structure:

	God_i	God_j	No God
Secular Life	–10	–10	0
Religious Life according to Faith_i	+9	–11	–1
Religious Life according to Faith_i	–11	+9	–1

Moreover, the evidence for the existence of each god is broadly symmetric, so Blaise's beliefs are similarly symmetric, and there are m faiths to choose from. To fix ideas, assume in fact that Blaise is "completely ignorant" about which god exists if any, so that he deems any logically well-defined prior reasonably admissible and suspends judgment between them.

Again, Blaise is strongly attracted by the apparently high upside potential and small downside risk of leading a life of faith, and prefers the life according to any faith to a secular life.[7] But there is a rub: all of the faiths are equally attractive, so which faith to choose seems completely arbitrary. After some dithering and distress, Blaise hits upon a way out: why not choose my faith randomly? This way I end up with the desired decision, but in a way that is less arbitrary as it gives each competing alternative an equal chance. Not fully at ease, Blaise again turns to his sagacious friend Denis, and explains to him how he intends to make up his mind.

Denis ponders Blaise's solution, but becomes sceptical considering the multiplicity of faiths on offer. A quick calculation in fact reveals that randomizing is not such a good idea: if god i exists in fact, Blaise's expected utility is $\frac{1}{m}9 + \frac{m-1}{m}(-11) = -11 + \frac{20}{m}$. Since there are more than 20 faiths to choose from, this is less than –10, the utility of a secular life if a god exists. And since the utility of a secular life is also greater if no god exists, randomization between faiths is unambiguously inferior to the choice of a secular life.

When Denis shares his computation with Blaise, the latter is quite surprised by the result and exclaims: "So randomization does not seem a good idea after all!" But, a little later, Blaise begins to have second

[7] This, indeed, is what typical instances of the classical Hurwicz rule recommend under complete ignorance; these put some weight α on the worst possible outcome, and complementary weight $1 - \alpha$ on the best possible outcome. Faith_i is preferred to a secular life whenever $\alpha < 0.95$. Savage's (1951) "minimax regret" criterion supports this preference as well.

thoughts, and starts to question whether the problem really lies with the randomization. After all, Denis' argument takes him back to square one, and provides no guidance on which faith to choose. "That choice remains arbitrary, and – if I am honest – how can my choice avoid being random? I might, for instance, choose according to my preference for the color of the priests' robes, but then there are all sorts of 'tie breakers' that I might use, and the choice of a tie breaker might itself be random. And even if can convince myself that robe color is *the* tie breaker, how would this way of choosing be *better than* random?"

"In fact, as I keep thinking about it, Denis' observation wants to tell me something. If I am arbitrarily plunging for one god, I am cheating myself. I am really acting as if I assigned at least 5% probability to its existence, and strictly less to the existence of the other gods. But everything about my choice between the faiths is symmetric, so I would really be making my choice on an idly inflated hope. And to put that much weight on the hope for *each* of these gods is self-contradictory. Hence, to put that much weight on the hope for *any* of these gods is self-defeating. I thus realize that, on balance, I have no sound and convincing grounds to choose faith (in any particular god), and should thus opt for a secular life."

We take Blaise's reasoning to constitute a paradigmatic instance of reasoning from robustness, and as the central point of departure for any fully developed theory of rational robust choice. A formalization of this "Basic Argument from Robustness" is given in section 4. This is preceded by the presentation of the general framework within which an account of rational robust choice can be developed.

As a challenge question, we invite the reader to ponder what the Basic Argument from Robustness implies about the choice between leading a secular life and living according to a particular faith i^* (e.g. a state religion), assuming that the other religious choices have been rendered infeasible (e.g. by a repressive political regime). In particular, can one invoke arguments from choice consistency to establish the superiority of the secular life in this restricted choice set as well? We shall return to this question in section 4.4 below.

3. GENERAL FRAMEWORK

We consider a decision maker who accepts the key precepts of SEU theory, Independence and Transitivity, as normatively *prima facie* desirable but who – due to imprecision of his probabilistic beliefs – is not prepared to commit himself to a complete ranking of acts. To describe our point of departure formally, let

- S denote a finite set (the "state space"),

- $f \in \mathbb{R}^S$ denote "acts", that is mappings from states to consequences described in utility payoffs ("utiles").
- $\trianglerighteq$ denote a partial ordering (transitive, reflexive relation) over $\mathbb{R}^S$ with asymmetric component $\triangleright$.

The partial ordering $\trianglerighteq$ is understood to provide an exhaustive description of all betterness judgments that satisfy the Independence axiom the DM is willing to commit to. These can be thought of as his *likelihood-based preferences*.

A priori, this description allows the DM to have further material preferences that do not respect the Independence axiom such as those exhibited in the Ellsberg paradox. Whether there is *rational* room for such additional preferences reflecting non-likelihood-based betterness judgments is a different matter, and one of the key issues raised by the present paper.

The specification of acts in terms of utility consequences is for simplicity of exposition only, and could easily be derived in the Anscombe and Aumann (1963) framework or from the notion of utility-based mixtures proposed in Ghirardato *et al.* (2003).

The preference relation $\trianglerighteq$ is *independent* if, for all $f, g, h \in \mathbb{R}^S$ and $\alpha \in (0, 1)$, $f \trianglerighteq g$ if and only if $\alpha f + (1-\alpha)h \trianglerighteq \alpha g + (1-\alpha)h$; $\trianglerighteq$ is *monotone* if $f \trianglerighteq g$ whenever $f_s \geq g_s$ for all $s \in S$ and $f \triangleright g$ whenever $f_s > g_s$ for all $s \in S$; finally, $\trianglerighteq$ is *coherent* if it is transitive, reflexive, independent and monotone.

Let Δ^S denote the set of non-negative vectors in $\mathbb{R}^S$ summing up to one, the probability simplex on S; its elements will be called *priors*. For $\pi \in \mathbb{R}^S$, E_π is the expectation operator with respect to π, with $E_\pi f = \sum_{s \in S} f_s \pi_s$. Bewley (1986, 2002) and Walley (1991) have shown results that essentially imply the following.

Theorem 1 *$\trianglerighteq$ on $\mathbb{R}^S$ is coherent if and only if there exists a (unique) closed, convex set $\Pi \subseteq \Delta^S$ such that, for all $f, g \in \mathbb{R}^S$,*

(1) $$f \trianglerighteq g \text{ if and only if } E_\pi f \geq E_\pi g \text{ for all } \pi \in \Pi.$$

Non-comparability as Suspension of Judgment If neither $f \trianglerighteq g$ nor $g \trianglerighteq f$, we will say that f and g are *non-comparable*. Under the assumed exhaustive interpretation of $\trianglerighteq$, non-comparability of f and g reflects suspension of judgment between them.

Since $\trianglerighteq$ is supposed to provide an adequate basis to determine the DM's choice behaviour, for $\trianglerighteq$ itself to be "rationally adequate" (relative to any "evidence" at the DM's disposal), the "burden of proof" must arguably be placed on the suspension of judgment rather than its assertion. In particular, in contrast to what may be the appropriate assertoric standard for a mere observer, especially a scientifically conscientious one, absence of definite evidence is not, by itself, enough to motivate suspension of judgment. That takes serious doubt. Thus, one may take the expression "$f \trianglerighteq g$" to describe an assertion of preference that (by the DM's lights) is "reasonably mandated" by his evidence, and accordingly the negation of "$f \trianglerighteq g$" as an assertion of the converse strict preference for g over f as "reasonably permitted". According to Theorem 1, under coherence, the latter is in turn equivalent to the existence of a prior $\pi \in \Pi$ that assigns greater expected utility to g than to f. The priors in Π can thus be viewed as (subjectively) "reasonable" complete probability assessments between which judgment has been suspended. Rather congenially, Levi (1980) refers to the priors in Π as "standards of serious possibility", without however deriving this set from an incomplete preference relation as done here.

Subjectivity of Beliefs Importantly, in contrast to a number of recent contributions which include Π among the primitives of the model, $\triangleright$ and Π are assumed to be interpreted *subjectively* as beliefs rather than *objectively* as information. Only in rather special, limiting cases can beliefs be identified with information, as it stands to reason that in almost any concrete situation, the DM's hunches and vague but substantial common-sense understanding entail beliefs that are far more specific (associated with far smaller sets Π) than what is "implied" by the data.

A *choice set* X is a compact convex set of acts. Compactness is assumed, as usual, to ensure existence of an optimal choice. Convexity is based on the assumption that randomization among feasible acts is always feasible. To see this, note that the convex combination of two acts in utility space $\alpha f + (1-\alpha)g$, with $0 < \alpha < 1$, can be viewed as resulting from a choice of f with probability α, and a choice of g resulting with probability $1-\alpha$, since this randomization entails a conditional expected utility in state s of $\alpha f_s + (1-\alpha)g_s = (\alpha f + (1-\alpha)g)_s$.

The canonical inclusion of mixtures may seem demanding at first, but it is hard to see what an exclusion of mixtures through randomization as infeasible could really mean; note in particular that such randomization might be subjective and take place in the DM's head. Furthermore, randomization would almost always seem to be feasible at a practical level. Choice sets will frequently be described as the convex hull of a set Y of non-randomized acts, $X = coY$.

A "*decision problem under uncertainty*" (d.p.u.) is given as a pair $(X, \trianglerighteq)$ where X is compact convex, and $\trianglerighteq$ is coherent, or, equivalently, as a pair (X, Π), with Π a closed (hence compact) convex set of priors. Let $\mathfrak{D}_S$ denote the set of all such pairs for the given state space S.

To study choices as a systematic, possibly even rational response to ambiguity, we shall consider the DMs choices when faced with different choice sets and when equipped with different beliefs. For reasons that will become clear later, it will also be important to formally allow for the possibility of describing the same decision problem as embedded within different state spaces. Formally, the state space will therefore be considered to be a variable as well.

To flesh out the formal framework further to this purpose, let Ω be an infinite ground set of "ultimate states", and Σ an algebra[8] of subsets of Ω. A "state space" S is now simply a finite partition of Ω into infinite sets s_i, $S = \{s_1, \ldots, s_n\}$ whose cells are viewed as states.[9] Let $\mathfrak{S}$ denote the family of such state spaces. Each d.p.u. will be assumed to be presented to the DM within the frame of a particular state-space $S \in S$. A *choice rule* C maps d.p.u.s $(X, \trianglerighteq) \in \mathfrak{D} := \cup_{S \in \mathfrak{S}} \mathfrak{D}$ to non-empty subsets of feasible acts $C(X, \trianglerighteq) \subseteq X$.

One basic requirement of rational choice under incompleteness is Admissibility. The act f is *admissible* ($f \in \mathcal{A}(X, \trianglerighteq)$) if there does not exist another feasible act g such that $g \rhd f$.

Axiom 2 (*Admissibility*) *For all* $(X, \trianglerighteq) \in \mathfrak{D} : C(X, \trianglerighteq) \subseteq \mathcal{A}(X, \trianglerighteq)$.

In the future, we will generally omit explicit universal quantifiers over d.p.u.s.

A number of choice rules considered in the literature satisfy only a weak form of Admissibility. To state it, say that f is strongly preferred to g ($f \rhd\rhd g$) if $f - h \trianglerighteq g$ for some $h \ll 0$ (i.e. that is strictly negative in every state). Hence say that the act f is *weakly admissible* ($f \in \mathcal{A}_{weak}(X, \trianglerighteq)$) if there does not exist another feasible act g such that $g \rhd\rhd f$. Note that if $\trianglerighteq$ is complete (respectively if Π is a singleton), Admissibility and Weak Admissibility coincide and imply SEU maximization.

4. ROBUSTNESS

Consider now a decision maker whose material preferences are likelihood-based, i.e. are exhaustively given by the PEU $\trianglerighteq$. This situation may arise because the DM normatively accepts the Independence axiom for material preferences, or may simply be due to an abstention from further material preferences judgments that would reflect ambiguity attitudes. How should such a DM choose? Is optimality tantamount to Admissibility, or does it demand more?

Well, we take this question to be already answered by our discussion of Pascal's Wager, which we take to demonstrate that *choice rationality* alone mandates more than Admissibility.[10] Not all admissible acts are

[8] An algebra is a family of sets that is closed under complementation and union and contains the universal event Ω.

[9] The requirement that each s_i be infinite does not play a role in the arguments presented here but is required for the full axiomatic characterization of some/many of the choice rules discussed below.

[10] The word "demonstrate" may ring preposterous in many reader's ears; but we only claim that an ideal discussion of this choice situation would deliver such a demonstration, while the actual discussion offered here is likely to fall short.

on par; some are "more choice-worthy" than others. What makes these acts more choice-worthy? Can one account for this choice-worthiness differential in terms of some unified *rationale*? Presumably, the choice-worthiness differential extends to other, less symmetric situations, and an understanding of it in terms of some unified rationale would help understanding and spelling out this differential more generally.

We submit that "robustness" can furnish such a unified rationale. This does not mean that one can rely on a well-defined, pre-existing notion of robustness that could be invoked more or less from the shelf to ground or define what choice rationality (optimality) is. Rather, we merely claim that the pertinent rationality conditions (as encoded in choice axioms) can be explicated in terms that tie into common intuitive notions of robustness. Or, at least, that such an explication should be possible for some (in particular: for the right) model of rationally robust choice.

On a broad, pre-formal understanding, we think of a choice as *robust* if it avoids one-sidedness and arbitrariness, if it integrates multiple viewpoints and gives all viewpoints "their due" – the "conflicting viewpoints" being identified here naturally with the different reasonable priors among which the DM has suspended judgment.

The two key considerations in the discussion of Pascal's Wager can be understood in terms of robustness. First, the argument for symmetry is clearly an argument for avoiding arbitrariness, namely for treating "informationally equivalent" choices equally. The argument for randomization can also be motivated by a desire to minimize arbitrariness. It presumes the existence of multiple competing, putatively optimal alternatives; committing to choose a particular one of them would break the symmetry of this equi-optimality; randomization among them ensures that this symmetry breaking occurs in a controlled, even-handed way.

4.1 Formalizing the basic argument from robustness

Formalizing the Basic Argument from Robustness requires a bit of technical machinery. Given $S, T \in \mathfrak{S}$ of equal cardinality, let φ be a bijection between S and T, and $\Phi : \mathbb{R}^S \to \mathbb{R}^T$ the associated linear transformation given by $\Phi(f)_t = f_{\varphi^{-1}(t)}$ for all $t \in T$. Here, φ maps states in S to states in T, and Φ transforms the associated objects (acts as payoff-vectors and priors as probability vectors) accordingly. Φ maps the feasible set X in $\mathbb{R}^S$ to the feasible set $\Phi\,(X := \{\Phi\,(f) : f \in X\}$, and the set of priors Π in Δ^S to the set of priors $\Phi\,(\Pi) := \{\Phi(\pi) : \pi \in \Pi\}$ in Δ^T. Likewise, the image of $\trianglerighteq$ under Φ in T denoted by $\Phi(\trianglerighteq)$ is given by

$$f\,\Phi\,(\trianglerighteq)\,g \quad \text{if and only if} \quad \Phi^{-1}\,(f) \trianglerighteq \Phi^{-1}\,(g) \text{ for } f, g \in \mathbb{R}^T.$$

X is symmetric under Φ if $\Phi(X) = X$; let $\mathfrak{G}_X$ denote the set of transformations Φ under which X is symmetric, the "symmetry group" of X. In like manner, let $\mathfrak{G}_{\trianglerighteq}$ and $\mathfrak{G}_{\Pi}$ denote the symmetry groups of $\trianglerighteq$ and Π, respectively; $\mathfrak{G}_{\trianglerighteq}$ and $\mathfrak{G}_{\Pi}$ are in fact easily seen to be equal to each other. The group $\mathfrak{G}_X \cap \mathfrak{G}_{\trianglerighteq}$ describes all transformations under which the d.p.u. $(X, \trianglerighteq)$ is symmetric as a whole. A real vector f in $\mathbb{R}^S$ is *invariant* in the d.p.u. $(X, \trianglerighteq)$ if $\Phi(f) = f$ for all $\Phi \in \mathfrak{G}_X \cap \mathfrak{G}_{\trianglerighteq}$. (This concludes the technically most tedious passage of the paper.)

A fundamental requirement of robustness is Invariance under state transformations Φ. This reflects the assumption that the decision problem is fully described by the feasible payoffs and the partial order $\trianglerighteq$, and that the state space S describing the DM's domain of uncertainty does not matter as such. The "full description" assumption rules out hidden relevant factors, first among them beliefs or any belief-like attitudes that might be decision-relevant but are not captured by $\trianglerighteq$. Thus, Invariance can be expected to hold only if $\trianglerighteq$, respectively Π, reflects the DM's exhaustively specified beliefs rather than some "objective information". As discussed above, only in very special cases will beliefs coincide with objective information. In addition, irrelevance of the state space rules out, for example, a taste-based "source preference" for responding differently to the same imprecision in beliefs.

Axiom 3 (*Invariance*) *For any $S, T \in \mathfrak{S}$ of equal cardinality and any state transformation $\Phi : \mathbb{R}^S \to \mathbb{R}^T$,*

$$C(\Phi(X), \Phi(\trianglerighteq)) = \Phi(C(X, \trianglerighteq)).$$

Since the set of priors associated with $\Phi(\trianglerighteq)$ is given by $\Phi(\Pi)$, Invariance can be restated in terms of sets of priors by replacing the symbol "$\trianglerighteq$" by the symbol "Π".

Invariance implies in particular the following Symmetry axiom that requires symmetries of the d.p.u. to be mirrored in symmetries of the choice. Symmetry was at the heart of our discussion of Pascal's Wager.

Axiom 4 (*Symmetry*) *For any $\Phi \in \mathfrak{G}_X \cap \mathfrak{G}_{\trianglerighteq}$,*

$$C(X, \trianglerighteq) = \Phi(C(X, \trianglerighteq)).$$

The second key ingredient of the argument in Pascal's Wager, the preservation of optimality under randomization, is captured by the following Randomization axiom.

Axiom 5 (*Randomization*) *$C(X, \trianglerighteq)$ is convex.*

Symmetry and Randomization together imply tight constraints on optimal choices. These are described in the following Proposition 6 which states that all optimal choices maximize expected utility with respect to the same invariant prior, and there is some optimal act that is invariant as well. Proposition 6 is proved in the Appendix.

Proposition 6 *If C satisfies Weak Admissibility, Symmetry and Randomization, there exist an invariant prior π^* and an invariant feasible act f^* such that*

i) for all $f \in X : E_{\pi^*} f^* \geq E_{\pi^*} f$, and
ii) for all $f \in C(X, \trianglerighteq)$, $E_{\pi^*} f^* = E_{\pi^*} f$.

Application to Pascal's Wager As an application, return to the discussion Pascal's Wager with rivalrous gods in section 2. The state space S has $m+1$ states, with state s_i corresponding to the existence of god$_i$, and state s_{m+1} denoting the non-existence of any god. Let $\mathfrak{G}$ denote the set of transformations leaving the $m+1$-st coordinate invariant. $\mathfrak{G}$ describes "symmetry with respect to gods", but not necessarily with respect to the existence or non-existence of some god in the first place. If beliefs are symmetric with respect to gods, i.e. if $\mathfrak{G} \subseteq \mathfrak{G}_\Pi$, then Proposition 6 can be applied since the feasible set X is symmetric under all $\Phi \in \mathfrak{G}$. Take an invariant prior $\pi^* = (\beta, \ldots, \beta, 1-m\beta)$ from Proposition 6. If $m > 20$, $\beta < 0.05$; the invariant act "secular life" then uniquely maximizes expected utility with respect to π^*, and is therefore the uniquely optimal choice.

The Principle of Insufficient Reason and its Critique Perhaps the most straightforward way to try to achieve robustness is to maximize SEU relative to some "compromise prior" $\rho(\Pi)$. By considering choice sets X with appropriate symmetries, it is easily seen that such a prior selection ρ is consistent with the assumptions of Proposition 6 if and only if $\rho(\Pi)$ is invariant under all symmetries of Π (i.e. $\Phi(\rho(\Pi)) = \rho(\Pi)$ for all $\Phi \in \mathfrak{G}_\Pi$). An example popular in statistics is the prior in Π that maximizes Shannon entropy (see e.g. Jaynes 1957, 2003). Another attractive possibility is the Steiner point of Π (a kind of geometric centre point of Π) as used by Gajdos *et al.* (2008). Such invariance-respecting selection functions ρ can be viewed as capturing an extended "Principle of Insufficient Reason". The classical Principle of Insufficient Reason can be seen as instance of this assigning a uniform prior under complete ignorance (with $\Pi = \Delta^S$).

Yet this straightforward approach to robustness does not really work. Indeed, dating back to at least von Kries (1886), the classical Principle of Insufficient Reason itself has come under heavy and justified criticism as overly reliant on a particular specification of the relevant state space. This problem is not really an unfortunate side-effect of the prior-selection approach that one could hope to dodge or mitigate, it shows that it is fundamentally on the wrong track. The prior-selection approach attempts to represent a state of partial or complete ignorance by a unique prior; but a unique prior can never represent a state of genuine ignorance, and therefore must misrepresent the true structure of underlying beliefs.[11]

4.2 Redescription Invariance

To ensure protection against such unwarranted description sensitivity, the following axiom requires invariance under appropriate alternative representations of the decision problem. Specifically, it will frequently be the case that the original specification of the state space is finer than what is needed to describe the feasible choice set. For example, in Pascal's Wager with a state religion as in the concluding paragraph of section 2, the choice between a secular life and the state-approved faith$_{i^*}$ can be described in terms of the state space $T = \{\text{God} = i^*, \text{God} \neq i^*, \text{No God}\}$. It stands to reason that the optimal decision should not depend on whether the uncertainty is described in the original state space S or in the coarsened space T. Ruling out such arbitrariness is clearly fundamental for achieving robustness.

This is formally captured by the following requirement of Redescription Invariance. For any coarsening T of S, one can view $\mathbb{R}^T$ as the subset of vectors (acts) in $\mathbb{R}^S$ that are constant within each cell of T. Thus, a choice set $X \subseteq \mathbb{R}^S$ can be described within the state space T if and only if $X \subseteq \mathbb{R}^T$.

For any PEU $\trianglerighteq$ on $\mathbb{R}^S$, let $\trianglerighteq^T$ given by the restriction of $\trianglerighteq$ to acts in $\mathbb{R}^T$.[12] Likewise, given underlying beliefs Π over S, the DM's marginal beliefs are given by the set of priors $\Pi^T := \{\pi^T : \pi \in \Pi\} \subseteq \Delta^S$, where $\pi_{|T}$ is the vector of marginal probabilities associated with the prior π.[13] It is not difficult to show that Π^T is the multi-prior representation of $\mathbb{R}^T$.

Axiom 7 (*Redescription Invariance*)

If $X \subseteq \mathbb{R}^T$, $C(X, \trianglerighteq) = C(X, \trianglerighteq^T)$, *and, equivalently*, $C(X, \Pi) = C(X, \Pi^T)$.

Redescription Invariance strengthens the conclusion of Proposition 6 significantly, since one can now exploit the possibly richer set of symmetries that emerge at coarser descriptions. To illustrate, let $\Pi = \Delta^S$ reflecting complete ignorance, and let $X = co(\{1_A, 1_B\})$, where 1_E is the indicator function of event $E \subseteq S$ paying off 1 utile in event E and 0 otherwise. Clearly $X \subseteq \mathbb{R}^T$ with $T = \{A \cap B, A \backslash B, B \backslash A, A^c \cap B^c\}$. Whenever both $A \backslash B \neq \emptyset$ and $B \backslash A \neq \emptyset$, X and $\Pi^T = \Delta^T$ are symmetric under the permutation of $A \backslash B$ and $B \backslash A$.

[11] To be sure, we don't mean to banish the use of symmetry or entropy arguments to come up with precise priors in certain situations. But these would need to be based on an appropriate balance of the evidence, not ignorance in the form of prior imprecision.

[12] More rigorously, $\trianglerighteq^T = \trianglerighteq \cap (\mathbb{R}^T \times \mathbb{R}^T)$.

[13] With $T = \{t_j\}$ and $t_j = \{s_{i,j}\}_{i=1,..,n_j}$, $\pi^T_{t_j} = \Sigma_{i=1,..,n_j} \pi_{s_{i,j}}$.

Thus, by Proposition 6, it is optimal to toss a coin between the two bets; that is, $\frac{1}{2}1_A + \frac{1}{2}1_B \in C(X, \Delta^S)$. In other words, neither A nor B are revealed to be more likely than the other, however "small" or "large" these events may be, as long as they are mutually non-inclusive. This conclusion is striking (and strikingly anti-probabilistic!), but sound as a reflection of the assumed complete ignorance of the DM.

A similar result is obtained when the DM is faced with a proper scoring rule, for example the logarithmic one.[14] In this case, when asked to announce a (scored) probability distribution p over some partition $T = \{t_j\}$ of S, the DM will provide the uniform distribution over T as his uniquely optimal estimate, whatever T he is faced with. Again, this is a striking yet convincing conclusion. Since in the context of a proper scoring rule, the prior supporting the optimal estimate p is simply p, it implies that the priors supporting the optimal choice must depend on the choice set X.

To accommodate the context-dependence of the supporting prior, Proposition 6 requires its symmetry only when the choice set X is appropriately symmetric as well. Proposition 6 can thus be viewed as stating a "Decision-theoretic Principle of Insufficient Reason".

4.3 Enter Independence

Does robustness entail further axiomatic restrictions on optimal choices? We submit that the answer is positive, and that robust choices need to satisfy an Independence axiom formulated for choice functions.

Axiom 8 (*Independence*) $C(\alpha\{x\} + (1-\alpha)X, \trianglerighteq) = \alpha\{x\} + (1-\alpha)C(X, \trianglerighteq)$.[15]

Independence has a compelling robustness motivation, since it can be read as an implication of the invariance requirement that the structure of choice should mirror the structure of preference. Specifically, let $\theta : \mathbb{R}^S \to \mathbb{R}^S$ be given as the map $z \mapsto \theta(z) = \alpha x + (1-\alpha)z$, which induces a one-to-one mapping between X and $\alpha\{x\} + (1-\alpha)X$. Since $\trianglerighteq$ satisfies the Independence axiom for preferences, one has, for all $x \in X$,

$$x \trianglerighteq y \text{ iff } \theta(x) \trianglerighteq \theta(y).$$

The following Lemma due to Chernoff (1954) illuminates the content of the Independence axiom; in particular, it shows that under Independence, utility levels do not matter.

[14] Here $X = \overline{co}(\{x^{p,T}\})$, where $x^{p,T}$ is the payoff vector resulting from an announcement of $p \in \Delta^T$ as the probability distribution over T; the logarithmic scoring rule, for example, is defined by $x_s^{p,T} = \ln p_{t_j}$ if $s \in t_j$.

[15] The Minkowski mixture $\alpha\{x\} + (1-\alpha)X$ denotes the set $\{\alpha x + (1-\alpha)y : y \in X\}$. Likewise, the sum $X + x$ denotes the set $\{y + x : y \in X\}$.

Lemma **9** *Independence is equivalent to the conjunction of the following two conditions:*

i) $C(X+x, \trianglerighteq) = C(X, \trianglerighteq) + x$ *(Translation Invariance), and*
ii) $C(\alpha X, \trianglerighteq) = \alpha C(X, \trianglerighteq)$ *(Homotheticity).*

A Direct Violation of Choice Consistency Remarkably, when added to other axioms, Independence comes in direct conflict with standard context-independent choice consistency.

Axiom 10 (*Choice Consistency*)

i) $Y \subseteq X$ *implies* $C(Y, \trianglerighteq) \supseteq C(X, \trianglerighteq) \cap Y$;
ii) $C(X, \trianglerighteq) \subseteq Y \subseteq X$ *implies* $C(Y, \trianglerighteq) = C(X, \trianglerighteq)$.

The first condition says that an optimal choice that is feasible in a smaller set remains optimal; the second says that the removal of unchosen alternatives does not change the set of optimal choices. These conditions are sometimes referred to as the Chernoff and Aizerman axioms, respectively. In standard settings in which the domain of the choice function consists of all non-empty finite subsets of some universal set of alternatives, their conjunction is equivalent to Plott's (1973) classical Path-Independence condition.

Choice Consistency implies the following version of the Weak Axiom of Revealed Preference.

Axiom 11 (*WARP*) $X, Y \supseteq \{f, g\}$ *and* $C(X) = \{f\}$ *implies* $g \notin C(Y)$.

Lemma 12 *Choice Consistency implies WARP.*[16]

The following example exhibits the direct conflict of the assumed robustness axioms with Choice Consistency.

Example 13 *Let S have three states, and assume complete ignorance* $(\Pi = \Delta^S)$. *Let* $f = (0, 0, 0)$ *and* $g = (3, -2-2)$, *and consider choices from the following three choice sets:*

1. $X_1 = co(\{f, g, (-2, 3, -2), (-2, -2, 3)\})$;
2. $X_2 = co(\{(-3, 2, 2), (0, 0, 0), (2, -3, -3)\})$; *and*
3. $X_3 = co(\{f, g, (5, -5, -5)\})$.

Then $C(X_1, \Delta^S) = \{f\}$ *and* $C(X_3, \Delta^S) = \{g\}$, *in violation of Choice Consistency.*

[16] Indeed, if $C(X) = \{f\}$, then $C(co\{f, g\}) = \{f\}$ by i) and ii) of Choice Consistency. Using i) again, by contraposition $g \notin C(Y)$.

Indeed, by the Basic Robustness Argument (Proposition 6),

$$C(X_1, \Delta^S) = \{f\}.$$

Likewise, by the Basic Robustness Argument and Redescription Invariance, $C(X_2, \Delta^S) = \{(0, 0, 0)\}$.

Now $X_3 = X_2 + (3, -2, -3)$. Hence by Translation Invariance/ Independence,

$$C(X_3, \Delta^S) = \{g\},$$

in violation of WARP, and thus, by Lemma 12, in violation of Choice Consistency.

4.4 In a pickle?

The recognition of a conflict between Independence and Choice Consistency is in fact nothing new, and appears to have contributed significantly to a loss of interest in the classical literature on Complete Ignorance of the 1950s. Probably representatively for that literature, Arrow (1960: 72) concluded that a rational solution to complete ignorance problems is impossible: "Perhaps the most nearly definite statement is that of Milnor (1954) who showed in effect that every proposed ordering principle contradicts at least one reasonable axiom."

The new twist in the present exposition is to constructively derive from the putative rationality axioms pairs of choices that directly contradict Choice Consistency. Thus, if You, the decision maker, maintain a commitment to these choices, You have *already* committed yourself to violating Choice Consistency, and cannot, *on pain of logical incoherence*, maintain a commitment to Choice Consistency as a normative desideratum on rational choice. Of course, You might still wish, for example, that life would be simpler, that every axiom that looks attractive can be taken on board without a second look etc., but here You are at a cross roads.

So, let us assume that you, dear reader, are in exactly this position: that you found the argument for the optimality judgments in Example 13 convincing, but are now disturbed by their conflict with Choice Consistency. Is there any way you can extricate yourself out of this pickle? Your first instinct may be to step back and retrace the earlier reasoning and see whether you can find a flaw, whether you have been taken in too quickly by some of the rhetorics; maybe you find a hidden assumption that you were tricked into accepting when you shouldn't have been. But the reasoning was pretty straightforward and transparent, and you had already considered it quite carefully. So, suppose you still do not find such a flaw, and still cannot see anything wrong with the choices in Example 13. Now, you will need to take a second look at Choice Consistency: is Choice

Consistency really as categorical and universal a rationality requirement as it may have appeared?

At this point, the following reasoning may be helpful. Choice Consistency conditions are normatively compelling if they concern choices or judgments of which alternative is "best", where "goodness" is understood as some attribute of an alternative *qua* the (possibly uncertain) future history that is brought about by it. An apple is an apple is an apple, whether chosen from a basket with a pear and a banana, or a pear and a mango. You cannot consistently deem the apple "best" in the first basket, but the pear "best" in the second.[17]

Yet, "choice based on robustness" does not express or constitute a judgment of context-independent bestness. The need for robustness considerations was based on the very premise that You have run out of usable betterness judgments, that $\trianglerighteq$ is all the material betterness that You can go on. Moreover, the "deductive closure" conditions of Transitivity and Independence that define the coherence of $\trianglerighteq$ are "designed" to keep track of all the betterness judgments that can be legitimately derived from others that You have directly endorsed. So it would be strange if, based on certain abstract considerations encoded in axioms, one could deduce something further about the betterness of alternatives beyond what is given by $\trianglerighteq$, i.e. beyond Admissibility.

It would thus be mistaken to read the proposition expressed by the symbols "$x \in C(X, \trianglerighteq)$" as "$x$ is best in X given $\trianglerighteq$", where "best" is understood in a context-independent manner, as an attribute of the chosen act X viewed as a performance (a state of affairs in the world brought about by the decision maker). Instead, the proposition expressed by the symbols "$x \in C(X, \trianglerighteq)$" is to be read here as about the *choice* of that act *from the set of alternatives* X, as in "it is best to *choose* X from $(X, \trianglerighteq)$". Yet the word "choose", while suggestive, still suffers from a certain logical opaqueness. As a more transparent and cogent substitute we offer the following: "The endorsement of X as the most choice-worthy act in X given $\trianglerighteq$ is the best supported". "Best supported", in turn, can be fleshed out as "most comprehensively", "most evenly" supported, and, in this sense "most robust". Robustness, then, is not an attribute of acts as a performance, but of an endorsement. The best, most robust endorsement, and thus the one that counts, simply *tells* the agent what he should do in a particular choice situation, without pronouncing on the value of x as a state of affairs. Since the rational selection of choices *qua* endorsements takes place on a different, "second-order" plane than the "first-order" plane of

17 This perspective can accommodate apparent context effects if those context effects can be explained as modifying the real future history that result from a given alternative, e.g. by the incorporation of a psychological context-dependendent consequence such as regret over what might have been chosen.

states of affairs, it relies on and has access to different grounds than first-order judgments of material betterness. In particular, there is no puzzle why second-order robustness judgments might be possible even when material betterness judgments have been exhausted.

Moreover, the *essential context-dependence* of robustness considerations is very natural and transparent when placed at a second-order level. Understood as a second-order axiom, Symmetry, for example, says that if one accepts the statement: "the endorsement of x as most choice-worthy in $(X, \trianglerighteq)$ is best supported", one is bound to equally accept, for any symmetry-preserving transformation Φ, the statement: "the endorsement of $\Phi(x)$ as most choice-worthy in $(X, \trianglerighteq)$ is best supported", since any consideration that might be adduced in favour of the former has a Φ-image in favour of the latter.

The essential context-dependence is also critical to the construal of the Basic Robustness Argument in Pascal's Wager as an argument from self-defeat. For, clearly, to claim that privileging of one faith over the others is self-defeating, one must rely on their actual feasibility; it is that actual feasibility (rather than mere conceivability), which creates the asymmetry between the choice of any particular faith and that of a secular life that the argument from self-defeat exploits.

If, based on some misconceived attachment to Choice Consistency, one was to then extrapolate this context-dependently inferred superiority of the secular option to a choice between it and one particular faith, one would clearly overstep the authority bestowed by the original argument. Or, to put this differently, an appeal to Choice Consistency is normatively incoherent here: on the one hand, the putative force of Choice Consistency presumably rests on the assumption that the comparative choice between two acts is materially the same, no matter within which set it is applied. But this assumption of material sameness is undercut by the essential context-dependence of the Basic Robustness Argument. Note the difference here between invoking Choice Consistency and Independence. While both imply further optimality judgments, the latter, in contrast to the former, does so on the basis of substantively new considerations, which, moreover, are entirely orthogonal to those involved in the Basic Robustness Argument itself.

4.5 Towards a full characterization of robust choice

The axioms invoked so far, including Independence, still leave the choice in many problems undetermined, and there is a fairly wide range of more or less sensible choice rules consistent with them. The issue whether there is a unique right rule, and which it is, is beyond the scope of this paper. It has in fact been addressed in an earlier paper (Nehring 2000), although we make no claim to this being the last word. In fact, while we are hopeful that the "Simultaneous Expected Utility" (SIMEU) rule proposed there will turn out to agree in overall character with the right robust choice rule (yet to be found), the SIMEU rule itself is likely to prove too simplistic.[18] In spite of these reservations, we shall briefly exhibit the SIMEU rule here to give concreteness and tangibility to the idea that robust choice may be well-defined and essentially determinate in any sufficiently well-behaved decision problem, not just in ones with a high degree of symmetry.

The SIMEU rule is inspired by a rather more specific notion of robustness than has been invoked so far. According to this notion, the choice of an act is deemed robust if it is as satisfactory as possible from all points of view between which judgment has been suspended relative to the range of acts that might have alternatively (admissibly) be chosen. To define the SIMEU rule formally, a special role is given to the set of extreme points $\mathcal{E}(\Pi)$ of the set Π which is assumed to be finite in number. D.p.u.s with this property will henceforth be referred to as *polyhedral*.[19]

The SIMEU choice rule is based on the notion of a "degree of implementation" $\lambda(f, \pi; X, \Pi)$ of an extremal prior π by act f in the d.p.u. (X,Π) defined as follows:

$$\lambda(f, \pi;\, X, \Pi) = \frac{E_\pi f - \inf_{g\in\mathcal{A}(X,\Pi)} E_\pi g}{\sup_{g\in\mathcal{A}(X,\Pi)} E_\pi g - \inf_{g\in\mathcal{A}(X,\Pi)} E_\pi g},$$

with $\lambda(f,\pi;X,\Pi) = 1$ if $\sup_{g\in\mathcal{A}(X,\Pi)} E_\pi g - \inf_{g\in\mathcal{A}(X,\Pi)} E_\pi g = 0$. The "degree of implementation" is the expected utility of f with respect to π appropriately renormalized so as set its maximally achievable value equal to 1, and its minimal possible value associated with some admissible act equal to 0.

The SIMEU rule evaluates acts according to a lexicographic maximin of sorts in degrees of implementation:

$$SIMEU(X, \Pi) = \{f \in X | \text{for all } g \in X : \\ : \min_{\pi\in\mathcal{E}(\Pi): E_\pi f \neq E_\pi g} \lambda(f, \pi;\, X, \Pi) \geq \min_{\pi\in\mathcal{E}(\Pi): E_\pi f \neq E_\pi g} \lambda(f, \pi;\, X, \Pi)\}.$$

As further explained in Nehring (2000), SIMEU choices can also be interpreted as the outcome of a cooperative bargaining among fictitious "alter egos" of the decision maker representing the different extremal priors among which judgment has been suspended. This outcome is determined by an adaptation of a lexicographic version of the Kalai and Smorodinsky (1975) solution; the strictly "egalitarian" character of that

[18] In particular, the axiomatization of the SIMEU rule is based on a fairly simplistic and suspiciously conventional context-dependent choice consistency property, the "Weak Axiom of Revealed Equivalent Preference"; see Nehring (2000).

[19] Since the extreme points of a polyhedron are contained in any set that generates this polyhedron as its convex hull, their set is the smallest set that generates $\rhd$ as its unanimity relation. In this sense, the incompleteness of $\rhd$ can be attributed to suspension of judgment among the priors in $\mathcal{E}(\Pi)$.

solution is no coincidence, but a direct consequence of the fundamental Redescription Invariance axiom.

Randomization in Pairwise Choices A noteworthy feature of SIMEU choices is their fifty-fifty randomization in the choice among two non-comparable acts. This may seem counterintuitive at first, since it rules out consideration of the stakes in deciding among non-comparables. In the first, mono-religious version of Pascal's Wager, the difference in stakes appeared to provide a natural criterion to decide among non-comparables in favour of the act "religious life". The SIMEU model rejects this as not sufficiently robust: after all, while non-comparability ensures the existence of some reasonable prior supporting a strict preference in either direction, it does not ensure a minimal weight to such a prior. On the other hand, if the DM really is convinced that the difference in stakes is a decisive consideration, he is able to express this conviction by recording a material likelihood-based preference for the favoured act; in the example just mentioned, it would take a lower probability of the existence of God of at least 5% to make the difference in stakes decisive. By contrast, reflecting the asymmetry in stakes, a clear-cut preference for a secular life would require assigning a lower probability of at least 95% to the non-existence of god.

4.6 A normative model as a repertoire of choice functions

We have so far referred to the partial order $\trianglerighteq$ as a "datum" for the robust choice problem reflecting the DM's imprecise probabilistic beliefs Π. While such beliefs may exist as a genuine *given* for decision making in special cases – as when they are based on some corpus of objective information which is deemed to represent all that the DM is willing to rely on – in the majority of cases such a datum would be hard to pin down. Indeed, even a DM himself will typically have a hard time to demarcate his own range of reasonable beliefs Π directly by introspection or reasoning.

But such a strong interpretation of $\trianglerighteq$ as an "absolute" datum is not really needed for the theory to apply. What *is* needed is an understanding of $\trianglerighteq$ as a hypothetical datum. On this understanding, the proposition expressed by the symbols "$f \in C(X, \trianglerighteq)$" asserts that *if* the DM were to assert $\trianglerighteq$ as his material preference ordering, it would be optimal to choose act f from the feasible set X. From this point of view, a normative model can be viewed as a *repertoire* of choice functions $\{C(., \trianglerighteq)\}_{\trianglerighteq \in \mathbf{PEU}}$, with the DM deciding which choice-function to adopt – which beliefs Π to form – in light of the choice implications associated with different partial orders.

In the context of the SIMEU model, these implications are particularly crisp and transparent. For within the SIMEU model, within which non-comparability translates into randomization, a weak preference of f over

g is equivalent to accepting a choice of f as optimal in the choice between f and g. That is, for any $\trianglerighteq \in \mathbf{PEU}$,

(2) $$f \trianglerighteq g \text{ iff } f \in SIMEU(co(\{f, g\}), \trianglerighteq).$$

Thus, within the SIMEU model, deeming two acts non-comparable gains robustness by ensuring an even-handed, fifty-fifty choice between them, but incurs a significant opportunity cost of failing to choose the potentially "better" act only half of the time. Contemplating such trade-offs between decisiveness and robustness may be essential in guiding the DM where to draw the line between the assertion and suspension of probabilistic belief.

Vagueness A somewhat unnatural and potentially restrictive feature of the adopted framework is the assumed sharp distinction between comparability and non-comparability inherent in a partial ordering. In the representation, this is reflected in an equally sharp distinction between reasonably admissible and inadmissible priors in the set Π. Intuitively, the notion of reasonableness seems graded rather than sharp. It might well be possible to develop generalizations of the SIMEU model that capture such graded distinctions between comparability and non-comparability by allowing for randomization in binary choices with unequal probabilities. Such tilted randomization could reflect a vague "leaning towards preference" in the absence of a full commitment to a preference.

4.7 Robustness and rationality

The conception of robust choice developed above assumed that the PEU ordering $\trianglerighteq$ exhaustively describes the DM's material preferences (betterness judgments). It assumed these preferences to satisfy Independence as a matter of fact, without relying on Independence being itself a normative requirement. Conceivably, Independence might reflect an attitude of ambiguity neutrality, as in the model of extensive ambiguity aversion sketched in section 5.4 below.

To argue that a normative model of robust choice such as the SIMEU model $\{SIMEU(., \trianglerighteq)\}_{\trianglerighteq \in \mathbf{PEU}}$ is categorically rational, one thus needs to argue for the categorical rationality of Independence of material preference[20], while rejecting the categorical rationality of Completeness. We will not attempt to fully defend this view, but briefly mention some pertinent arguments in its support.

[20] We distinguish categorical from prima-facie normative desirability of the Independence axiom. A view of Independence as prima-facie but not categorically rational as exemplified by Levi (1980) may admit ambiguity-averse choice behaviour as a tie-breaking of sorts.

To begin with, Completeness of material preference is arguably not mandated by rationality. The normative claims of Completeness have been challenged before; see Aumann (1962) and Sen (1973) among others. This challenge is fortified here through the conception of robust choice as an optimizing response to incompleteness, as it shows that incompleteness of preference does not render the notion of optimality indeterminate. Furthermore, if the notion of robust choice is fully fleshed out as in the SIMEU model, there is a well-defined notion of *revealed incomplete preference* given by (2).

In the context of robust choice, rather than being a universal normative requirement, Completeness of preference reflects the DM's lack of a desire to suspend judgment. It will naturally obtain in special situations in which there is sufficient "objective" information to pin down reasonable probability estimates uniquely. Conversely, incompleteness is a natural consequence of the judged inadequacy or "ambiguity" of relevant evidence.

By contrast, if one accepts Independence in the presence of objective probabilities, it is much less clear that ambiguity offers a compelling reason to abandon it. Indeed, most explicit critiques of the Independence axiom entail a wholesale rejection of it which condones violations of Independence with objective probabilities such as the Allais paradox as well; see, for example, McClennen (2009). To the best of our knowledge hitherto the only fully worked-out attempt to demarcate violations of Independence due to ambiguity from others has been made in Nehring (2007).[21]

Moreover, it is important to note that once the normative permissibility of incompleteness is accepted, and the natural mirroring of epistemic ambiguity in incompleteness understood, there is no need to make room for ambiguity by relaxing Independence.

5. AMBIGUITY AVERSION AS A RESPONSE TO AMBIGUITY

5.1 The epistemic deficit of the ambiguity-aversion literature

In section 4, we assumed that the DM's full material preference relation satisfied Independence but was incomplete, and we determined the content of rational (robust) choice on the basis of this incomplete preference relation $\trianglerighteq$. We will now allow the DM to express further material preferences (betterness judgments) beyond those based on expected utility comparisons captured by $\trianglerighteq$, preferences that reflect in particular the DM's

[21] The contributions of Schmeidler (1989), Gilboa and Schmeidler (1989), and Ghirardato Maccheroni, and Marinacci (2004) among others can be viewed as implicitly appealing to the notion of Bernoullian rationality proposed there.

ambiguity attitude. This perspective is in line with the ambiguity-aversion literature at the centre of the Al-Najjar and Weinstein critique.[22]

ANW reproach this literature for not providing an adequate notion of "beliefs". As just argued, the conception of rational robust choice sketched above does not fall prey to their critique. On the other hand, we believe that the ANW critique has a valid point in that, on the whole, the ambiguity-aversion literature to date has not adequately supported a view of Ellsberg choices as rational responses to ambiguity. The purpose of the present section is to address this "epistemic deficit" in the literature by piggy-backing on and reinterpreting the choice rule framework developed above.

This epistemic deficit is well-illustrated by the MEU model as axiomatized by Gilboa and Schmeidler (1989), with a representation of the form

$$f \succsim g \text{ iff } \min_{\pi \in \Psi} E_\pi f \geq \min_{\pi \in \Psi} E_\pi g$$

for some closed convex set of priors (decision weights) Ψ.

In this model, the set of decision weights Ψ is often interpreted as the decision maker's beliefs. While this interpretation is a significant part of the rhetorical allure of the model, its conceptual underpinnings are not so clear. On the one hand, as with any representation theorem, there is always the possibility that a DM merely happens to behave in a certain way, and that an explanation in terms of beliefs would be gratuitous. But this skepticism may be diffused by discounting the difference between "true belief" and mere "as-if belief", or even denying its meaningfulness. A stronger argument points to the substantial elements of arbitrariness in taking the set Ψ as the set of as-if beliefs. On the one hand, as pointed out in particular by Marinacci (2002), MEU preferences may well be "probabilistically sophisticated" in the sense of Machina and Schmeidler (1992), in which cases a strong argument can be made for attributing the canonically revealed probability to the DM as his as-if beliefs. In these cases, the set Ψ exaggerates the ambiguity of the DM's beliefs. On the other hand, the set Ψ may also underestimate their ambiguity. In particular, Siniscalchi (2006) has shown that in many cases a DM's MEU choices based on a set of priors Ψ can be replicated by so-called "alpha-MEU" choices behaviour based on a larger set of priors Ψ'; alpha-MEU choosers rank acts according to a weighted combination of the lowest and the highest admissible expected utility, i.e. according to

$$V(f) = \alpha \min_{\pi \in \Psi'} E_\pi f + (1 - \alpha) \max_{\pi \in \Psi'} E_\pi f.$$

[22] In the light of the discussion in section 4.7, such decision makers may be viewed as "semi-rational".

It seems clear from these and similar criticisms that one cannot be expected to obtain a transparent account of the role of beliefs as a determinant of ambiguity-averse behaviour without introducing beliefs as distinct entities as done in the present paper.

5.2 Belief-based ambiguity aversion and exact MEU

Normative Axioms In line with the ambiguity-aversion literature, we will now assume that all completeness gaps in the DM's likelihood-based preference relation $\trianglerighteq$ are now filled by assertions of material preferences that incorporate the DM's aversion to ambiguity. We shall go beyond the ambiguity-aversion literature by not allowing these gaps to be filled in freely, but requiring them to be a response to the ambiguity associated with $\trianglerighteq$. Such ambiguity-averse choice behaviour will be referred to as *belief-based*.

With choice now based on maximization of preference by assumption, Choice Consistency is assumed a fortiori. In the presence of Choice Consistency, the convex-valuedness requirement on choice can now no longer be interpreted as even-handedness, since that interpretation was based on a second-order interpretation of choice judgments. It is now simply a choice-functional counterpart of the usual ambiguity aversion axiom due to Schmeidler (1989) and Gilboa and Schmeidler (1989), and will thus be referred to in this context as Ambiguity Aversion.

The Symmetry and Redescription Invariance axioms now express the belief-basedness of the DM's choice behaviour; they ensure that choices "respond to", "track" the underlying ambiguity Π associated with $\trianglerighteq$.

Admissibility retains its meaning and force.

Choices pinned down uniquely Somewhat remarkably, under complete ignorance, these axioms together pin down choices uniquely. Specifically, they imply the following version of the "lexicographic" maximin rule *LM*:

$$LM(X) = \{f \in X | \text{for all } g \in X \backslash f : \min_{s \in S: f_s \neq g_s} f_s \geq \min_{s \in S: f_s \neq g_s} g_s\}.$$

Let $\mathfrak{D}^{CI}$ denote the family of all complete ignorance decision problems, i.e. of all d.p.u. (X,Π) with $\Pi = \Delta^S$.

Proposition 14 *C on $\mathfrak{D}^{CI}$ satisfies Admissibility, Choice Consistency, Ambiguity Aversion and Strong Symmetry if and only if $C = LM$.*

This Proposition is an immediate consequence of Theorem 2 in Nehring (2000) and Lemma 12 above. There are a number of related results in the literature, the first result of them being Milnor's (1954) characterization of the maximin rule.[23]

Using the results of Nehring (1991, 1992), Proposition 14 can be extended to a characterization of the Leximin Expected Utility rule *LM EU* on the class of polyhedral d.p.u.s as follows:

$$LMEU(X,\Pi)=\{f\in X|\text{for all } g\in X: \min_{\pi\in\mathcal{E}(\Pi):E_\pi f\neq E_\pi g} E_\pi f \geq \min_{\pi\in\mathcal{E}(\Pi):E_\pi f\neq E_\pi g} E_\pi g\}.$$

Modifying Proposition 14, the "*exact MEU*" rule

$$MEU(X,\Pi)=\{f\in X|\text{for all } g\in X: \min_{\pi\in\diamond} E_\pi f \geq \min_{\pi\in\Pi} E_\pi g\}.$$

is obtained if Admissibility is weakened to Weak Admissibility, while an appropriate continuity assumption is added.

Is this too "extreme"? Exactness results such as Propositions 14 may appear surprising – and rather irritating – because they appear to derive an extreme degree (intensity) of ambiguity aversion from the merely qualitative axiom of Ambiguity Aversion.[24] The surprise may be diminished by viewing exact MEU choices as varying in the *extent* rather than *intensity* of ambiguity averse behaviour, that extent being fully commensurate and exactly mirroring the underlying ambiguity in beliefs, as captured by the identity of Ψ and Π.

Yet the irritation is not likely to diminish much as a result of this explanation, since it stems directly from the behaviour under the exact MEU choice rule according to which acts are evaluated at the lowest reasonable prior $\pi\in\Pi$. This is sometimes viewed as "extreme pessimism" or even "paranoia". Neither of these characterizations is really on target, since they take this supporting prior(s) $\arg\min_{\pi\in\Pi} E_\pi f$ as a genuine belief, while the DM's belief is and remains ambiguous, being given by the underlying set of priors Π.

MEU's robustness deficit Any proper critique of exact MEU maximization must target its choice behaviour, not putatively implied beliefs. And, thus reconceived, the irritation with MEU has a point, namely its lack of robustness. As a simple example, consider choice sets $X\subseteq\mathbb{R}^S$ with $S=\{s_1,s_2\}$, and $f_{s_1}\leq f_{s_2}$ for all $f\in X$. Here $\pi=\arg\max_{\pi\in\Pi}\pi_{s_1}$ is a supporting prior of the optimal choice under MEU, and is, for many choice sets X, the unique such prior. Putting all decision weight on the worst prior appears intuitively is very "non-robust"; after all, this prior represents merely one possible one-sided probabilistic assessment of the

[23] For further references, see Nehring (2000).

[24] Gilboa *et al.* (2008) derive exact MEU from an axiom of Caution which states this "extreme" character upfront. Their Caution axiom requires that, for any act f and constant act g, $g\succsim f$ whenever not $f\trianglerighteq g$.

situation. And, of course, our discussion of robust choice in section 4 makes clear that such choice behaviour is indeed far from maximally robust.

The robustness deficit of MEU preferences can be attributed to an extreme feature of them, namely their completeness. For completeness amounts to assuming here that *any* comparability gap due to belief imprecision is filled by an assertion of ambiguity aversion, *leaving no room at all for robustness concerns to enter*. Before briefly sketching a model of incompleteness preferences that allows for intermediate extents of ambiguity-aversion, we will discuss a more conventional and more mainstream response to the extremity suspicions against MEU choices based on the notion of "intensive" degrees of ambiguity aversion.

5.3 Degrees of ambiguity aversion: the intensive approach

The most interesting and explicit attempt to remove the stigma of extreme ambiguity aversion or pessimism from MEU choice behaviour has been made in the recent paper of Gajdos *et al.* (2008), henceforth GHTV.[25] Their model is particularly relevant since it captures crisply widespread intuitions about ambiguity neutrality and degrees of ambiguity aversion.

Adapted to the present framework, GHTV assume that, for every belief Π, the DM's choice behaviour maximizes a MEU preference ordering $\succsim$ with set of decision weights $\Psi = \psi(\Pi)$ whose "fatness" reflects the intensity of the DM's ambiguity aversion. This is made more specific in the fully developed version of their model, which contains as a key ingredient the notion of a "neutral" response to ambiguity captured by SEU maximization based on the Steiner point $\Pi^{st}(\Pi)$. In this model, the set of decision weights Ψ associated with the DM's MEU preference is given as convex combination of the Steiner point and the underlying belief set Π, i.e.

$$\psi_\beta(\Pi) := (1-\beta)\{\pi^{st}(\Pi)\} + \beta\Pi,$$

where $\beta \in [0, 1]$ is the (intensive) degree of the DM's ambiguity aversion. Evidently, a greater degree of ambiguity aversion is associated with a fatter set $\Psi = \psi_\beta(\Pi)$.

The GHTV model boils down to a choice rule C_{ψ_β} given by

$$C_{\psi_\beta}(X, \Pi) = \arg\max_{f\in X}\min_{\pi\in\psi_\beta(\Pi)} E_\pi f.$$

Clearly, C_{ψ_β} violates Redescription Invariance whenever $\beta < 1$, for just the same reason that the limiting case C_{ψ_0} of putatively ambiguity-neutral behaviour (given by $\Psi = \{\pi^{st}(\Pi)\}$) does. Thus C_{ψ_β}-choices fail to be belief-based.

We submit that this shows that the intuitive idea of degrees of ambiguity aversion reflected in the "fatness" of the set Ψ and numerically

[25] Their framework is somewhat different but broadly related.

measured by the parameter β is not well-conceived,[26] since there is really no such thing as a properly ambiguity-neutral *SEU maximizing* response to the ambiguous beliefs Π. That would require a viable Principle of Insufficient Reason to exist, which, as pointed out above, fails to be the case.

5.4 Degrees of ambiguity aversion: an extensive approach

Indeed, from the point of view adopted in this paper, it is natural to identify ambiguity neutrality with robust choice based on the PEU $\trianglerighteq$. One can then conceive of a continuum of ambiguity attitudes, with ambiguity neutrality associated with the partial order $\trianglerighteq$ as the least extensive material preference ordering on one end, and exact MEU preferences as the most extensively ambiguity averse preferences on the other. The following one-dimensional continuum of incomplete preferences $\succsim_\gamma$ with *extensive* ambiguity-aversion parameter $\gamma \in [0, 1]$ captures this range of attitudes naturally: let

$$f \succsim_\gamma g \text{ if and only if } \gamma \left(\min_{\pi \in \Pi} E_\pi f - \min_{\pi \in \Pi} E_\pi g \right) + (1-\gamma) \left(\min_{\pi \in \Pi} (E_\pi f - E_\pi g) \right) \geq 0.$$

Clearly, $\succsim_0$ is $\trianglerighteq$, $\gamma \geq \gamma'$ iff $\succsim_\gamma \supseteq \succsim_{\gamma'}$, and $\succsim_\gamma$ is complete iff $\gamma = 1$ in which case $\succsim_\gamma$ is the exact MEU order associated with $\triangleright$. A fully developed model of such "moderately ambiguity-averse" choice would require an account of robust choice with respect to $\succsim_\gamma$ which is not given here.

To illustrate the role of the extensive ambiguity-aversion parameter γ, consider the preference comparison between an arbitrary act f and a constant act $g = c\mathbf{1}$. It is easily verified that

$$f \succsim_\gamma c\mathbf{1} \quad \text{iff} \quad \min_{\pi \in \Pi} E_\pi f \geq c, \text{ and}$$

$$c\mathbf{1} \succsim_\gamma f \quad \text{iff} \quad c \geq \gamma \min_{\pi \in \Pi} E_\pi f + (1-\gamma) \max_{\pi \in \Pi} E_\pi f.$$

In particular, $\succsim_\gamma$ will exhibit strict preferences as displayed by the ambiguity averse choices in the 2- and 3-colour Ellsberg paradoxes if and only if $\gamma > \frac{1}{2}$. Overall, then, allowing for intermediate degrees of ambiguity aversion γ permits the incorporation of basic patterns of ambiguity aversion while taking the sting out of the extremism charge against the exact MEU preferences.

[26] Belief-basedness plays no role in GHTV's own discussion.

6. CONCLUSION

In section 5, we have provided an account of ambiguity aversion as a "semi-rational" response to ambiguity. Roughly speaking, ambiguity-averse behaviour results from projecting robustness onto the plane of choice consistency/preference maximization. Somewhat less flatteringly, ambiguity aversion can be viewed as *faux robustness* resulting from a category error of sorts. Asserting this does entail that, from the normative point of view, there is a serious mismatch in some of the ambiguity-aversion literature between the epistemic motivations related to model uncertainty (as in the work of Hansen and Sargent) or ignorance about parameters (as in the foundations of statistics) on the one hand and the psychological concerns with the "discomfort" of lacking reliable information that drive Ellsberg choices.

At the same time, asserting the superior rationality of robust choices does not entail a claim of the superiority of robust choice as an *economic* model of decision making under ambiguity. For that, robustness exudes too strong an extraterrestrial flavour, at least in unmitigated form: not only does the computational complexity of ideally robust choices (as, for example, in the SIMEU model) appear to be daunting, robust choice presupposes a degree of detachment and *ataraxia* that seems hard to come by for humans. That being said, if robustness is as convincing normatively as claimed, one would expect it to leave its imprint on earthlings' attempts to cope with ambiguity. Ambiguity aversion is unlikely going to be the whole story.

7. APPENDIX

Proof of Proposition 6

Let $\mathfrak{G}_{(X,\trianglerighteq)} := \mathfrak{G}_X \cap \mathfrak{G}_{\trianglerighteq}$ denote the entire family of transformations Φ under which the problem $(X, \trianglerighteq)$ is symmetric. For any $f \in \mathbb{R}^S$, define the associated "mean" (under the transformations $\mathfrak{G}_{(X,\trianglerighteq)}$) $\overline{f}$ as follows:

$$\overline{f} := \sum_{\Phi \in \mathfrak{G}_{(X,\trianglerighteq)}} \frac{1}{\#\mathfrak{G}_{(X,\trianglerighteq)}} \Phi(f).$$

Elementary arguments show that, for any f, $\overline{f}$ is invariant, and that $f = \overline{f}$ iff f is invariant under $\mathfrak{G}_{(X,\trianglerighteq)}$. Furthermore, a straightforward invariance argument yields the following Lemma.

Lemma 15. *For all $f \in \mathbb{R}^S$ and $\pi \in \Delta^S$, $E_{\overline{\pi}} f = E_{\overline{\pi}} \overline{f} = E_{\pi} \overline{f}$.*

Using a standard separation argument, Weak Admissibility and Randomization imply that the existence of some $\pi \in \Pi$ and $g \in X$ such that i) for all $f \in X : E_\pi g \geq E_\pi f$, and ii), for all $f \in C(X, \trianglerighteq)$, $E_\pi g = E_\pi f$.

By Symmetry, $\Phi(g) \in C(X, \trianglerighteq)$ for all $\Phi \in \mathfrak{G}_{(X,\trianglerighteq)}$, hence by Randomization, $\overline{g} \in C(X, \trianglerighteq)$.

We claim that the pair $(\overline{\pi}, \overline{g})$ has the desired properties.

Take any $f \in C(X, \trianglerighteq)$. By Symmetry, $\Phi(f) \in C(X, \trianglerighteq)$ for all $\Phi \in \mathfrak{G}_{(X,\trianglerighteq)}$, hence by ii), $E_\pi f = E_\pi \Phi(f) = E_\pi \overline{f}$. In particular, $E_\pi g = E_\pi \overline{g}$. Since by the Lemma $E_\pi \overline{f} = E_{\overline{\pi}} f$ and $E_{\overline{\pi}} \overline{g} = E_{\overline{\pi}} g$, and since $E_\pi f = E_\pi g$ using ii) again, we obtain $E_{\overline{\pi}} f = E_{\overline{\pi}} \overline{g}$, establishing the second part of the proposition.

Take now any $f \in X$. From i) and the invariance of X under any $\Phi \in \mathfrak{G}_{(X,\trianglerighteq)}$, $E_\pi g \geq E_\pi \Phi(f)$ for all $\Phi \in \mathfrak{G}_{(X,\trianglerighteq)}$, and thus $E_\pi g \geq E_\pi \overline{f}$. Since we have already shown that $E_\pi g = E_{\overline{\pi}} \overline{g}$ and since $E_\pi \overline{f} = E_{\overline{\pi}} f$ by the Lemma, we obtain $E_{\overline{\pi}} \overline{g} \geq E_{\overline{\pi}} f$, establishing the first part of the proposition.

REFERENCES

Al-Najjar, N. and J. Weinstein. 2009. The ambiguity-aversion literature: a critical assessment. Economics and Philosophy 25: this volume.

Anscombe, F. J. and R. J. Aumann. 1963. A definition of subjective probability. *Annals of Mathematical Statistics* 34: 199–205.

Arrow, K. J. 1960. Decision theory and the choice of a level of significance for the t-test. In *Contributions to probability and statistics: Essays in the honor of Harold Hotelling*, ed. I. Olkin *et al.*, 70–78. Stanford: Stanford University Press.

Arrow, K. J. and L. Hurwicz. 1972. An optimality criterion for decision-making under ignorance. In *Uncertainty and expectations in economics*, ed. D.F. Carter and F. Ford, Oxford: Oxford University Press.

Aumann, R. J. 1962. Utility theory without the completeness axiom. *Econometrica* 30: 445–62.

Bewley, T. F. 1986. *Knightian decision theory, Part I.* Cowles Foundation Discussion Paper No. 807.

Bewley, T. F. 2002. Knightian decision theory, Part I. *Decisions in Economics and Finance* 25: 79–110.

Chernoff, H. 1954. Rational selection of decision functions. *Econometrica* 22: 422–43.

Chew, S.-H. and J. Sagi. 2008. Small worlds: modeling attitudes toward sources of uncertainty. *Journal of Economic Theory* 139: 1–24.

Ellsberg, D. 1961. Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics* 75: 643–69.

Ergin, H. and F. Gul. 2009. A theory of subjective compound lotteries. *Journal of Economic Theory* 144: 899–929.

Gajdos, T., T. Hayashi, J.-M. Tallon, and J.-C. Vergnaud. 2008. Attitude toward imprecise information. *Journal of Economic Theory* 140: 23–56.

Ghirardato, P., F. Maccheroni, M. Marinacci and M. Siniscalchi. 2003. A subjective spin on roulette wheels. *Econometrica* 71: 1897–908.

Ghirardato, P., F., Maccheroni and M. Marinacci. 2004. Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory* 118: 133–73.

Gilboa, I. and D. Schmeidler. 1989. Maxmin expected utility with a non-unique prior. *Journal of Mathematical Economics* 18: 141–53.

Gilboa, I., F. Maccheroni, M. Marinacci, and D. Schmeidler. 2008. *Objective and Subjective Rationality in a Multiple Prior Model*. Mimeo.

Gul, F. and W. Pesendorfer. 2008. The case for mindless economics. In *Foundations of Positive and Normative Economics*, ed. A. Caplin and A. Schotter, New York: Oxford University Press.

Hacking, I. 1975. *The Emergence of Probability*. Cambridge: Cambridge University Press.

Hajek, A. 2008. Pascal's wager. In *Stanford Encyclopedia of Philosophy*, available at http://plato.stanford.edu/entries/pascal-wager/.

Hansen, L. and T. Sargent. 2001a. Acknowledging misspecification in macroeconomic theory. *Review of Economic Dynamics* 4: 519–35.
Hansen, L. and T. Sargent. 2001b. Robust control and model uncertainty. *American Economic Review* 91: 60–6.
Kalai, E., M. Smorodinsky. 1975. Other solutions to Nash's bargaining problem. *Econometrica* 43: 513–18.
Jaynes, E. T. 1957. Information theory and statistical mechanics. *Physical Review* 106: 620–30.
Jaynes, E. T. 2003. *Probability theory: the logic of science*. Cambridge: Cambridge University Press.
Levi, I. 1980. *The enterprise of knowledge*. Cambridge, MA: MIT Press.
Luce, R. D. and H., Raiffa. 1957. *Games and decisions*. New York: John Wiley.
Machina, M. and D. Schmeidler. 1992. A more robust definition of subjective probability. *Econometrica* 60: 745–80.
Marinacci, M. 2002. Probabilistic sophistication and multiple priors. *Econometrica* 70: 755–64.
Milnor, J. 1954. Games against nature. In *Decision processes*, ed. Thrall, Coombs and Davis, 49–59. New York: John Wiley.
McClennen, E. 2009. The normative status of the independence principle. In *Rational and social choice*. ed. P. Anand, P. Pattanaik, and C. Puppe. New York: Oxford University Press.
Moulin, H. 1985. Choice functions over a finite set: a summary. *Social Choice and Welfare* 2: 147–60.
Nehring, K. 1991. *A theory of rational decision with vague beliefs*. Ph.D. dissertation, Harvard University.
Nehring, K. 1992. Foundations for the theory of rational choice with vague priors. In *Decision making under uncertainty - new models and empirical findings*, ed. J. Geweke, 231–242. Dordrecht: Kluwer.
Nehring, K. 2000. Rational choice under ignorance. *Theory and Decision* 48: 205–40.
Nehring, K. 2007. *Bernoulli without Bayes: A Theory of Utility-Sophisticated Preference*. Mimeo, UC Davis.
Nehring, K. 2009. Imprecise probabilistic beliefs as a context for decision-making under ambiguity. *Journal of Economic Theory* 144: 1054–91.
Plott, C. R. 1973. Path independence, rationality and social choice. *Econometrica* 41: 1075–91.
Rigotti, L. and C. Shannon. 2005. Uncertainty and risk in financial markets. *Econometrica* 73: 203–43.
Savage, L. J. 1951. The theory of statistical decision. *Journal of the American Statistical Association* 46: 55–67.
Savage, L. J. 1954. *The foundations of statistics*. New York: Wiley. Second edition 1972, Dover.
Schmeidler, D. 1989. Subjective probability and expected utility without additivity. *Econometrica* 57: 571–87.
Sen, A. 1973. Behavior and the concept of preference. *Economica* 40: 241–59.
Siniscalchi, M. 2006. Electronic supplement to A behavioral characterization of plausible priors. *Journal of Economic Theory* 128: 91–135.
Tversky, A. and P. Wakker. 1995. Risk. attitudes and decision weights. *Econometrica* 63: 1255–80.
Von Kries, J. 1886. *Principien der Wahrscheinlichkeitsrechnung, eine logische Untersuchung*.
Walley, P. 1991. *Statistical reasoning with imprecise probabilities*. London: Chapman and Hall.