EXTENSION RULES
OR
WHAT WOULD THE SAGE DO?

EHUD LEHRER† AND ROEE TEPER‡

August 22, 2011

Abstract. Quite often, decision makers face choices that involve new aspects and alternatives never considered before. Scenarios of this sort may arise, for instance, as a result of technological progress, or from individual circumstances such as information acquisition and improved awareness. In such situations, simple inference rules, past experience and knowledge about historic choice problems, may prove helpful in determining what would be a reasonable action to take vis-a-vis a new problem. In the context of decision making under uncertainty, we introduce and study an extension rule, that enables the decision maker to extend a preference order defined on a restricted domain. We show that utilizing this extension rule results in Knightian preferences, when starting off with an expected-utility maximization model confined to a small world.

Keywords: Extension rule, growing awareness, restricted domain, prudent rule, Knightian preferences, partially-specified probability, lenient rule.

The authors wish to thank Eddie Dekel, Tzachi Gilboa, Edi Karni, Gil Riella, Luca Rigotti and Marie-Louise Vierø for their comments and remarks.

†School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel and INSEAD, Bd. de Constance, 77305 Fontainebleau Cedex, France. e-mail: lehrer@post.tau.ac.il

‡Department of Economics, University of Pittsburgh. e-mail: rteper@pitt.edu.

1. Introduction

Quite often, a decision maker faces choices involving new aspects and alternatives never considered before. Situations of this sort may arise, for instance, as a result of technological progress, or from individual circumstances such as information acquisition and improved awareness. In these novel situations, a decision maker could harness simple inference rules that would help her determine present preferences and actions, based on past experience and historic choices.

1.1. **Inference rules.** Referral to previous cases is prevalent in the judicial system. Courts often find that a current dispute is fundamentally distinct from all previous cases, and judges have the authority and duty to establish new laws based on precedents. When considering judicial disputes, several stages of research and analysis are required to determine judgement. In order to make a ruling a judge must locate any relevant statutes and relevant past decisions and integrate them by simple inference rules and common sense. One heuristic example of an inference is the Latin legal term *maiore ad minus*, meaning 'from greater to smaller'. The *maiore ad minus* principle suggests that existing confidence in a proposition argues for a second proposition implicit in, or implied by, the first. In particular, the second proposition deserves even more confidence than is placed on the first proposition.

Another example comes from propositional logic. Suppose we know that the propositions 'if x is a raven then it is black' and 'the bird currently perched on my window ledge is a raven' are true. We may then infer that the proposition 'the bird currently perched on my window ledge is black' is also true. By using an inference rule, this time a rule known as *modus ponens* (or forward chaining), we expand the set of propositions to which we assign the valuation 'true'.

To illustrate the extension rule central to this paper, suppose you are trying to choose a book to take with you on vacation. You are hesitating between two books: X and Y. A week ago you read a newspaper column by your favorite literary critic L.C. (with whom you often agree) comparing two other books, Z and W. The critic wrote that (a) Z resembles X and is written in the style of the well-known author A, whereas W resembles Y and is written in author B's style. As both you and L.C. agree that (b)

B's style is preferred over that of A, you were surprised to read that L.C. concluded that (c) Z is a better than W.

Through a comparison between another pair of books, Z and W, you can infer what would be the judgment of L.C. regarding X and Y. You make the following inference: L.C. obviously prefers the style of B over that of A. Despite that Z is written in A's style, L.C. concludes that it is better than W. This suggests that the ingredients of Z associated with X are what makes Z better in L.C.'s eyes, and clearly not A's style. In other words, if (c) is true despite (a) and (b), this must stem from the fact that L.C. prefers X over Y. In such a case we would say that Z and W, over which the preferences of L.C. are explicitly stated, testify that X is better than Y.

1.2. **Extension of partial preferences.** In the examples above, the rulings in some of the situations have been previously recognized, established, or canonized. In the third example, for instance, B's style is clearly preferred to that of A and Z is better than W. Based on these preferences, the decision maker infers that between X and Y the former must be better. A natural question arises as to what would be a reasonable inference rule that could enable one to project from past decisions to choices between alternatives that have never been considered before.[1]

In order to answer this question, we take a decision theoretic approach and adopt the Anscombe and Aumann [1] framework of choice between alternatives called *acts*. An act assigns an outcome to every possible state of the world that might be realized. In our setup, as in the examples, the preference relation is specified only over some alternatives, but definitely not over all of them. Hence, the primitive is a *preference relation over a subset of alternatives*. This preference relation is referred to as the *sage*'s preference and can be thought of as the one that was canonized or is already well established.

One particularly interesting instance of our framework is the case where the decision maker's growing awareness enables her to distinguish between states, which the sage

[1]For instance, Beraita of Rabbi Yishmael (apparently from the 8th century) refers to thirteen principles of elucidating the Torah (Bible) and making deductions from it. These principles enable one to extend biblical laws to cases that were never considered in the Bible itself. Curiously, the first principle is *maiore ad minus*.

could not distinguish between. Due to growing awareness, the decision maker encounters a greater state space and acts that she never consider before. In this case, the sage's preferences could be interpreted as the decision maker's own preferences, prior to the world's expansion. After the expansion she wishes to project from past choices to decision problems involving new alternatives that belong to the expanded world.

Another plausible scenario to which our framework is applicable, is a process in which an agent gradually obtains information over time. During this time she accumulates more and more information regarding the true underlying probability distribution that governs the states of the world. In the middle of the process, the agent can rank only the acts about which she has a full information. However, it might happen that before all information is revealed, she is required to decide between acts over which she has only partial information. In such cases she could utilize some extension rules that would help her make inferences.

Consider a decision maker who encounters a decision problem she has never contemplated before; for instance a choice between f and h. Opting to benefit from the guidance of the sage, she recalls that the sage preferred f to g and that h yields a worse outcome than g in every possible state of the world. Employing *maiore ad minus*, she intelligently deduces that the sage would have preferred f over h, in case he – the sage – had to choose between these two alternatives. *Maiore ad minus* is thus used to extend the preference relation from a known domain to *terra incognita*.

The goal of the current paper is to introduce a simple and natural extension rule that generalizes *maiore ad minus*. A decision maker is required to choose between two acts, say f and g, that she has never encountered before. She knows that the sage made his choice between two other acts – say b the *better act* and w the *worse act*. She knows also that a coin toss (a *mixture*) between f and w yields an act that is universally (i.e., in every state of the world) better than the act resulting from tossing the same coin to decide between g and b. In other words, mixing w and b with f and g reverses the preference relation: while b is preferred to w, the mixture of w and f is dominant over the mixture of b and g. The acts b and w provide 'evidence' that f is better than g. The reasoning is that if w is worse than b, but when mixed with f it becomes undoubtedly better than b when mixed with g, then there must be one fact responsible for this reversal: w is mixed with the better act, f. In light of this evidence it is reasonable to adjudicate that f is preferred to g. This rule serves us to extend an existing preference relation from a small set of alternatives to a larger one. This extension rule is dubbed the *prudent rule.*

1.3. **The main results.** Two primary questions are addressed in this paper:

(1) In case the sage's preferences adhere to standard behavioral properties such as expected utility maximization, what are the resulting preferences when employing the prudent rule?
(2) Do the resulting preferences and those of the sage always agree (on the partial domain over which the sage's preferences are well defined)? If they do not always agree, which properties of the sage's preferences guarantee agreement?

To answer the first question, assume that the sage's preferences cohere to maximization of expected utility. Unlike the standard expected-utility maximization case, the prior that represents the sage's preferences need not be unique, because the domain of his preferences is restricted to a subset of acts. That is, there might be several prior distributions over the state space, with respect to which the expected utility maximization yields the same preference order. Nevertheless, in terms of expected value, all these distributions agree on the restricted domain of acts. It turns out that utilizing the prudent rule results in the following preference order: an act f is preferred to g if and only if f is preferred to g with respect to every prior consistent with the sage's expected utility preferences. In other words, a decision maker utilizing the prudent inference rule will prefer f to g if according to *every* belief consistent with the sage's preferences, f is preferred to g. Such preferences are known also as Knightian preferences (Bewley [2]).

The classical Knightian preferences model is established upon a set of priors. An act f is preferred to g if consensus exists among all priors that this is the case. An explanation as to where these priors come from is usually not provided. Here we provide a plausible narrative to the origin of the set of priors employed by Knightian preferences. We commence with conservative preferences that adhere to expected-utility maximization over a restricted domain of acts. These, in turn, induce a set of priors that agree on this domain, which is equivalent to inducing a partially-specified probability (see Lehrer [8]). Extending the Anscombe and Aumann (sage's) preferences by utilizing the prudent rule, results in Knightian preferences that are based on the priors consistent with the sage's preferences.

As to the second question, the answer is negative, namely the resulting preferences and those of the sage do not always agree. For example, in case the sage's preference order is incomplete even over a restricted domain, the extension could possibly be complete. However, given standard assumptions regarding the sage's preferences, the prudent rule extension coincides with them if and only if the the sage's preferences are Knightian. In this case, the prudent rule yields Knightian preferences (over the entire domain of acts) that are based on the set of priors representing the sage's preferences.

The last result also points to a shortcoming of the prudent rule: if consistency with the sage's preferences is an essential requirement, then the prudent rule can be employed only for a relatively narrow family of preferences. However, following Nehring [10] and Ghirardato et al. [4] we show that, given a standard assumption, from *any preference relation* defined over a restricted domain, one can extract a (maximal) Knightian preference component. Given a preference relation, this Knightian component is the unambiguous portion that lends itself to extension: the prudent rule can be reasonably applied to this component.

We conclude with a brief description of an alternative to the prudent rule, the *lenient* rule. Given the sage's preferences, the lenient rule declares that an act f is preferred to g if there is no evidence that g is strictly preferred to f. In a sense, the lenient rule is dual to the prudent one. The prudent rule declares that f is preferred to g providing there is evidence suggesting that this is the case. On the other hand, the lenient rule declares that f is preferred to g by default, namely that f is preferred to g unless there is evidence that reinforces the opposite option. Clearly, the lenient rule is more permissive than the prudent one; when f is preferred to g according to the prudent rule, it is so according to the lenient rule as well.

We discuss the lenient rule in case the sage is an expected-utility maximizer over a restricted domain. Under justifiable preferences (Lehrer and Teper [9]) f is preferred to g if and only if the expected value of f is higher than that of g, with respect to at least one prior in a pre-specified set of priors. It turns out that applying the lenient rule in the case the sage is an expected-utility maximizer over a restricted domain, yields a particular case of justifiable preferences.

1.4. **Related literature.** Gilboa et al. [5] and Karni and Vierø [7] are the most closely related papers to the current one. Unlike the primitive of our model, Gilboa et al. [5] consider a pair of preference relations, Knightian and maxmin. The former is being more reliable but incomplete, while the latter is more precarious, but complete. Gilboa et al. [5] present axioms that tie the two preferences together, in the sense that both are based on an identical set of priors.

Karni and Vierø [7] consider a Bayesian agent and allow the set of alternatives to expand, that is, they let the agent become aware of new alternatives, outcomes and states of the world. Similarly to Gilboa et al. [5], they assume the structure of the agent's preferences also in the post-expansion world. Karni and Vierø postulate that the agent remains Bayesian. Two aspects in which the current paper differs from theirs should be noted. First, in our framework a Bayesian agent need not stay Bayesian after the expansion. Consider, for instance, an Ellsberg type of a decision maker, who obtains information about the probability of certain events, but not of all events. Over a restricted domain, namely over those acts about which she in fully informed, the decision maker maximizes expected utility. Upon letting the world of alternatives expand to the entire domain, however, she may exhibit a behavior pattern other than Bayesian; she may become a Knightian decision maker, for instance. Second, our framework differs from that of Karni and Vierø in that we do not pre-determine the final set of acts between which the decision maker is able to compare. In this paper, the current set of acts depends on the original preference order over the restricted domain and is left for the rule to determine the scope of the extension. The main emphasis of the current paper is directed to the extension rule and its implications.

1.5. **Organization.** The following section presents the framework and primitive of our model. Section 3 then introduces the problem we consider and formally defines the prudent rule. Section 4 studies the example in which the sage is an expected-utility maximizer. In addition, it discusses the implication of the main results to the particular case in which the decision maker gradually accumulates information regarding the state space. The issue of consistency is dealt with in Section 5. Lastly, a discussion regarding the lenient extension rule appears in Section 6. All Proofs are presented in Section 7.

2. The set-up

2.1. **Domain.** Consider a decision making model in an Anscombe–Aumann [1] setting, as reformulated by Fishburn [3]. Let Z be a non–empty finite set of *outcomes*, and let $L = \Delta(Z)$ be the set of all *lotteries*,[2] that is, all probability distributions over Z. In particular, if $l \in L$ then $l(z)$ is the probability that the lottery l assigns to the outcome z.

Let S be a finite non–empty set of *states of nature*. Now, consider the collection $\mathcal{F} = L^S$ of all functions from states of nature to lotteries. Such functions are referred to as *acts*. Endow this set with the product topology, where the topology on L is the relative topology inherited from $[0,1]^Z$. By $\mathcal{F}_c$, we denote the collection of all constant acts. Abusing notation, for an act $f \in \mathcal{F}$ and a state $s \in S$, we denote by $f(s)$ the constant act that assigns a lottery $f(s)$ to every state of nature. Mixtures (convex combinations) of lotteries and acts are performed pointwise. In particular, if $f, g \in \mathcal{F}$ and $\alpha \in [0,1]$, then $\alpha f + (1-\alpha)g$ is the act in $\mathcal{F}$ that attains the lottery $\alpha f(s) + (1-\alpha)g(s)$ for every $s \in S$.

2.2. **Preferences.** In our framework, a decision maker is associated with a binary relation $\succeq$ over $\mathcal{F}$ representing his preferences. We say that $f, g \in \mathcal{F}$ are $\succeq$*-comparable* (to each other) if either $f \succeq g$ or $g \succeq f$. $\succ$ is the asymmetric part of the relation. That is $f \succ g$ if $f \succeq g$ but it is not true that $g \succeq f$. $\sim$ is the symmetric part, that is $f \sim g$ if $f \succeq g$ and $g \succeq f$. Any such binary relation induces a monotonicity relation $\succeq^S$ over $\mathcal{F}$. For two acts $f, g \in \mathcal{F}$, we denote $f \succeq^S g$ if $f(s) \succeq g(s)$ for every $s \in S$.

Our main focus throughout this paper will be preferences over a subset of acts, which we will denote by $\mathcal{H}$. In particular, if f, g are $\succeq$-comparable, then $f, g \in \mathcal{H}$.

Assumption 1. *$\mathcal{H}$ is a convex set of acts that contains the constants $\mathcal{F}_c$.*

The following are well-known properties of preferences, which will be referred to throughout the paper. We say that $\succeq$ is *non-degenerate* if there are $f, g \in \mathcal{H}$ such that $f \nsim g$; $\succeq$ is *reflexive* if $f \sim f$ for every $f \in \mathcal{H}$; $\succeq$ is *complete over constants* if l, l' are $\succeq$-comparable, for every $l, l' \in \mathcal{F}_c$; $\succeq$ is *transitive* if for every $f, g, h \in \mathcal{H}$, $f \succeq g$ and $g \succeq h$ imply $f \succeq h$; $\succeq$ is *monotonic* if $f, g \in \mathcal{H}$ such that $f \succeq^S g$ then $f \succeq g$;

[2] Given a finite set A, $\Delta(A)$ denotes the collection of all probability distributions over A.

$\succeq$ is *continuous* if the sets $\{f \in \mathcal{H} : f \succeq g\}$ and $\{f \in \mathcal{H} : f \preceq g\}$ are closed, for every $g \in \mathcal{H}$; $\succeq$ is *lottery independent* if for every $l, l', r \in \mathcal{F}_c$ and $\alpha \in [0,1]$, $l \succeq l'$ if and only if $\alpha l + (1-\alpha)r \succeq \alpha l' + (1-\alpha)r$; and $\succeq$ satisfies *independence* if for every $f, g \in \mathcal{H}$, $\alpha \in [0,1]$ and $h \in \mathcal{H}$, $f \succeq g$ if and only if $\alpha f + (1-\alpha)h \succeq \alpha g + (1-\alpha)h$. All the properties above are standard, but considered only on a subset of acts $\mathcal{H}$ (as opposed to the entire domain of acts $\mathcal{F}$). In the special case in which $\mathcal{H} = \mathcal{F}$, all these properties are translated to the standard ones over the entire domain of acts.

3. Sage's preferences and the prudent extension rule

Consider a sage's preferences $\succeq$, where the only relevant alternatives at the time of determination were those in $\mathcal{H}$, a sub-collection of acts. In particular, $\succeq$-comparable acts consist only of couples from $\mathcal{H}$. Also consider a decision maker who currently needs to specify his preferences over the entire domain of acts $\mathcal{F}$. Opting to use the guidance of the sage, whose preferences were limited, over $\mathcal{H}$ their preferences coincide. But how can he decide between alternatives that are not in $\mathcal{H}$? The decision maker may try to guess what a reasonable preference between such alternatives would be, given its specification over $\mathcal{H}$. In other words, the decision maker may try to assess what would have been plausible if the sage had to state his preferences over the entire domain.

Leading example. As in Ellsberg's urn, there are three colors and the probability of drawing red is known to be $\frac{1}{3}$. In order to make this example fit our sage's story, suppose that $S = \{R, B, W\}$ and $\mathcal{H}$ consists of the acts of the form (a, b, b). This restricts the sage's domain to acts whose expected utility is known. Assume that the sage prefers (a, b, b) over (a', b', b') iff $\frac{(a-a')}{3} + \frac{2(b-b')}{3} \geq 0$. In other words, (a, b, b) is preferred over (a', b', b') iff the expected value of the former is greater than that of the latter, provided that the probability of the first state is $1/3$. The sage's preferences over $\mathcal{H}$ conform to expected-utility maximization with respect to the prior $p = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. Of course, any other prior that assigns the probability $1/3$ to the first state could represent the sage's preferences as well. The question is how could one extend the sage's preferences so as to compare, for instance, $g = (.5, .4, .2)$ and $h = (.7, .2, .1)$, which are beyond his restricted domain.

Definition 1. *We say that* $\succeq^*$ *is an* extension *of* $\succeq$ *if* $f_1 \succeq f_2$ *implies* $f_1 \succeq^* f_2$*.*

Here we suggest an extension rule to a binary relation $\succeq$. The rule we introduce uses only pointwise dominance, better known as monotonicity, and the relation $\succeq$ itself.

Definition 2. *Fix two acts* $g, h \in \mathcal{F}$*. Two* $\succeq$*-comparable acts,* f_1 *and* f_2*,* testify that h is preferred over g *if there exists* $\alpha \in (0,1)$ *such that* $\alpha h + (1-\alpha) f_1 \succeq^S \alpha g + (1-\alpha) f_2$ *and* $f_2 \succeq f_1$*.*

The $\succeq$-comparable acts f_1 and f_2 testify that g is preferred over h, if despite the fact that f_2 is preferred over f_1, the mixture $\alpha h + (1-\alpha) f_1$ dominates $\alpha g + (1-\alpha) f_2$. That is, f_1 and f_2 testify that this reversed dominance (note, f_1 is on the dominating side), must stem from g being preferred over h.

The following definition is the heart of the paper. Starting with $\succeq$ defined over some pairs of acts, we expand it and define the preference order $\succeq_P$. The rule that enables us to expand the domain of the preference order is called an extension rule. The prominent extension rule is the prudent rule that follows.

PR: The prudent rule. $g \succeq_P h$ if and only if there are $\succeq$-comparable acts testifying that g is preferred over h.

According to the prudent rule, in order to declare that g is preferred to h it is necessary to have an explicit testimony that this is the case. Restating the prudent rule, $g \not\succeq_P h$ if and only if for every $\succeq$-comparable acts and $\alpha \in (0,1)$, $\alpha h + (1-\alpha) f_1 \preceq^S \alpha g + (1-\alpha) f_2$ then $f_2 \succ f_1$. When $\alpha h + (1-\alpha) f_1 \preceq^S \alpha g + (1-\alpha) f_2$ always implies $f_2 \succ f_1$, it is strong evidence that $g \not\succ_P h$.

Proposition 1. $\succeq_P$ *is an extension of* $\succeq$*.*

Proof. If $f \succeq g$ then $\frac{1}{2} f + \frac{1}{2} g = \frac{1}{2} g + \frac{1}{2} f$. In particular, f and g testify that f is preferred to g and, according to the prudent rule, we have that $f \succeq_P g$. □

Leading example, continued. The acts $f_1 = (-0.2, 0.1, 0.1)$ and $f_2 = (0,0,0)$ are in $\mathcal{H}$, and their expected values are 0. That is, the sage is indifferent between f_1 and f_2. Mixing g with f_2 and h with f_1, with equal probabilities yields $\frac{1}{2} g + \frac{1}{2} f_2 = (0.25, 0.2, 0.1)$ and $\frac{1}{2} h + \frac{1}{2} f_1 = (0.25, 0.15, 0.1)$. The former dominates the latter. In particular, f_1, f_2 testify for the preference of g over h.

The prudent rule in this case dictates that the decision maker prefers g to h. We will see below that there is no pair from $\mathcal{H}$ that testifies for the preference of h over g, and thus, the conclusion of the prudent rule in this case is that there is a strict preference of g over h.

4. Expected-utility maximization on a restricted domain

In this section, we consider an example in which the sage is an expected-utility maximizer, and wish to study the implication of our prudent extension rule.

4.1. Expected-utility maximization and partially-specified probabilities (PSP).

Assume that the sage's preferences over $\mathcal{H}$ adhere to expected-utility maximization. In particular, there exists an affine[3] vN–M utility u over lotteries and a prior distribution $p \in \Delta(S)$ over the state space such that, for every $f, g \in \mathcal{H}$, $f \succeq g$ if and only if $E_p(u(f)) \geq E_p(u(g))$. Given the standard Anscombe–Aumann assumptions on preferences over $\mathcal{H}$, such a representation is guaranteed by Proposition 2.2 in Siniscalchi [11]. Since $\mathcal{H}$ is a subset of $\mathcal{F}$, the prior distribution p that represents the preferences over $\mathcal{H}$ need not be unique. Nevertheless, each such prior identically projects on $\mathcal{H}$. To make this statement formal, denote $\mathcal{P}(u, \mathcal{H}, p)$ the collection of priors over S that are consistent with p on the acts in $\mathcal{H}$. That is,

$$\mathcal{P}(u, \mathcal{H}, p) = \{q \in \Delta(S) \ : \ \mathbb{E}_q(u(f)) = \mathbb{E}_p(u(f)) \text{ for all } f \in \mathcal{H}\}.$$

For example, consider the simple case where $\mathcal{H}$ consists only of constants. Then, $\mathcal{P}(u, \mathcal{H}, p) = \Delta(S)$ and, for every $l \in \mathcal{F}_c$ and prior p, we have that $E_p(u(l)) = u(l)$.

Proposition 2. *$E_q(u(f)) = E_{q'}(u(f))$ for every $q, q' \in \mathcal{P}(u, \mathcal{H}, p)$ and every $f \in \mathcal{H}$.*

Proposition 2 states that although several priors may exist over the state space representing the preferences, they all agree on $\mathcal{H}$. This implies that expected utility maximization over $\mathcal{H}$ is identical to partially-specified probabilities (Lehrer [8]), where the information available about the probability is partial and is given by the expected values of some, but typically not all, random variables. In other words, Anscombe–Aumann preferences over a restricted domain induce a partially-specified probability.

[3] $u : L \to \mathbb{R}$ is said to be affine if $u(l) = \sum_{z \in Z} l(z)u(z)$ for every $l \in L$.

Formally, a *partially-specified probability* is given by a list of pairs $\mathcal{X} = \{(X, \mathbb{E}_p(X))\}$ where X are random variables and $\mathbb{E}_p(X)$ is the expectation of X with respect to a probability distribution p. By $\mathcal{X}_1$ we denote the set of random variables X such that $(X, \mathbb{E}_p(X)) \in \mathcal{X}$. This $\mathcal{X}_1$ consists of all random variables the expectations of which become known. That is, the information about p is obtained only through the expectations of the random variables in $\mathcal{X}_1$. For instance, if the variables in $\mathcal{X}_1$ are all constants, then the expectations are necessarily the corresponding constants, meaning that no information is obtained about p. However, if $\mathcal{X}_1$ is the set of all random variables, then p is fully specified. Without loss of generality, we may assume that $\mathcal{X}_1$ is a subspace, for when the expectations of two random variables are known then so is the expectation of any linear combination of them. Note that in case p is not fully specified, there exist other distributions q that are consistent with the information obtained. That is, $\mathbb{E}_q(X) = \mathbb{E}_p(X)$ for every $X \in \mathcal{X}_1$. In particular, the collection of distributions that are consistent with a partially-specified probability is closed and convex.

Getting back to the sage's preferences $\succeq$ over $\mathcal{H}$, due to Proposition 2 we have that $\{(u(f), \mathbb{E}_p(u(f))) \ : \ f \in \mathcal{H}\}$ (where p is any prior representing $\succeq$ over $\mathcal{H}$) is a partially-specified probability (henceforth, *PSP*). Throughout the paper, we denote by $\succeq_{(u,\mathcal{H},p)}$ a binary relation $\succeq$ over $\mathcal{H}$ that coheres to expected-utility maximization with respect to u and p, or alternatively conforms to a PSP.

4.2. **Knightian preferences.** Before investigating the implications of the prudent extension rule to a preference relation that conforms to PSP, we need to make formal the notion of Knightian preferences.

Bewley [2] introduces a status quo approach towards uncertainty and axiomatizes *Knightian preferences.* Under Knightian preferences the decision maker prefers g to h if g dominates h in the sense that according to all priors in a given multiple-prior set, the expected utility induced by g is greater than that induced by h. Formally, for every two acts $g, h \in \mathcal{F}$, g is preferred over h, with respect to the utility function u and the convex and closed set $\mathcal{P}$ of probability distributions over S, if and only if

$$\forall p \in \mathcal{P}, E_p(u(g)) \geq E_p(u(h)).$$

Leading example, continued. Denote the collection of priors, consistent with the sage's information, by $\mathcal{P} = \{q \in \Delta(\{R, B, W\}) \ : \ q(R) = \frac{1}{3}\}$. Let the utility function u be the identity function. Thus, $u(g) - u(h) = g - h$. Note that $E_q(g - h) = E_q(-0.2, 0.2, 0.1) \geq 0$ for every $q \in \mathcal{P}$. The minimum over all such q's is 0, and is obtained when $q \in \mathcal{P}$ assigns probability $\frac{2}{3}$ to the state W. On the other hand, there is a prior in $\mathcal{P}$, for instance $q = (\frac{1}{3}, \frac{2}{3}, 0)$, such that $E_q(h - g) = E_q(0.2, -0.2, -0.1) < 0$. Thus, it is not true that for every $q \in \mathcal{P}$, $E_q(h - g) \geq 0$.

4.3. **The lower integral.** Let $\mathcal{P}$ be a closed set of probability distributions and $\mathcal{A}$ be a set of real-valued functions defined over the state space S. The following functional is defined on the set of random variables over S:

$$\underline{U}_{(\mathcal{A},\mathcal{P})}(X) = \min_{q \in \mathcal{P}} \max_{\substack{Y \in \mathcal{A} \\ X \geq Y}} \mathbb{E}_q(Y). \tag{1}$$

In words,[4] $\underline{U}_{(\mathcal{A},\mathcal{P})}(X)$ is the best possible evaluation of X by random variables in $\mathcal{A}$ that are dominated from below by X itself. In case of PSP, denote by $\mathcal{A}(u, \mathcal{H})$ the linear subspace of random variables spanned by $\{u(f) \ : \ f \in \mathcal{H}\}$. Then, for every $Y \in \mathcal{A}(u, \mathcal{H})$ and every $q \in \mathcal{P}(u, \mathcal{H}, p)$, $\mathbb{E}_q(Y) = \mathbb{E}_p(Y)$. Thus,

$$\underline{U}_{(\mathcal{A}(u,\mathcal{H}),\mathcal{P}(u,\mathcal{H},p))}(X) = \min_{q \in \mathcal{P}(u,\mathcal{H},p)} \max_{\substack{Y \in \mathcal{A}(u,\mathcal{H}) \\ X \geq Y}} \mathbb{E}_q(Y) = \max_{\substack{Y \in \mathcal{A}(u,\mathcal{H}) \\ X \geq Y}} \mathbb{E}_p(Y).$$

Leading example, continued. To illustrate the lower integral, let $\mathcal{A} = \text{span}\{(1, 1, 1), (1, 0, 0)\}$. That is, $\mathcal{A}$ is the collection of all linear combinations of alternatives that have a know expected value. The largest element in $\mathcal{A}$ dominated by $g - h = (-0.2, 0.2, 0.1)$ is $(-0.2, 0.1, 0.1)$, the expected value of which is 0. Thus, $\underline{U}(-0.2, 0.2, 0.1) = \mathbb{E}_p(-0.2, 0.1, 0.1) = 0$. On the other hand, the largest element in $\mathcal{A}$ dominated by $h - g = (0.2, -0.2, -0.1)$ is $(0.2, -0.2, -0.1)$, the expected value of which is strictly smaller than 0. Thus, $\underline{U}(0.2, -0.2, -0.1) = \mathbb{E}_p(0.2, -0.2, -0.1) < 0$.

4.4. **What would the sage do?** In the leading example we have seen that (a) the prudent rule yields a preference for g over h, (b) the lower integral of $g - h$ is non-negative, and (c) the Knightian preferences, with respect to the set of priors consistent with the sage's information, prefers g to h as well. The main result of this section shows that the fact that (a) , (b) and (c) occur simultaneously is not a coincidence.

[4]By $X \geq Y$ we mean that $X(s) \geq Y(s)$ for every $s \in S$.

To formulate this result, fix a preference relation $\succeq_{(u,\mathcal{H},p)}$ that conforms to a PSP as described in Section 4.1.

Theorem 1. *Consider a binary relation $\succeq_{(u,\mathcal{H},p)}$. The following are equivalent:*

*(1) **PR** is applied. That is, $\succeq_P$ extends $\succeq_{(u,\mathcal{H},p)}$;*

(2) For every $g, h \in \mathcal{F}$,

$$g \succeq_P h \text{ if and only if } \underline{U}_{(\mathcal{A}(u,\mathcal{H}),\mathcal{P}(u,\mathcal{H},p))}(u(g) - u(h)) \geq 0;$$

(3) For every $g, h \in \mathcal{F}$,

$$g \succeq_P h \text{ if and only if } \mathbb{E}_q(u(g)) \geq \mathbb{E}_q(u(h)) \; \forall q \in \mathcal{P}(u,\mathcal{H},p).$$

Corollary 1. *Applying the prudent extension rule to $\succeq_{(u,\mathcal{H},p)}$ generates Knightian preferences, where the set of priors representing the preferences consists of all priors consistent with the PSP (i.e., $\mathcal{P}(u,\mathcal{H},p)$).*

In particular, when applying the prudent extension rule, the decision maker would prefer g over h iff all priors consistent with the sage's preferences agree that g should be ranked higher than h.

In (3) the Knightian preference order is expressed in the traditional way: g is preferred to h iff there is a consensus among all priors in $\mathcal{P}$ that this is the case. In statement (2), however, the Knightian preference order is expressed differently. It is expressed in terms of the lower integral: g is preferred to h iff the lower integral of $u(g) - u(h)$ is at least zero. Recall that the lower integral is defined by the set of priors $\mathcal{P}(u,\mathcal{H},p)$ and the set of variables $\mathcal{A}(u,\mathcal{H})$. However, in the case of a binary relation $\succeq_{(u,\mathcal{H},p)}$, all priors in $\mathcal{P}(u,\mathcal{H},p)$ agree on the variables in $\mathcal{A}(u,\mathcal{H})$. Thus, there is no need to consider all priors; one prior, for instance p, is enough; meaning that g is preferred to h iff

$$\max_{\substack{Y \in \mathcal{A}(u,\mathcal{H}) \\ u(g)-u(h) \geq Y}} \mathbb{E}_p(Y) \geq 0. \tag{2}$$

Statement (2) of Theorem 1 and Eq. (2) imply that the Knightian preference order can be expressed in terms of the variables in $\mathcal{A}(u,\mathcal{H})$, while the set of priors is playing no role. In other words, when starting up with Anscombe–Aumann preferences over a restricted domain (or alternatively, with PSP specified only on the acts in $\mathcal{H}$), there is no need to find all the priors that agree with p on $\mathcal{H}$. Rather, only the expected utility associated with every act in $\mathcal{H}$, like in the preferences themselves, are needed in order to find out whether or not g is preferred over h (i.e., to determine whether or not Eq. (2) is satisfied).

An interesting question is what would happen if the extension **PR** would be applied twice, first to $\succeq$ and then to $\succeq_P$. Would it extend $\succeq_P$ further or leave it unchanged. It turns out that in case the original preference relation is induced by PSP, one step captures the entire extension: applying **PR** to $\succeq_P$ leaves it unchanged.

Proposition 3. *Assume that $\succeq$ are Knightian preferences over $\mathcal{F}$ represented by a utility u and a collection of priors $\mathcal{P}(u, \mathcal{H}, p)$. Then, the following are equivalent:*

*(1) **PR** is applied. That is, $\succeq_P$ extends $\succeq$;*

(2) For every $g, h \in \mathcal{F}$,

$$g \succeq_P h \text{ if and only if } \mathbb{E}_q(u(g)) \geq \mathbb{E}_q(u(h)) \ \forall q \in \mathcal{P}(u, \mathcal{H}, p).$$

We will later see that this result is a particular case and that extension is being made to the full extent during the first application in a more general family of preferences.

4.5. **Growing awareness.** So far, we presented a model having a fixed state space. However, this setup can accommodate informational acquisition and growing awareness. It allows for an expanding state space that might result from an improved ability to distinguish between different states.

Suppose that the decision maker's information about the state space is represented by a partition. That is, the decision maker can distinguish between states in different parts of the partition, but she cannot do so between states within the same part. A particular case of our model is where $\mathcal{H}$ consists of all acts measurable with respect to a partition of the state space. From the decision maker's perspective, $\mathcal{H}$ consists of all acts possible.

When the partition representing the decision maker's information becomes finer, it means that the decision maker can distinguish between more states than she could before. This might be a result of a growing awareness, for instance. Letting the partition become finer is equivalent to letting the state space become larger.

In order to illustrate this point, assume that a decision maker's preferences adhere to expected-utility maximization as she faces a state space S consisting of two states,

R and BW.[5] The probability of R is $\frac{1}{3}$, while that of BW is $\frac{2}{3}$. As a result of growing awareness the decision maker realizes that BW is in fact two separate states, B and W. She is now in the position of the decision maker in the leading example. She is facing the state space $S = \{R, B, W\}$ with the probability of R being $\frac{1}{3}$.

Splitting one state into several others, makes a fully-specified probability, like in an expected utility model, partially-specified. From a fully-specified probability defined over a state space, it becomes a partially-specified probability over an expanded state space.

At this point, the decision maker is able to compare between any two acts of the type (a, b, b), but not between other acts. A preference order that conforms to PSP is precisely the situation dealt with in Theorem 1. Applying the prudent rule, the decision maker can extend her preferences, currently defined on a restricted domain of acts, and obtain Knightian preferences over the set of all acts defined on the expanded state space.

5. Consistency

Proposition 1 states that no matter what the preference $\succeq$ over $\mathcal{H}$ is, the prudent rule generates preferences $\succeq_P$ that extend $\succeq$. An important question is whether $\succeq_P$ is consistent with $\succeq$. Throughout this section we will make the following assumption.

Assumption 2. *$\succeq$ over $\mathcal{H}$ is non-degenerate, reflexive, transitive, complete over constants, monotonic, continuous and lottery independent.*

Definition 3. *$\succeq_P$ is* consistent *with $\succeq$ if, for every $f, g \in \mathcal{H}$, $f \succeq_P g$ if and only if $f \succeq g$.*

Before attempting to answer the question about consistency, consider the following observation. If $\succeq$ satisfies lottery independence, then for any $f_1, f_2 \in \mathcal{H}$ and $e \in \mathcal{F}$,

$$\alpha g + (1-\alpha) f_1 \succeq^S \alpha h + (1-\alpha) f_2$$

if and only if

$$\alpha(\gamma g + (1-\gamma)e) + (1-\alpha) f_1 \succeq^S \alpha(\gamma h + (1-\gamma)e) + (1-\alpha) f_2.$$

[5]No, it is not a typo. The color is Blite.

That is, f_1 and f_2 in $\mathcal{H}$ testify (recall Definition 2) that g is preferred to h if and only if they testify that $\gamma g + (1-\gamma)e$ is preferred to $\gamma h + (1-\gamma)e$. In other words,

Corollary 2. *If $\succeq$ satisfies lottery independence, the extension $\succeq_P$ satisfies independence.*

This statement implies that when the original order does not satisfy independence, the prudent rule does not yield a consistent extension. The following proposition summarizes this conclusion (for the sake of completeness we provide the proof as well).

Proposition 4. *If $\succeq$ satisfies lottery independence but not independence, then $\succeq_P$ is not consistent with $\succeq$.*

Proof. If $\succeq$ does not satisfy independence, then there are $f, g \in \mathcal{H}$, $h \in \mathcal{F}$ and $\alpha \in [0,1]$ such that $\alpha f + (1-\alpha)h$, $\alpha g + (1-\alpha)h \in \mathcal{H}$, were $f \succeq g$ but $\alpha f + (1-\alpha)h \not\succeq \alpha g + (1-\alpha)h$. $f \succeq_P g$ since $\succeq_P$ extends $\succeq$, and since $\succeq_P$ satisfies independence we have that $\alpha f + (1-\alpha)h \succeq \alpha g + (1-\alpha)h$. □

Thus, independence is a necessary condition for consistency. The question is whether independence is a sufficient condition. The following theorem answers this question on the affirmative.

Theorem 2. *Given Assumption 2, $\succeq_P$ is consistent with $\succeq$ if and only if $\succeq$ satisfies independence.*

Theorem 2 is implied by Proposition 4 above and the following two propositions.

Proposition 5. *Let $\succeq$ be a preference over $\mathcal{H}$ satisfying Assumption 2. Then the following are equivalent:*

1. $\succeq$ satisfies independence;

2. There exist a non-constant affine utility function u over L and a closed convex set of priors $\mathcal{P} \subseteq \Delta(S)$, such that, for every $f, g \in \mathcal{H}$,

$$f \succeq g \;\Leftrightarrow\; \forall p \in \mathcal{P}, E_p(u(f)) \geq E_p(u(g)).$$

Furthermore, (i) u is unique up to a positive linear transformations; and (ii) if $\mathcal{P}'$ represents $\succeq$ as well, then $\mathcal{P}' \subseteq \mathcal{P}$.

The preferences obtained by independence are standard Knightian preferences, but restricted to $\mathcal{H}$. Unlike the case were the preferences are defined over the entire set of acts, here the representing set of priors need not be unique. However, this case owns a maximal set (with respect to inclusion) that contains all the sets of priors that represent the preferences. To conceptualize that, consider the example in which $\mathcal{H} = \mathcal{F}_c$. Every subset of priors induces the same preferences which, in fact, adhere to expected-utility maximization.

Proposition 6. *Consider Knightian preferences over $\mathcal{H}$ represented by a utility u and maximal $\mathcal{P}$. The following are equivalent:*[6]

(1) ***PR*** *is applied. That is, $\succeq_P$ extends $\succeq$;*

(2) For every $g, h \in \mathcal{F}$,

$$g \succeq_P h \text{ if and only if } \underline{U}_{(\mathcal{A}(u,\mathcal{H}),\mathcal{P})}(u(g) - u(h)) \geq 0;$$

(3)) For every $g, h \in \mathcal{F}$,

$$g \succeq_P h \text{ if and only if } \mathbb{E}_q(u(g)) \geq \mathbb{E}_q(u(h)) \ \forall q \in \mathcal{P}.$$

Theorem 2 suggests that the use of the prudent extension rule is quite narrow, simply because it generates a consistent extension only if the original order satisfies independence. We argue that this is not the case though. Given a preference relation over $\mathcal{H}$ (that need not satisfy independence), we follow Nehring [10] and Ghirardato et al. [4] and show that one can extract "the Knightian component" out of it. Given a preference $\succeq$ over $\mathcal{H}$, we derive an unambiguous preference $\succeq'$ defined as follows. $f \succeq' g$ iff

$$\alpha f + (1-\alpha)h \succeq \alpha g + (1-\alpha)h$$

for every $\alpha \in [0, 1]$ and $h \in \mathcal{H}$.

Proposition 7. *Assume that $\succeq$ satisfies constant independence and Assumption 2. Then $\succeq'$ over $\mathcal{H}$ is a Knightian preference.*

[6]The proposition need not assume that $\mathcal{P}$ is the maximal set of priors representing the Knightian preferences over $\mathcal{H}$. If $\mathcal{P}' \subseteq \mathcal{P}$ represents these prferences as well, then terms (2) and (3) in the proposition would still be formulated with the maximal set $\mathcal{P}$.

This proposition holds true due to Proposition 4 in Ghirardato et al. [4] and Proposition 5 above.[7] It implies that if one wishes to employ the prudent rule in a consistent manner, one should apply it to the preferences' Knightian component.

Two points should be noted. First, the Knightian component is the largest preference order contained in $\succeq$ that satisfies independence. That is, every preference relation contained in $\succeq$ and satisfies independence is also contained in the Knightian component. Second, the constant-independence assumption (Gilboa and Schmeidler [6]) is a standard and natural assumption in the literature of decision making under uncertainty. Nevertheless, here it is a necessary assumption for the Knightian component to be non-empty; that is, able to compare, at least, between two constant acts.

Finally, a corollary of the results in this section is a generalization of Proposition 3. Whenever the sage's preferences are Knightian, the prudent rule will have full affect after the first iteration. A second iteration would leave the preferences as is.

6. The lenient extension rule

Definition 2 indicates when two acts testify that h is preferred over g. Employing this testimony in order to extend the preference order to the pair h and g, results in a weak preference order: according to the extended preference h is weakly preferred over h.

Suppose, however, that there are two $\succeq$-comparable acts f_1 and f_2 such that (i) f_2 is strictly preferred over f_1; and (ii) $\alpha h + (1-\alpha) f_1 \succeq^S \alpha g + (1-\alpha) f_2$ (note, f_1 is on the dominating side). It means that a mixture of f_1 with h dominates a mixture of f_2 with g. In other words, mixing with g and h reverses the strict preference over f_2 over f_1. Such a reversal must stem from g being strongly preferred over h. That is, when h and g reverse a strong preference among f_1 and f_2, this is a testimony that h is strictly preferred over g

We make it formal in the following definition.

[7]Ghirardato et al. [4] prove this result when $\mathcal{H} = \mathcal{F}$; however, the same proof holds in general as well, even if the preferences over $\mathcal{H}$ are incomplete.

Definition 4. *Fix two acts $g, h \in \mathcal{F}$. Two $\succeq$-comparable acts, f_1 and f_2,* testify that h is strongly preferred over g *if $f_2 \succ f_1$ and there exists $\alpha \in (0, 1)$ such that $\alpha h + (1 - \alpha) f_1 \succeq^S \alpha g + (1 - \alpha) f_2$.*

Definitions 2 and 4 introduce two types of testimonies. While Definition 2 defines what constitutes a testimony about g being (weakly) preferred over h, Definition 4 defines what constitutes a testimony about g being strongly preferred over h. The prudent rule utilizes a testimony of the first type. In order to declare that g is preferred over h, it requires the existence of a testimony that this is the case. The upcoming extension rule, dubbed the *lenient rule*, utilizes a testimony of the second type. The lenient rule declares that g is preferred over h unless otherwise is suggested, that is, unless there is an explicit testimony that h is strictly preferred over g. In other words, the lenient rule declares that g is preferred over h by default, unless there is an explicit testimony suggesting the opposite.

LR: The lenient rule. $g \succeq_L h$ if and only if there are no $\succeq$-comparable acts testifying that h is strongly preferred over g.

According to the lenient rule, $g \succeq_L h$ by default. That is, this rule declares that $g \succeq_L h$, unless there is a testimony attesting that its negation (i.e., $h \succ_L g$) is true. The comparison between the lenient and the prudent rules is made easier if the latter is phrased as follows: $g \not\succeq_P h$ if and only if for every two acts $f_1, f_2 \in F$ and for every $\alpha \in (0, 1)$, $\alpha h + (1 - \alpha) f_1 \preceq^S \alpha g + (1 - \alpha) f_2$ implies $f_1 \succ f_2$.

The lenient rule, **LR**, captures together monotonicity and continuity reasonings. Suppose that $\alpha h + (1 - \alpha) f_1 \succeq^S \alpha g + (1 - \alpha) f_2$. Monotonicity suggests that a reasonable decision maker would prefer $\alpha h + (1 - \alpha) f_1$ over $\alpha g + (1 - \alpha) f_2$. Now suppose that $h \succ_L g$. Continuity suggests that for a sufficiently small α, $\alpha h + (1 - \alpha) f_1 \succ \alpha g + (1 - \alpha) f_2$. But if there is no such α, it attests that $h \not\succ_L g$. Stated differently, when $\alpha h + (1 - \alpha) f_1 \succeq^S \alpha g + (1 - \alpha) f_2$ it will always imply that $f_1 \succeq f_2$, and thus serves as evidence that $h \not\succ_L g$, meaning that $g \succeq_L h$.

In order to formulate a theorem regarding the preferences obtained as a result of applying the lenient rule, we present notions dual to those discussed in Section 3. In a recent work, Lehrer and Teper [9] propose a mild approach to uncertainty and axiomatize *justifiable preferences.* According to justifiable preference-relation, the decision maker prefers f over g if there exists a prior in a multiple-prior set, with respect to which the expected utility of f is at least as high as that of g. More formally, for every $f, g \in \mathcal{F}$

$$f \succeq^* g \ \Leftrightarrow \ \exists p \in \mathcal{P}; E_p(u(f)) \geq E_p(u(g)),$$

where $\mathcal{P}$ is a convex closed set of probability distributions over S.

The functional $\underline{U}$ is defined as minmax (see Eq. (1)). Its dual, the *upper integral*, is the following maxmin:

$$\overline{U}_{(\mathcal{A},\mathcal{P})}(X) = \max_{q \in \mathcal{P}} \min_{\substack{Y \in \mathcal{A} \\ Y \geq X}} \mathbb{E}_q(Y).$$

$\overline{U}_{(\mathcal{A},\mathcal{P})}(X)$ is the evaluation of X by random variables in $\mathcal{A}$ that dominate it from above. This evaluation can be interpreted as an optimistic one. The evaluation of X is obtained through approximations of X by variables that dominate it, using the prior that assigns the highest value.

We are now ready to formulate a result regarding the implication of the lenient extension rule when applied to preferences that adhere to PSP.

Theorem 3. *Consider the PSP relation $\succeq_{(u,\mathcal{H},p)}$. Let $\mathcal{P} = \mathcal{P}(u, \mathcal{H}, p)$ and $\mathcal{A} = \mathcal{A}(u, \mathcal{H})$. The following are equivalent:*

(1) ***LR*** *is applied. That is, $\succeq_L$ extends $\succeq_{(u,\mathcal{H},p)}$;*

(2) For every $g, h \in \mathcal{F}$,

$$g \succeq_L h \textit{ if and only if } \overline{U}_{(\mathcal{A},\mathcal{P})}(u(g) - u(h)) \geq 0;$$

(3) For every $g, h \in \mathcal{F}$,

$$g \succeq_L h \textit{ if and only if } \ \exists q \in \mathcal{P} \text{ such that } \mathbb{E}_q(u(g)) \geq \mathbb{E}_q(u(h)).$$

Corollary 3. *Applying the lenient extension rule generates justifiable preferences.*

7. Proofs

7.1. **Proof of Proposition 2.** Consider two prior distributions p, q over S that represent $\succeq$ over $\mathcal{H}$. For every $f, g \in \mathcal{H}$, we have that

$$E_p(u(f)) \geq E_p(u(g)) \;\Leftrightarrow\; E_q(u(f)) \geq E_q(u(g)). \tag{3}$$

Now, to the contrary of the postulate, assume that E_p and E_q do not coincide over $\mathcal{H}$. That is, there exists $h \in \mathcal{H}$ such that, without loss of generality, $E_p(u(h)) > E_q(u(h))$. In particular, there exists $c \in \{u(l) \;:\; l \in L\}$, such that $E_p(u(h)) > c > E_q(u(h))$. Pick any $l \in L$ with $u(l) = c$. Then, since $\mathcal{H}$ contains the constants, we have that

$$E_p(u(h)) > c = E_p(u(l)) \;\text{ and }\; E_q(u(l)) = c > E_q(u(h)),$$

in contradiction to Eq. (3). □

7.2. **Proof of Theorem 1.** Before turning to the proof of Theorem 1 we need the following lemma.

Lemma 1 (The key lemma)**.** *Consider a Knightian preference order $\succeq$ over $\mathcal{H}$ represented by a utility u and a maximal set of priors $\mathcal{P} = \mathcal{P}(\succeq, u, \mathcal{H})$. Denote $\mathcal{A} = \mathcal{A}(u, \mathcal{H})$. Then,*
(i) $\underline{U} \leq \overline{U}$;
(ii) $\underline{U}$ *and* $\overline{U}$ *are concave and convex functions, respectively;*
(iii) $\underline{U}(X) = \min_{q\in\mathcal{P}} \mathbb{E}_q(X)$ *and* $\overline{U}(X) = \max_{q\in\mathcal{P}} \mathbb{E}_q(X)$;
(iv) *If* $\underline{U}(X) \geq \overline{U}(W)$, *then for every* $q \in \mathcal{P}$, $\mathbb{E}_q(X) \geq \mathbb{E}_q(W)$;
(v) *If* $\overline{U}(X) \geq \underline{U}(W)$, *then there exist* $q, q' \in \mathcal{P}$ *such that* $\mathbb{E}_q(X) \geq \mathbb{E}_{q'}(W)$.

Proof. Assertions (i) and (ii) are easy to prove. We now prove (iii).

$\mathcal{P}$ is closed and convex. For every $p \in \mathcal{P}$, $\mathbb{E}_p(X)$ is linear. Thus, $\underline{V}(X) := \min_{p\in P} \mathbb{E}_p(X)$ and $\overline{V}(X) := \max_{p\in\mathcal{P}} \mathbb{E}_p(X)$ are well defined and are convex and concave, respectively. Moreover, $\overline{V} \geq \underline{V}$.

For every $q \in \mathcal{P}$, $\mathbb{E}_q \geq \underline{U}$. To see this claim, fix a random variable X and suppose that $\underline{U}(X) = \mathbb{E}_p(Y)$, where $X \geq Y$, $Y \in \mathcal{A}$, and $\mathbb{E}_p(Y) \leq \mathbb{E}_q(Y)$ for every $q \in \mathcal{P}$. Thus, $\mathbb{E}_q(X) \geq \mathbb{E}_q(Y) \geq \mathbb{E}_p(Y) = \underline{U}(Y)$ for every $q \in \mathcal{P}$. Hence, $\mathbb{E}_q \geq \underline{U}$, implying that $\underline{V} \geq \underline{U}$, while $\underline{V}(Y) = \underline{U}(Y)$ for every $Y \in \mathcal{A}$.

We show now that $\underline{V} = \underline{U}$. Fix a random variable X. There is a probability distribution r such that $\mathbb{E}_r$ supports $\underline{U}$ at $\underline{U}(X)$. That is, $\mathbb{E}_r(X) = \underline{U}(X)$ and $\mathbb{E}_r \geq \underline{U}$. In particular, $\mathbb{E}_r(Y) \geq \underline{U}(Y)$, for every $Y \in \mathcal{A}$. The reason why such $\mathbb{E}_r$ exists is that $\underline{U}$ is concave. Thus, it is the minimum of linear functions. But since $\underline{U}$ is homogenous, these functions are of the form $r \cdot X$ (the dot product of r and X), where r is a vector in $\mathbb{R}^S$. Since $\underline{U}$ is monotonic w.r.t. to $\geq$, r is non-negative. Furthermore, since the constants (positive and negative) are in $\mathcal{A}$, and $\underline{U}(\mathbf{c}) = c$, and therefore r is a probability distribution.

We claim that $r \in \mathcal{P}$. Otherwise there are $f, g \in \mathcal{H}$ such that $f \succeq g$, but $\mathbb{E}_r(u(g)) > \mathbb{E}_r(u(f))$. Thus, $0 > \mathbb{E}_r(u(f)) - \mathbb{E}_r(u(g)) = \mathbb{E}_r(u(f) - u(g))$. But since $\mathbb{E}_r \geq \underline{U}$, $\mathbb{E}_r(u(f) - u(g)) \geq \underline{U}(u(f) - u(g))$. We obtain,

$$\underline{U}(u(f) - u(g)) < 0. \tag{4}$$

However, $u(f) - u(g) \in \mathcal{A}$ and $f \succeq g$. Thus, for every $p \in \mathcal{P}$, $\mathbb{E}_p(u(f) - u(g)) \geq 0$, implying that $\underline{U}(u(f) - u(g)) = \min_{p \in \mathcal{P}} \mathbb{E}_p(u(f) - u(g)) \geq 0$, in contradiction to Eq. (4). We conclude that $r \in \mathcal{P}$. Thus, $\underline{U}(X) = \mathbb{E}_r(X) \geq \min_{p \in \mathcal{P}} \mathbb{E}_p(X) = \underline{V}(X)$, implying $\underline{U}(X) = \underline{V}(X)$. Hence, $\underline{V} = \underline{U}$, as desired.

The second assertion of (iii) is proven in a similar way. (iv) and (v) are immediate consequences of (iii). $\square$

Proof of Theorem 1. (1) $\Leftrightarrow$ (2). Define a binary relation $\succeq^*$ over all acts as follows: for every two acts g and h, $g \succeq^* h$ if and only if $\underline{U}(u(g) - u(h)) \geq 0$. Let us show that $\succeq^*$ and $\succeq_P$ coincide. Suppose that $g \succeq^* h$ and assume on the contrary, that $g \not\succeq_P h$. **PR** implies that $f_2 \succ f_1$ for every two $\succeq$-comparable acts f_1, f_2 and $\alpha \in (0,1)$ such that $\alpha h + (1-\alpha) f_1 \preceq^S \alpha g + (1-\alpha) f_2$. In particular $\mathbb{E}_p(u(f_2)) > \mathbb{E}_p(u(f_1))$. However, $\alpha h + (1-\alpha) f_1 \preceq^S \alpha g + (1-\alpha) f_2$ implies that $\frac{1-\alpha}{\alpha}[u(f_1) - u(f_2)] \leq u(g) - u(h)$. Thus, $\underline{U}(u(g) - u(h)) < 0$, which contradicts the assumption that $g \succeq^* h$.

On the other hand, assume that $g \succeq_P h$. By **PR**, $\underline{U}(u(g) - u(h))$ is at least 0, which implies that $g \succeq^* h$.

(2) $\Rightarrow$ (3) is due to (iv) of Lemma 1 and (2) $\Leftarrow$ (3) is a consequence of (iii) of Lemma 1 above. $\square$

7.3. **Proof of Proposition 5.** The necessity of independence is obvious. We prove that it is sufficient. The hypothesis of the von Neumann–Morgenstern theorem hold. Thus, this theorem assures the existence and uniqueness, up to a positive linear transformation, of a non-constant and affine function $u : L \to [0,1]$, which represents the preferences restricted to $\mathcal{F}_c$. The existence of such a utility function u induces a preference relation $\hat{\succeq}$ over $\{u(h) \,:\, h \in \mathcal{H}\}$, defined as follows. For $f, g \in \mathcal{H}$, $u(f) \hat{\succeq} u(g)$ if and only $f \succeq g$. To see that $\hat{\succeq}$ is well defined, assume that there are $f, f', g, g' \in \mathcal{H}$ such that $f \succeq g, u(f) = u(f')$ and $u(g) = u(g')$. The monotonicity assumption yields that $f \sim f'$ and $g \sim g'$, and in turn, transitivity implies $f' \succeq g'$.

Now,[8] let $\mathcal{D} = \text{conv}\{\text{u(f)} - \text{u(g)} \,:\, \text{f}, \text{g} \in \mathcal{H},\ \text{u(f)} \hat{\succeq} \text{u(g)}\}$. We show that $\mathcal{D}$ does not intersect the negative orthant of $\mathbb{R}^{|S|}$. Assume it does. This means that there are non-negative $\{\lambda_i\}_{i=1}^n$ that some up to 1 and $\{\langle f_i, g_i\rangle\}_{i=1}^n$, each in $\mathcal{H}\times\mathcal{H}$, where $f_i \succeq g_i$ for all $i \in \{1, ..., n\}$, and $\sum_{i=1}^n \lambda_i(u(f_i) - u(g_i)) < 0$. This, however, is impossible due to the independence and monotonicity of $\succeq$. Indeed, both $\sum_{i=1}^n \lambda_i f_i$ and $\sum_{i=1}^n \lambda_i g_i$ are in $\mathcal{H}$, due to convexity of $\mathcal{H}$. In addition, independence implies that $\sum_{i=1}^n \lambda_i f_i \succeq \sum_{i=1}^n \lambda_i g_i$. Lastly, monotonicity guarantees that $\sum_{i=1}^n \lambda_i f_i \not\preceq^S \sum_{i=1}^n \lambda_i g_i$.

Since $\mathcal{D}$ is a convex set that does not intersect the negative orthant, it is standard to show that $\mathcal{D}$ can be separated by a hyper-plane associated with a prior probability over S. That is, there exists $p \in \Delta(S)$ such that $E_p(x) \geq 0$ for every $x \in \mathcal{D}$ (because $0 \in \mathcal{D}$). Such p, however, need not be unique. Denote the collection of all priors that separate $\mathcal{D}$ from the negative orthant by $\mathcal{P}$. Thus, if $x \in \mathcal{D}$, then $E_p(x) \geq 0$ for every $p \in \mathcal{P}$. In other words, if $f, g \in \mathcal{H}$ such that $f \succeq g$, then $E_p(u(f)) \geq E_p(u(g))$ for every $p \in \mathcal{P}$.

Assume now that there are $f, g \in \mathcal{H}$ such that $E_p(u(f)) \geq E_p(u(g))$ for every $p \in \mathcal{P}$. If $f \succeq^S g$, then $f \succeq g$ by monotonicity of $\succeq$ over $\mathcal{H}$. Otherwise, assume that $E_p(u(f)) > E_p(u(g))$ for every $p \in \mathcal{P}$. This means that there exist $f', g' \in \mathcal{H}$ and $0 < \lambda < 1$, such that $u(f') - u(g') \in \mathcal{D}$ and $u(f') - u(g') = \lambda(u(f) - u(g))$. By independence, we have that $f \succeq g$. Finally, if $E_p(u(f)) \geq E_p(u(g))$ for every $p \in \mathcal{P}$, then since $\mathcal{H}$ is convex and contains $\mathcal{F}_c$, we obtain that $\varepsilon f + (1-\varepsilon)l$ and $\varepsilon g + (1-\varepsilon)l'$ are both in $\mathcal{H}$, for every $l, l' \in L$. If $l \succ l'$, then $E_p(\varepsilon f + (1-\varepsilon)l) > E_p(\varepsilon g + (1-\varepsilon)l')$

[8]conv is the convex hull operator.

for all $p \in \mathcal{P}$. Thus $\varepsilon f + (1-\varepsilon)l \succeq \varepsilon g + (1-\varepsilon)l'$. From continuity we obtain that $f \succeq g$.

We conclude that for every $f, g \in \mathcal{H}$, $f \succeq g$ if and only if $E_p(u(f)) \geq E_p(u(g))$ for every $p \in \mathcal{P}$.

7.4. **Proof of Proposition 6.** The proof that (1) $\Leftrightarrow$ (2) follows exactly the same line of arguments as in the proof of Theorem 1. (2) $\Leftrightarrow$ (3) due to Lemma 1 (similar to the corresponding implications of Theorem 1).

7.5. **Proof of Proposition 3.** This is implied by Proposition 6.

7.6. **Proof of Theorems 3.** (1) $\Leftrightarrow$ (2). Define a binary relation $\succeq^*$ over all acts as follows: for every two acts g and h, $g \succeq^* h$ if and only if $\overline{U}(u(g)-u(h)) \geq 0$. We claim that $\succeq^*$ and $\succeq_L$ coincide. Suppose that $g \succeq^* h$ and assume to the contrary of the claim, that $g \not\succeq_P h$. **LR** implies that there exist two $\succeq$-comparable acts f_1, f_2 and $\alpha \in (0,1)$ such that $\alpha h + (1-\alpha)f_1 \succeq^S \alpha g + (1-\alpha)f_2$, and at the same time, $f_2 \succ f_1$. In particular $\mathbb{E}_p(u(f_2)) > \mathbb{E}_p(u(f_1))$. However, $\alpha h + (1-\alpha)f_1 \succeq^S \alpha g + (1-\alpha)f_2$ implies that $\frac{1-\alpha}{\alpha}[u(f_1) - u(f_2)] \geq u(g) - u(h)$. Thus, $\overline{U}(u(g)-u(h)) < 0$, which contradicts the assumption that $g \succeq^* h$.

On the other hand, assume that $g \succeq_L h$. By **LR**, $\overline{U}(u(g)-u(h))$ is at least 0, which implies that $g \succeq^* h$.

(2) $\Rightarrow$ (3) is due to (v) of the previous Lemma 1, where $X = u(f) - u(g)$ and $W = 0$. (2) $\Leftarrow$ (3) is a consequence of (iii) of Lemma 1 above.

References

[1] F. J. Anscombe, and R. J. Aumann, A definition of subjective probability, *Annals of Mathematical Statistics* **34** (1963), 199-295.

[2] T. F. Bewley, Knightian decision theory. Part 1, *Decisions in Economics and Finance* **25** (2002), 79-110.

[3] P. C. Fishburn, *Utility Theory for Decision Making.* New York: Wiley, (1970).

[4] P. Ghirardato, F. Maccheroni, and M. Marinacci, Differentiating ambiguity and ambiguity attitude, *Journal of Economic Theory* **118** (2004), 133-173.

[5] I. Gilboa, F. Maccheroni, M. Mariniacci, and D. Schmeidler, Objective and subjective rationality in a multiple prior model, *Econometrica* **78** (2010), 755-770.

[6] I. Gilboa, and D. Schmeidler, Maxmin expected utility with a non-unique prior, *Journal of Mathematical Economics* **18** (1989), 141-153.

[7] E. Karni, and M. L. Vierø, "Reverse Bayesianism": A choice-based theory of growing awareness, mimeo (2011), Johns Hopkins University.

[8] E. Lehrer, Partially-specified probabilities: decisions and games, (2007), to appear in American Economic Journal: Microeconomics.

[9] E. Lehrer, and R. Teper. Justifiable preferences, *Journal of Economic Theory* **146** (2011), 762-774.

[10] K. Nehring, Ambiguity in the context of probabilistic beliefs, mimeo (2001), UC Davis.

[11] M. Siniscalchi, A behavioral characterization of plausible priors, *Journal of Economic Theory* **128** (2006), 91-135.