Economic Theory **20**, 83–92 (2002)

Economic Theory

Revisiting Savage in a conditional world*

Paolo Ghirardato

Division of the Humanities and Social Sciences, California Institute of Technology,
Pasadena, CA 91125, USA (e-mail: paolo@hss.caltech.edu)

Received: December 1, 2000; revised version: February 26, 2001

Summary. I present an axiomatization of subjective expected utility and Bayesian updating in a conditional decision problem. This result improves our understanding of the Bayesian standard from two perspectives: 1) it uses a set of axioms which are weak and intuitive; 2) it provides a formal proof to results on the relation between dynamic consistency, expected utility and Bayesian updating which have never been explicitly proved in a fully subjective framework.

Keywords and Phrases: Subjective expected utility, Bayes's rule, Dynamic consistency, Consequentialism.

JEL Classification Numbers: D81.

Introduction

In the *Foundations of Statistics* [12], L. J. Savage shows that if a decision maker (DM)'s preference $\succcurlyeq$ on state-contingent payoffs ("acts" in his terminology) satisfies six axioms,[1] there exist a utility function u on the prize space $\mathscr{X}$ and a probability measure P on the state space Ω such that the preference can be represented by **subjective expected utility** (SEU): For all acts f and g,

$$f \succcurlyeq g \qquad \text{if and only if} \qquad \int_\Omega u(f(\omega))\,P(d\omega) \geq \int_\Omega u(g(\omega))\,P(d\omega). \tag{1}$$

* I am grateful to an anonymous referee, Matt Jackson, Peter Klibanoff, Tom Sargent, three generations of graduate students of Social Sciences at Caltech (Serena Guarnaschelli in particular) for helpful comments and encouragement.

[1] I restrict my attention to finite-valued acts. Savage's axiomatization comprised seven axioms, but the last is only used in obtaining the representation for infinite-valued acts.

Savage's axiomatization applies only to preferences in a *static* decision problem. That is, a problem in which the DM only takes a *single* decision as to which act is optimal. However, any dynamic decision problem can be reduced to a static problem by representing the DM's choice problem as that of choosing a contingent strategy for the dynamic problem. In this respect, it is useful to consider what we could call the **conditional** decision problem.[2] Suppose that the DM faces a static problem, either because the problem is static in its nature, or because she reduced a dynamic problem to a static one. Suppose that she can receive reliable information as to what is the real state of nature, in the form of an event A which contains such state, and that she is given a chance to review her preferences over acts in view of such information to obtain a **conditional** preference relation $\succcurlyeq_A$. In this problem, the DM is described by a *class* of conditional preferences, one for each possible event A. Her problem is conditional because her information at the time of (the single) choice is possibly finer than just the knowledge of Ω.

Given the conditional decision problem, assume that the *ex-ante* preference $\succcurlyeq_\Omega$, that I just denote $\succcurlyeq$, satisfies Savage's six axioms. Suppose that the conditional preferences satisfy the following condition: There is an act $\bar{h} \in \mathscr{F}$ such that for every event A and every pair of acts f and g,

$$f \succcurlyeq_A g \qquad \text{if and only if} \qquad f A \bar{h} \succcurlyeq g A \bar{h}, \tag{2}$$

where $f A \bar{h}$ (resp. $g A \bar{h}$) denotes the act which is equal to f (resp. g) on A and to $\bar{h}$ on $A^c \equiv \Omega \setminus A$. The choice of the act $\bar{h}$ is inconsequential as Savage's second axiom (P2) implies that the preference on the r.h.s. holds whenever $\bar{h}$ is substituted with any other h. In this case it is well-known that each preference $\succcurlyeq_A$ conditional on an event A such that $P(A) > 0$ can be represented as follows: For every pair of acts f and g,

$$f \succcurlyeq_A g \qquad \text{if and only if} \qquad \int_\Omega u(f(\omega))\, P_A(d\omega) \geq \int_\Omega u(g(\omega))\, P_A(d\omega), \tag{3}$$

where the measure P_A is the Bayesian update of P conditional on A. Thus, when the DM's conditional preference family satisfies Savage's axioms and the additional dynamic restriction imposed by Eq. (2), it is "as if" the DM constructs her *ex-post* preferences by updating her prior beliefs using Bayes's rule (and keeping her utility function fixed). We have thus obtained a simple axiomatic foundation to what I call the **Bayesian model**: The DM has SEU unconditional preferences and her conditional preferences are the result of Bayesian updating of her subjective prior.

The result just described crucially depends on the fact that the *ex-ante* preference satisfies Savage's axioms, in particular his axiom P2. Failing that, it is

[2] A conditional decision problem in this sense should not be confused with Luce and Krantz [8]'s "conditional decision structure", which is more general in allowing the state space to depend on the act the DM chooses. Also, their framework is fundamentally static in nature: Their DM does not envision the possibility of receiving information on the true state.

well-known that $\succcurlyeq$ does not necessarily have an SEU representation. For instance, it can have a "Choquet expected utility" representation in the sense of Schmeidler [13], where the probability P is not necessarily additive. Moreover, the conditional preferences constructed via Eq. (2) do not have to satisfy SEU and the posterior beliefs are not necessarily obtained by Bayesian updating of the prior beliefs (see Gilboa and Schmeidler [4]).

In this note, I show that the Bayesian model can be obtained even if we significantly weaken the requirements imposed on $\succcurlyeq$, provided that we impose stronger conditions on the relation between $\succcurlyeq$ and the conditional preferences $\succcurlyeq_A$. Precisely, I show that the result follows if the following two conditions are imposed: 1) conditional preferences are dynamically consistent in a very weak and intuitive sense; 2) for every A the preference conditional on A only depends on the behavior of acts on A (a property I call "consequentialism"). I modify Savage's axioms by removing his P2 axiom, and formulating a more intuitive dynamic version of his "state-independence" axiom (P3). Moreover, I do not impose the Bayesian updating property of Eq. (2). This provides a different axiomatic foundation for the Bayesian model, one which is more explicitly dynamic in spirit than the one outlined in the previous paragraphs.

The result I just described is likely to be unsurprising to experts. The fact that Savage's static axiom P2 can be substituted with a dynamic consistency property and consequentialism has been known at least ever since the seminal work of Arrow [1] (see also Myerson [11] and Hammond [5]). Granted that, it will be clear that I make no claim to the originality of the idea behind this result. The purpose of the analysis in this note is twofold: 1) to present a set of axioms which is, to the best of my knowledge, the most general and intuitive existing in the literature; 2) to prove in the general Savage set-up some results on dynamic consistency which, though simple and well-known to experts, have not been explicitly proved in the literature.

The note proceeds as follows. Section 1 spells out the set-up, the axioms and the main result. Section 2 contains the proof of the main result, which builds on three simple lemmas of some independent interest.

1 Set-up, axioms and representation theorem

The set-up of the analysis is that used by Savage [12]. I assume that a DM is faced with a decision problem whose state space is Ω, equipped with a σ-algebra $\mathscr{A} \subseteq 2^{\Omega}$. The set of possible **consequences** (outcomes of the decision problem) is $\mathscr{X}$. The objects of choice are **acts**, finite-valued (i.e., simple) measurable maps from Ω into $\mathscr{X}$, and the set of all such functions is labelled $\mathscr{F}$. As customary, I abuse notation and denote x the **constant act** yielding $x \in \mathscr{X}$ in every state of the world. Given event $A \in \mathscr{A}$ and acts $f, g \in \mathscr{F}$, I denote by $f\,A\,g$ the act h that is $h(\omega) = f(\omega)$ for $\omega \in A$ and $h(\omega) = g(\omega)$ for $\omega \in A^c$.

As explained in the Introduction, the decision problem thus described might be the static reformulation of a dynamic decision problem, where each strategy that the DM can implement corresponds to an act $f \in \mathscr{F}$. As customary in decision theory, we assume that the DM can express preferences over the comprehensive set $\mathscr{F}$; that is, she can judge counterfactual strategies.

The DM has a class of conditional preferences $\succcurlyeq_A$, one for each event in $\mathscr{A}$. There are two possible interpretations that we can give to each $\succcurlyeq_A$: 1) $\succcurlyeq_A$ represents how the DM thinks she would choose among the strategies in $\mathscr{F}$ if she were informed that only $\omega \in A$ can obtain; 2) $\succcurlyeq_A$ represents the *actual* preference that she would have over $\mathscr{F}$ if she was informed that only $\omega \in A$ can obtain. In both cases the *ex-ante* preference $\succcurlyeq$ represents the current preference of the DM over $\mathscr{F}$. All the axioms to be presented impose either within-preference restrictions on each $\succcurlyeq_A$, or on the relation between $\succcurlyeq_A$ and $\succcurlyeq$. Thus, they are compatible with either interpretation (in particular, no axiom deals with $\succcurlyeq_A$ and $\succcurlyeq_{A^c}$, so that counterfactual situations are never involved in checking an axiom). Consistently with my definition of the conditional decision problem, in either interpretation I assume that the DM chooses an act $f \in \mathscr{F}$ *only once*, after being informed of an event $A \in \mathscr{A}$.

From a strictly dynamic perspective, the set-up I use imposes a major restriction on preferences. Given acts $f, g \in \mathscr{F}$ and an event $A \in \mathscr{A}$, consider the act $f\,A\,g$. From an *ex-ante* perspective, this act mimics the following contingent strategy: Choose f if informed of A, and g otherwise. The set-up does not allow the DM to see the act and the contingent strategy as two different options, even though it is possible that she would. More generally, whatever the way the dynamic problem is originally presented, the set-up presupposes that she reduces it to a static problem. This implies that the DM must satisfy a subjective version of the “reduction of compound lotteries” axiom in the von Neumann-Morgenstern set-up. In that set-up, results like the ones to follow do not hold if reduction of compound lotteries is violated. Similarly, they do not necessarily hold in a subjective set-up if the DM does not reduce a dynamic problem to a static problem (say because she cares about the timing of the resolution of uncertainty, see Kreps and Porteous [7]).

The first axiom requires that each conditional preference be a **weak order** on $\mathscr{F}$. As usual, we then use $\succ_A$ (resp. $\sim_A$) to denote the asymmetric (resp. symmetric) component of a weak order $\succcurlyeq_A$.

Axiom 1 (Weak Order) *For each $A \in \mathscr{A}$, $\succcurlyeq_A$ is a complete and transitive binary relation on $\mathscr{F}$.*

Recall Savage’s definition of a null event [12]: We say that $A \in \mathscr{A}$ is **null** (w.r.t. $\succcurlyeq$) if for every $f, g, h, h' \in \mathscr{F}$,

$$f\,A^c\,h \succcurlyeq g\,A^c\,h' \qquad \text{if and only if} \qquad f \succcurlyeq g.$$

Let $\mathscr{A}'$ denote the subset of $\mathscr{A}$ containing all the non-null events.

The following two axioms impose *dynamic* restrictions on the DM’s conditional preferences which involve only events in $\mathscr{A}'$. (Those outside $\mathscr{A}'$ ultimately do not matter to the DM.) They are key to the representation. The first says that preferences are dynamically consistent whenever a non-null event obtains.

Axiom 2 (Dynamic Consistency) *For all* $A \in \mathscr{A}'$ *and* $f,g \in \mathscr{F}$*, both the following conditions hold:*

(a) *If* $f \succsim_A g$ *then* $f\,A\,g \succsim g$*;*
(b) *If* $f\,A\,g \succsim g$ *then* $f \succsim_A g$*.*

Part (a) of the axiom is interpreted as follows: Suppose that the DM (thinks she) would prefer f to g if she was told A. Consider the act $f\,A\,g$. As discussed earlier, this act mimics what the DM (thinks she) would be able to achieve if she could postpone her choice between f and g to after knowing whether A or A^c. The axiom requires that, from an *ex-ante* perspective, if her default option is to choose g, the possibility of thus "postponing" her choice does not make her worse off. In this sense, part (a) says that *information is valuable* to our DM. Part (b) is a (forward-looking) requirement of *consistency of implementation* of preferred strategies: If $f\,A\,g \succsim g$ and A is non-null, then if the DM observes A she must prefer her contingent plan f over g. Many stronger versions of properties (a) and (b) have been proposed in the literature. To the best of my knowledge, Arrow [1] (his "dominance" axiom) and Myerson [11] (his "substitution" axioms) were among the earliest proposers.

As its name makes obvious, the next dynamic axiom just says that the DM's "ordinal" preferences over consequences (i.e. constant acts) are identical over the non-null events.

Axiom 3 (Ordinal Preference Consistency) *For all* $A \in \mathscr{A}'$ *and* $x,y \in \mathscr{X}$*,*

$$x \succsim y \qquad \text{if and only if} \qquad x \succsim_A y.$$

The next three axioms impose some mild structure on the *ex-ante* preference $\succsim$.

Axiom 4 (Likelihood Payoff Independence) *For all* $A,B \in \mathscr{A}$ *and all* x,x', $y,y' \in \mathscr{X}$ *such that* $x \succ y$ *and* $x' \succ y'$*,*

$$x\,A\,y \succsim x\,B\,y \qquad \text{if and only if} \qquad x'\,A\,y' \succsim x'\,B\,y'.$$

This is Savage's axiom P4. As the name says, the axiom implies that the likelihood relation (a weak order itself) on events that we can derive by looking at the DM's over bets of the form $x\,A\,y$ (for $x \succ y$) is independent of the exact "amounts" x and y (as long as the payoff for A is better than that for A^c). The next axiom, Savage's axiom P5, is dispensable, but nothing would be gained conceptually by doing so.

Axiom 5 (Nontriviality) *There are* $x,y \in \mathscr{X}$ *such that* $x \succ y$*.*

Finally, we have a continuity axiom:

Axiom 6 (Archimedean) *If* $f,g \in \mathscr{F}$ *are such that* $f \succ g$ *and* $x \in \mathscr{X}$ *then there is a finite partition* $\mathscr{H}$ *of* Ω *such that, for every* $H \in \mathscr{H}$*:*

(i) $x\,H\,f \succ g$*,*

(ii) $f \succ x\, H\, g$.

The last axiom is also static. Differently from the previous three, it restricts every preference $\succcurlyeq_A$ conditional on a non-null A, rather than just $\succcurlyeq$.

Axiom 7 (Consequentialism) *For any $A \in \mathscr{A}'$ and all $f, g \in \mathscr{F}$, $f(\omega) = g(\omega)$ for each $\omega \in A$ implies $f \sim_A g$.*

This axiom says that the preference conditional on non-null A should not depend on how the strategy f behaves in the counterfactual states of A^c (in other words, it should only depend on the truncation $f|_A$). The name comes from Hammond [5], even though Hammond's notion is conceptually stronger than axiom 7 (it is much closer to Arrow's [1] "conditional preference" axiom).

It is perhaps helpful at this point to summarize the conditions that I have imposed on the different preferences. I have assumed that the *ex-ante* preference $\succcurlyeq$ satisfies weak order, likelihood payoff independence, nontriviality and the archimedean axiom. These axioms by themselves impose little structure on the representation of $\succcurlyeq$. For instance, they allow $\succcurlyeq$ to be a Choquet expected utility (CEU) preference in the sense of Schmeidler [13]. I have assumed that every preference $\succcurlyeq_A$ conditional on non-null A satisfies weak order and consequentialism. Again, this imposes very little structure on such preferences. For instance, any conditional preference obtained from a CEU preference by one of the updating rules discussed in Gilboa and Schmeidler [4] satisfies these properties. Finally, we have imposed two dynamic axioms: dynamic consistency and ordinal preference consistency. The latter is natural in a situation where we expect state-dependence not to be an issue. The former is crucial in forcing all the preferences to have an expected utility representation, and it is the one with the most significant empirical content.

Recall now that a probability **charge** is a finitely additive normalized set-function, and that a probability charge Q is **convex-ranged** if for every $A \in \mathscr{A}$ and every $\alpha \in [0, Q(A)]$, there is $B \subseteq A$ such that $Q(B) = \alpha$. We can now state the main result of this note:

Theorem 1 *Suppose that a DM's conditional preferences are represented by a class of binary relations $\{\succcurlyeq_A\}_{A \in \mathscr{A}}$. Then the following are equivalent:*

(i) The class $\{\succcurlyeq_A\}_{A \in \mathscr{A}}$ satisfies axioms 1–7;
(ii) There is a non-constant utility index $u : \mathscr{X} \to \mathbf{R}$, unique up to a positive affine transformation, and a unique convex-ranged probability charge $P : \mathscr{A} \to [0, 1]$ such that for all $f, g \in \mathscr{F}$, Eq. (1) holds. Moreover, for each $A \in \mathscr{A}$ such that $P(A) > 0$, $\succcurlyeq_A$ is represented as in Eq. (3), with P uniquely replaced by $P_A : \mathscr{A} \to [0, 1]$, a convex-ranged probability charge defined as follows: For each $B \in \mathscr{A}$,

$$P_A(B) = \frac{P(A \cap B)}{P(A)}. \tag{4}$$

Thus, axioms 1–7 imply that the *ex-ante* preference $\succcurlyeq$ has an SEU representation with beliefs P, and that for each $A \in \mathscr{A}'$, the conditional preference $\succcurlyeq_A$ is obtained from $\succcurlyeq$ by updating P by Bayes's rule (which applies, since $P(A) > 0$ for each non-null A).

Remark 1 Epstein and Le Breton [2] present a result analogous to Theorem 1 which provides an axiomatic foundation to the weaker probabilistic sophistication model of Machina and Schmeidler [10]. A version of their result can be stated in the notation of this paper. Suppose that we strengthen axiom 4 as follows: For every $C \in \mathscr{A}'$, all $A, B \in \mathscr{A}$ such that $A \cup B \subseteq C$, all $x, x', y, y' \in \mathscr{X}$ such that $x \succ y$ and $x' \succ y'$, and all $h \in \mathscr{F}$,

$$\begin{bmatrix} x, & A \\ y, & C \setminus A \\ h, & C^c \end{bmatrix} \succcurlyeq_C \begin{bmatrix} x, & B \\ y, & C \setminus B \\ h, & C^c \end{bmatrix} \Leftrightarrow \begin{bmatrix} x', & A \\ y', & C \setminus A \\ h, & C^c \end{bmatrix} \succcurlyeq_C \begin{bmatrix} x', & B \\ y', & C \setminus B \\ h, & C^c \end{bmatrix}.$$

This just says that every conditional preference should satisfy a version of the payoff independence property.[3] The result is: Suppose that a class $\{\succcurlyeq_A\}_{A \in \mathscr{A}}$ satisfies axiom 1, a slightly different dynamic consistency property (see Lemma 1 below), axiom 10, the stronger version of axiom 4, a slightly stronger version of 5, and 6. Then $\succcurlyeq$ is probabilistically sophisticated, in the sense that its likelihood relation can be represented by a probability charge P. Moreover the conditional preference $\succcurlyeq_A$, for $A \in \mathscr{A}'$, also induces a likelihood relation which is representable by the probability P_A.

2 Proof of Theorem 1

The proof uses three lemmas which are of independent interest. The first is the simple observation that in the presence of axiom 7, axiom 2 is equivalent to a slightly weaker form of the dynamic consistency property that is commonly used in the literature (see, e.g., Epstein and Le Breton [2]). The latter differs in restricting also the preferences conditional on null events. It implies that every such preference must be trivial, an unnecessary restriction.

Lemma 1 *Let the class $\{\succcurlyeq_A\}_{A \in \mathscr{A}}$ satisfy axioms 1 and 7. Then it satisfies axiom 2 if and only if for any $A \in \mathscr{A}'$ and all $f, g, h \in \mathscr{F}$,*

$$f A h \succcurlyeq g A h \qquad \text{if and only if} \qquad f A h \succcurlyeq_A g A h.$$

That is, axiom 2 is tantamount to requiring that the preference conditional on non-null A between acts which are identical off A conform to the *ex-ante* preference.

Proof "Only if": Let $g' = g A h$ and $f' = f A h$ and notice that axiom 2 and $f' \succcurlyeq_A g'$ imply $f' A g' \succcurlyeq g'$. This shows that $f A h \succcurlyeq_A g A h$ implies $f A h \succcurlyeq g A h$. We need to show that the converse implication holds. By the definition of g', axiom 2 and axiom 7, we have

[3] [2] contains a different version of this axiom. The present one is seen to be equivalent to theirs and it is is more directly interpretable this way.

$$f\,A\,h \succcurlyeq g\,A\,h \Longleftrightarrow f\,A\,g' \succcurlyeq g' \quad \Longrightarrow \quad f \succcurlyeq_A g' \\ \Longleftrightarrow \quad f\,A\,g' \succcurlyeq_A g' \Longleftrightarrow f\,A\,h \succcurlyeq_A g\,A\,h.$$

"If": Follows from the observation that under axiom 7, $f \succcurlyeq_A g$ iff $f\,A\,h \succcurlyeq_A g\,A\,h$. □

The second lemma contains a result which really belongs to the "folk wisdom" of decision theory (cf. Karni and Schmeidler [6] and Epstein and Le Breton [2]). It shows that axioms 2 and 7 imply that $\succcurlyeq$ satisfies Savage's axiom P2, that he called the "sure-thing principle", and that the preferences conditional on $A \in \mathscr{A}'$ are obtained from $\succcurlyeq$ by "Bayesian updating". Formally:

Axiom 8 (Sure Thing Principle) *For any $A \in \mathscr{A}$ and all $f, g, h, h' \in \mathscr{F}$,*

$$f\,A\,h \succcurlyeq g\,A\,h \qquad \text{if and only if} \qquad f\,A\,h' \succcurlyeq g\,A\,h'.$$

Axiom 9 ($\bar{h}$-Bayesian Updating) *There exists $\bar{h} \in \mathscr{F}$ such that for all $A \in \mathscr{A}'$ and all $f, g \in \mathscr{F}$, Eq. (2) holds.*

As I discussed in the Introduction, Bayesian updating is thus called because in the presence of Savage's axioms it implies that the conditional preference can be constructed by looking at the DM's posterior beliefs. Savage used it as a definition of a relation that he suggested *interpreting* as "conditional preference" (though he only used it as a technical construct). Gilboa and Schmeidler [4] suggested it as an updating criterion for general non-SEU preferences.

Lemma 2 *Let the class $\{\succcurlyeq_A\}_{A\in\mathscr{A}}$ satisfy axioms 1 and 7. Then $\{\succcurlyeq_A\}_{A\in\mathscr{A}}$ satisfies axiom 2 if and only if $\succcurlyeq$ satisfies axiom 8, and $\{\succcurlyeq_A\}_{A\in\mathscr{A}}$ satisfies axiom 9.*

Proof "Only if": If we replace axiom 2 with the equivalent property stated in Lemma 1, the proof of the first statement follows immediately from adjoining axioms 2, 7 and 2. The second statement follows from adjoining axioms 2 and 7 (notice also that *every* $h \in \mathscr{F}$ can be used as $\bar{h}$). "If": Suppose that $\{\succcurlyeq_A\}_{A\in\mathscr{A}}$ satisfies axiom 9 and $\succcurlyeq$ satisfies axiom 8. It then follows that for every $A \in \mathscr{A}'$ and $f, g, h \in \mathscr{F}$ we have:

$$f\,A\,h \succcurlyeq g\,A\,h \Longleftrightarrow f \succcurlyeq_A g \Longleftrightarrow f\,A\,h \succcurlyeq_A g\,A\,h.$$

We then invoke again Lemma 1 to conclude the proof. □

Though trivial, this result has great conceptual relevance: It implies that many non EU preferences cannot be applied to a dynamic framework in a way which satisfies both consequentialism and dynamic consistency, since they are interesting only if they violate the sure-thing principle (for a discussion, see Machina [9]).

The last lemma, which follows immediately from axiom 9, shows that the axioms imply that Savage's 'monotonicity' axiom P3 holds for $\succcurlyeq$. In the framework of this paper, the axiom is stated as follows:

Axiom 10 (Eventwise Monotonicity) *For all* $A \in \mathscr{A}'$ *and all* $x, y \in \mathscr{X}$, $h \in \mathscr{F}$,

$$x\,A\,h \succcurlyeq y\,A\,h \qquad \text{if and only if} \qquad x \succcurlyeq y.$$

The lemma is:

Lemma 3 *Let the class* $\{\succcurlyeq_A\}_{A\in\mathscr{A}}$ *satisfy axioms 1, 2, and 7. Then it satisfies axiom 3 if and only if it satisfies axiom 10.*

We can finally pull all the strands together and complete the proof of Theorem 1.

Proof of Theorem 1. The proof that $(ii) \Rightarrow (i)$ is straightforward. I now prove that $(i) \Rightarrow (ii)$: The previous lemmas have shown that $\succcurlyeq$ satisfies Savage's axioms P2 and P3. The other axioms — P1, P4, P5 and P6 — are respectively implied by axioms 1, 4, 5 and 6. We can therefore apply Savage's theorem (see Fishburn [3, Section 14.2 and 14.3]) to show that $\succcurlyeq$ has a SEU representation as in (1), with utility u and probability P. This proves the first statement of (ii). As for the second, suppose that $A \in \mathscr{A}'$. First, we observe that $P(A) > 0$, so that P_A is well-defined. (Its range convexity follows from that of P.) By axiom 9, we have that

$$\begin{aligned} f \succcurlyeq_A g \quad &\Longleftrightarrow \quad f\,A\,\bar{h} \succcurlyeq g\,A\,\bar{h} \\ &\Longleftrightarrow \quad \int_A u(f(\omega))\,P(d\omega) \geq \int_A u(g(\omega))\,P(d\omega) \\ &\Longleftrightarrow \quad (1/P(A)) \int_\Omega u(f(\omega))\,P_A(d\omega) \geq (1/P(A)) \int_\Omega u(g(\omega))\,P_A(d\omega). \end{aligned}$$

Thus, $\succcurlyeq_A$ is represented by SEU with utility u and beliefs P_A, as claimed. □

References

1. Arrow, K. J.: Exposition of a theory of choice under uncertainty. In: Essays in the theory of risk-bearing, chapter 2. Amsterdam: North-Holland 1974. (Part of the Yriö Jahnssonin Säätio lectures in Helsinki, 1965)
2. Epstein, L. G., Le Breton, M.: Dynamically consistent beliefs must be bayesian. Journal of Economic Theory **61**, 1–22 (1993)
3. Fishburn, P. C.: Utility theory for decision making. New York: Wiley 1970
4. Gilboa, I., Schmeidler, D.: Updating ambiguous beliefs. Journal of Economic Theory **59**, 33–49 (1993)
5. Hammond, P.J.: Consequentialist foundations for expected utility. Theory and Decision **25**, 25–78 (1988)
6. Karni, E., Schmeidler, D.: Atemporal dynamic consistency and expectded utility theory. Journal of Economic Theory **54**, 401–408 (1991)
7. Kreps, D. M., Porteus, E. L.: Temporal resolution of uncertainty and dynamic choice theory. Econometrica **46**, 185–200 (1978)
8. Luce, R. D., Krantz, D. H.: Conditional expected utility. Econometrica **39**, 253–271 (1971)
9. Machina, M. J.: Dynamic consistency and non–expected utility models of choice under uncertainty. Journal of Economic Literature **27**, 1622–1668 (1989)
10. Machina, M. J., Schmeidler, D.: A more robust definition of subjective probability. Econometrica **60**, 745–780 (1992)

11. Myerson, R. B.: An axiomatic derivation of subjective probability, utility, and evaluation functions. Theory and Decision **11**, 339–352 (1979)
12. Savage, L. J.: The Foundations of statistics. New York: Wiley 1954
13. Schmeidler, D.: Subjective probability and expected utility without additivity. Econometrica **57**, 571–587 (1989)