Available online at www.sciencedirect.com
ScienceDirect
Journal of Mathematical Economics 43 (2007) 888–899
JOURNAL OF Mathematical ECONOMICS
www.elsevier.com/locate/jmateco

Updating Choquet beliefs

Jürgen Eichberger [a], Simon Grant [b,*], David Kelsey [c]

[a] Universität Heidelberg, Germany
[b] Rice University, United States
[c] Exeter University, United Kingdom

Received 22 February 2006; received in revised form 30 March 2007; accepted 3 April 2007
Available online 22 May 2007

Abstract

We apply Pires's coherence property between unconditional and conditional preferences that admit a CEU representation. In conjunction with consequentialism (only those outcomes on states which are still possible can matter for conditional preference) this implies that the conditional preference may be obtained from the unconditional preference by taking the Full Bayesian Update of the capacity.
© 2007 Elsevier B.V. All rights reserved.

JEL Classification: D81

Keywords: Updating ambiguous beliefs; Full Bayesian Updating; Choquet Expected Utility

1. Introduction

In recent years there has been a mounting challenge to the standard model of decision under uncertainty, the subjective expected utility model. In particular, models in which beliefs cannot be represented by a single probability measure over the set of events, have been introduced as a way to formalize the distinction drawn by Knight (1921), Keynes (1921) and others between situations of risk (where probabilities are based upon an extensive data base of past relevant cases or can be readily gleaned from the structure and nature of the situation) and uncertainty (where probabilities are not well-known or agreed upon). For example, in the Choquet Expected Utility model (Schmeidler, 1989) beliefs are represented by a capacity, a measure that is not necessarily additive, while in the multiple-prior model (Gilboa and Schmeidler, 1989), as its name suggests, beliefs are represented by a set of probability measures.

* Corresponding author. Tel.: +1 713 348 3332; fax: +1 713 348 5278.
E-mail address: sgrant@rice.edu (S. Grant).

0304-4068/$ – see front matter © 2007 Elsevier B.V. All rights reserved.
doi:10.1016/j.jmateco.2007.04.005

There has been some success in applying these static models to individual decision making problems as well as many-agent settings including those involving strategic interaction (see for example, Dow and Werlang, 1992; Mukerji, 1998; Mukerji and Tallon, 2001, 2004; Eichberger and Kelsey, 2000, 2002). But in order to be able to apply non-additive beliefs models to sequential or dynamic settings requires a theory of how preferences are updated as new information arrives.

In subjective expected utility, the almost universally applied rule is Bayesian updating. For non-additive beliefs there have been two major approaches in the literature. The first is a statistical approach that considers for different updating rules the statistical properties of the updated beliefs that are derived from such rules. Examples include Denneberg (1994, 2002), Jaffray (1992), Lapied and Kast (2005), Lehrer (2005) and Shafer (1976). The other approach is decision-theoretic. The updating rule arises from axioms on the preferences both unconditional and conditional: a non-exhaustive list of examples includes Epstein and Schneider (2003), Gilboa and Schmeidler (1993), Hanany and Klibanoff (2006), Pires (2002), Siniscalchi (2006), Sarin and Wakker (1998), Walley (1991) and Wang (2003). This paper follows the decision-theoretic approach but we shall show in the sequel how the rule we obtain is the Full Bayesian Updating rule of Jaffray (1992).

For purposes of setting the benchmark against which alternatives will be introduced, motivated and developed, let us first consider how the Bayesian Updating rules for subjective expected utility preferences arises from the following natural method to deduce the conditional preferences from the prior preferences. Suppose we wish to deduce from the prior preference relation the preference between two acts, say f and g, conditional on the event E having obtained. Savage invoking his ‘sure-thing principle’ would argue it is enough to look at the unconditional preference between any pair of acts, say f' and g' which *agree* on states outside the conditioning event E, whereas for states in E, f' agrees with f and g' agrees with g. Since f' and g' agree on the states outside the conditioning event, Savage argues it is reasonable to assume that only how they differ on states in E will (or should) be decisive in determining the (unconditional) preference between them. If f' is preferred to g' unconditionally, then conditional on knowing that E has obtained, the individual should also prefer f to g (and indeed she should also still prefer f' to g'). For the unconditional preference, as f' and g' agree on states outside E, should E not obtain, then both acts will lead to the same outcome. For the conditional preference relation, knowing that E has obtained makes what f and g might have led to on states outside of E, immaterial.

Such reasoning embodies two properties:

1. consequentialism—only those outcomes on states which are still possible can matter for preference; and
2. dynamic consistency—if for an unconditional preference relation we have one act is preferred to another, then conditional on knowing that the complement of an event on which the two acts agree has obtained, the conditional preference should also have the latter act preferred to the former.

So in our example, consequentialism requires that if conditional on knowing that E has obtained, f is preferred to g, then the conditional preference relation should also have f' be preferred to g'. Dynamic consistency requires that if the unconditional preference relation has f' preferred to g', then the conditional preference relation of the individual after she learns that E has obtained, should also have f' preferred to g'.

In the extant literature, Hanany and Klibanoff (2006), Siniscalchi (2006) and Sarin and Wakker (1998) drop consequentialism and retain dynamic consistency. Epstein and Schneider (2003) show it is possible to retain both if one restricts the domain of acts and conditioning events (or more precisely, filtrations) over which preferences are defined. Wang (2003) casts his analysis in a more complicated setting of consumption–information profiles which have no direct counter-part in a standard Savage act framework, but effectively he is imposing similar restrictions to those of Epstein and Schneider on the domain of admissible problems.[1] We shall maintain, however, an unrestricted domain of acts and conditioning events and follow Gilboa and Schmeidler (1993), Pires (2002) and Walley (1991) in retaining consequentialism and dropping dynamic consistency.

Our reason for retaining consequentialism and dropping dynamic consistency is because, to paraphrase Baron and Frisch (1988), we feel ambiguity arises in a fundamental sense from uncertainty about probability created by missing information that is relevant and could be known. Hence once an event is known to have obtained, the only remaining ambiguity the individual faces relates to uncertainty about the probabilities of subevents of that event. Past (or borne) uncertainty one may have had about the probability of counterfactual event and its subsets are no longer relevant. But such uncertainty might have been relevant to the individual at the time when she did not know whether the event or its complement had obtained, and so such ambiguity that she perceived there to have been ex ante, may well have had an impact on her unconditional preferences. As we shall illustrate by way of example in Section 3 below, this may well lead to violations of dynamic consistency.

Our aim in this paper is to axiomatize an updating rule for a Choquet Expected Utility maximizer (that is, the extension of subjective expected utility where beliefs are represented by a capacity rather than an [additive] probability measure). After establishing the analytical framework in Section 2, we introduce, in Section 3, the axiom Pires (2002) proposed to link the unconditional and conditional preferences. This axiom, which we dub, *Conditional Certainty Equivalent Consistency*, has a similar intuitive appeal to Savage's sure-thing principle but is weak enough to accommodate standard Ellsberg type behavior. The representation result that we derive for the case where the unconditional and conditional preferences are all members of the Choquet Expected Utility family of preferences, shows that the utility function over outcomes is invariant to updating, while the updated capacity may be obtained using Jaffray's (1992) Full Bayesian Updating rule (or Walley's, 1991, Generalized Bayesian Updating rule) on the unconditional capacity.

2. Setup

We present our analysis in the context of a framework of purely subjective uncertainty. We take the uncertainty a decision maker faces to be described by a finite set of *states*, denoted by S. Associated with the set of states is the set of events, taken to be the set of subsets of S, denoted by $\mathcal{E}$. For each $E \in \mathcal{E}$, E^c shall denote its complement.

Let X, the set of outcomes, be a connected and separable topological space. An *act* is a function $f : S \to X$. $\mathcal{F}$ denotes the set of such acts and is endowed with the product topology induced by the topology on X. We shall identify each $x \in X$ with the constant act, $f(s) = x$ for all $s \in S$. For any pair of acts f, g in $\mathcal{F}$ and any event $E \in \mathcal{E}$, $f_E g$ will denote the act $h \in \mathcal{F}$, formed by 'splicing' the two acts f and g, in which $h(s)$ equals $f(s)$ if $s \in E$, and equals $g(s)$ if $s \notin E$. In general, for any finite partition $\{E_1, \ldots, E_n\}$ of S and any list of n acts $(f^1, \ldots, f^n)$, let $f^1_{E_1} f^2_{E_2} \ldots f^{n-1}_{E_{n-1}} f^n$ be the act that yields $f^i(s)$ if s is in E_i.

[1] Eichberger et al. (2005) also restrict preferences over information structure for a fixed filtration. But for the particular family of non-additive measures they consider for beliefs, they show a necessary and sufficient condition for dynamic consistency is that beliefs be additive over the final stage in the filtration.

We assume that the decision maker is characterized by a family of conditional preference relations on $\mathcal{F}$. For each event $E \in \mathcal{E}$, let $\succsim_E$ denote the preferences over acts *given* E. That is, we shall interpret $\succsim_E$ as the agent's preferences if she knew that E had obtained. As usual $\succ_E$ and $\sim_E$ will denote the asymmetric and symmetric parts of $\succsim_E$, respectively. The relation $\succsim$ shall denote the individual's unconditional preference relation on $\mathcal{F}$ (that is, $\succsim \, = \, \succsim_S$).

We say f and g are comonotonic if for every pair of states s and s' in S, $f(s) \succ f(s')$ implies $g(s) \succsim g(s')$. Given a preference relation $\succsim_E$, an event $A \in \mathcal{E}$ is $\succsim_E$*-null* if $f_A g \sim_E g$ for all pairs of acts $f, g \in \mathcal{F}$. Let $\mathcal{N}_E$ denote the set of $\succsim_E$*-null* events (and $\mathcal{N}$ denote the set of [unconditional] null events, that is, $\mathcal{N} = \mathcal{N}_S$).

For ease of exposition (and without any essential loss of generality) we shall assume the existence of a best and worst outcome, namely that there exist outcomes 0 and M in X, such that $M \succ 0$ and $M \succsim x \succsim 0$ for all $x \in X$.

3. Connecting conditional and unconditional preferences

The first property we require for the conditional preferences is consequentialism. That is, the conditional preferences are 'forward-looking' in the sense that what happens off the conditioning event should not be able to affect the conditional preference between any pair of acts.

Axiom 1 *(Consequentialism).* Fix an event $E \in \mathcal{E}$. The event E^c is $\succsim_E$*-null*. That is, $f_E g \sim_E f$ for all $f, g \in \mathcal{F}$.

This is a particularly desirable property to have in order to keep the preference model tractable in applications, since it means for a given conditional preference relation $\succsim_E$, we do not need explicitly to keep track of what outcomes might have resulted from an act had a state outside of E obtained.[2]

The next two axioms connect the conditional to the unconditional preferences. The first simply requires that the ordering of outcomes be the same across events. Notice that in conjunction with the existence of an unconditionally best outcome and worst outcome, this entails that for every event E, $\succsim_E$ is non-degenerate.

Axiom 2 *(State Independence).* For any pair of outcomes x, y in X, and any event $E \in \mathcal{E}$, $x \succsim y$ if and only if $x \succsim_E y$.

Our third axiom is adapted from Pires (2002). It says that if conditional on E obtaining, the decision maker is indifferent between the act f and the outcome x, then her unconditional preferences should also express indifference between the outcome x and the act $f_E x$, that is the act that agrees with f on E and agrees with x on the complement of E. Conversely, if her unconditional preferences express indifference between the outcome x and the act $f_E x$, then conditional on E obtaining, she should remain indifferent between the act f and the outcome x.

Axiom 3 *(Conditional Certainty Equivalent Consistency).* For any unconditionally non-null event $E \notin \mathcal{N}$ any outcome x in X, and any act f in $\mathcal{F}$, $f \sim_E x$ if and only if $f_E x \sim x$.

[2] The fact that by definition the conditional preference relation $\succsim_E$ does not depend on the act the individual may have chosen ex ante, already embodies much of the force of consequentialism. More properly, consequentialism should be viewed as the joint hypothesis that $\succsim_E$ does not depend on the act the individual chose ex ante and that the event E^c is $\succsim_E$-null. For example, Hanany and Klibanoff (2006) relax consequentialism by *dropping* the former while retaining the latter.

It is reminiscent of Savage's sure-thing principle. If on knowing that E had obtained, an individual would be indifferent between the act f and the outcome x, then from the perspective of his unconditional preferences he should be indifferent between x for sure, and the act $f_E x$, that is, the act that coincides with f on E and yields x should a state outside of E obtain. From the ex ante perspective, one might justify this with the following reasoning: suppose the decision maker runs the thought experiment in which he imagines he is going to learn whether the state of the world is in E or is not. If he learns the state is in E, then he knows he will be indifferent between $f_E x$ and x. On the other hand, if he learns it is not in E, then he knows he will receive x for sure. Hence, since he anticipates he will be either in a situation in which he is indifferent between $f_E x$ and the outcome x or in a situation in which he receives x for sure, he reasons he should be indifferent between $f_E x$ and x now, when he does not know whether the state is in E or its complement.

A stronger requirement is the following property that appears in Skiadas (1997).

Strict coherence: For any non-null event $E \notin \mathcal{N}$ and any pair of acts g and h in $\mathcal{F}$:

1. $g \succsim_E h$ and $g \succsim_{E^c} h$ implies $g \succsim h$.
2. $g \succ_E h$ and $g \succcurlyeq_{E^c} h$ implies $g \succ h$.

Strict coherence precludes any 'hedging' benefits (or costs) that might be associated with the unconditional preferences. Indeed, in conjunction with consequentialism, strict coherence essentially entails that the unconditional preferences are additively separable across states (see Skiadas, 1997, Section 4 for details). Conditional Certainty Equivalent Consistency does not rule out non-neutral attitudes towards 'hedging' because when we consider the implication for the unconditional preference, we do so only for acts that are constant and agree on the complementary event. Roughly speaking, Conditional Certainty Equivalent Consistency is the restriction of strict coherence to pairs of acts g and h, where $h = x$ is a constant act, and $g(s) = x$ for all $s \notin E$.

To see why strict coherence might not hold in the context of ambiguous beliefs, consider the following 'Ellsberg-type urn example. The urn contains 100 balls numbered 1–200. The balls numbered 1 to 66 are red. The balls numbered from 67 to $200 - 2n$ are black and the remainder (that is, those numbered from $201 - 2n$ to 200) are white. The only information the decision maker has about n is that it is an integer and that $1 \leq n \leq 66$.

Let O (respectively, E) be the event that the ball drawn from the urn has an odd (respectively, even) number on it. Let R (respectively, B, W) be the event that the ball drawn is red (respectively, black, white) in color. Let OR be the event that the ball drawn from the urn has an odd number and its color is red, and so on. Consider the pair of acts: $g = M_R 0$ and $h = M_{OB} M_{EW} 0$. Given her information, the individual knows that she will win with g if any one of the 66 red balls in the urn is drawn. Similarly, she knows that she will win with h, if any one of the 67 balls that are odd and black or that are even and white is drawn. Thus we would expect for her unconditional preferences she would have $h \succ g$.

But what about her conditional preferences relations $\succsim_O$ and $\succsim_E$? If she knows that the ball drawn has an odd number on it, then out of the 100 balls with odd numbers, she knows 33 are red, but out of the remaining 67 balls all she knows is that at least one of them is black and at least one of them is white. If she is averse to bets with ambiguous odds, we may well expect for conditional preference relation to have $g \succ_O h$, and by similar reasoning $g \succ_E h$. A violation of strict coherence.

On the other hand, if she expressed conditional on knowing the number of the ball drawn was odd, an indifference between the act $f = M_B 0$ (that is, betting on the color of the ball drawn being black) and say the amount for sure $x = 0.30 \times M$, that is, $f \sim_O x$, then Conditional Certainty

Equivalent Consistency requires her to express an unconditional indifference between the act $f_O x = M_{OB} 0_{OR} 0_{OW} x$ and getting x for sure, that is $f_O x \sim x$. But this by itself does not seem to contradict any of our intuitions for the individual's conditional or unconditional preferences.

4. The representation result

Pires conducted her analysis in the context of the multiple priors model of Gilboa and Schmeidler (1989). That is, she imposed the axioms of Gilboa and Schmeidler (1989) so that each preference relation $\succsim_E$ admitted a representation of the form

$$f \succsim_E g \Leftrightarrow \min_{p \in \Delta_E} \int u \circ f \,\mathrm{d}p \geq \min_{p \in \Delta_E} \int u \circ g \,\mathrm{d}p,$$

where Δ_E is a convex set of probability measures over S, with the property: $p \in \Delta_E$ implies $p(E) = 1$. Let $\Delta = \Delta_S$, that is, the set of priors for the unconditional preference relation $\succsim$. Pires shows that if such a family of multiple prior preferences satisfy *Conditional Certainty Equivalent Consistency* then for any event E and any $p \in \Delta$, such that $p(E) > 0$, there exists $q \in \Delta_E$, where q is the Bayesian Update of p, that is,

$$q(A) = \frac{p(A \cap E)}{p(E)}, \quad \text{for any} \quad A \in \mathcal{E}.$$

That is, for any prior probability that gives positive weight to the conditioning event, its Bayesian Update is an element of the updated set of multiple priors. Furthermore, if for every $p' \in \Delta$, $p'(E) > 0$, then for every $q \in \Delta_E$, there exists $p \in \Delta$, such that q is the Bayesian Update of p. That is, if every prior probability gives positive weight to the conditioning event, then the set of updated priors is obtained by updating all the prior probability measures using Bayes' rule.

In this section, we explore what the axioms introduced in the previous section imply for a family of preferences in which each preference relation $\succsim_E$ admits a Choquet Expected Utility (CEU) representation. The Choquet Expected Utility of an act f, is taken with respect to a utility index over outcomes, $u : X \to \mathbb{R}$, and a normalized and monotonic set function (or capacity) $\nu : \mathcal{E} \to [0, 1]$, that satisfies $\nu(\emptyset) = 0$, $\nu(S) = 1$ and $A \subset B \Rightarrow \nu(A) \leq \nu(B)$, for all $A, B \in \mathcal{E}$.

Definition 1. Fix a capacity $\nu : \mathcal{E} \to [0, 1]$. The c onjugate capacity, denoted $\bar{\nu}$, is defined as $\bar{\nu}(E) = 1 - \nu(E^c)$.

Since every act is finite-ranged, for each act f, we can find a finite partition $\{E_1^f, \ldots, E_n^f\}$ of S and a list of n outcomes $(x_1^f, \ldots, x_n^f)$, such that $u(x_i^f) \geq u(x_{i+1}^f)$, $i = 1, \ldots, n-1$ and $f = (x_1^f)_{E_1^f}(x_2^f)_{E_2^f} \ldots (x_{n-1}^f)_{E_{n-1}^f}(x_n^f)$. Formally, the Choquet Expected Utility of an act f with respect to the utility index u and the capacity ν may be defined as:

$$\int u \circ f \,\mathrm{d}\nu = \nu(E_1^f)u(x_1^f) + \sum_{i=2}^{n} [\nu(\cup_{j=1}^{i} E_j^f) - \nu(\cup_{j=1}^{i-1} E_j^f)]u(x_i^f). \tag{1}$$

In particular, the Choquet Expected Utility of the simple bet $x_E y$, where $x \succsim y$, is given by

$$\int u \circ (x_E y) \,\mathrm{d}\nu = \nu(E)u(x) + \bar{\nu}(E^c)u(y).$$

Definition 2 *(CEU preferences).* The set of conditional preference relations $\langle \succsim_E \rangle_{E \in \mathcal{E}}$ is said to constitute a collection of CEU preferences, if for each $\succsim_E$, there exists a capacity ν_E on $\mathcal{E}$ and a continuous non-constant real-valued function u_E on X such that for all $f, g \in \mathcal{F}$

$$f \succsim_E g \Leftrightarrow \int u_E \circ f \, d\nu_E \geq \int u_E \circ g \, d\nu_E.$$

As is well-known, a CEU preference relation admits a multiple prior representation if (and only if) the capacity is convex, that is, for all pairs of events A and B,

$$\nu(A \cup B) \geq \nu(A) + \nu(B) - \nu(A \cap B).$$

Thus the intersection of the two models is non-empty, but not all multiple prior preferences admit a CEU representation and clearly, since not all capacities are convex, not all CEU preferences admit a multiple prior representation. So neither family is a special case of the other. Furthermore, non-convex capacities can be used to model individuals who have both optimistic as well as pessimistic attitudes towards ambiguity (see Wakker, 2001).

We can now state and prove our main representation result: a family of CEU preferences satisfy the three axioms above if and only if the all the utility indices are the same and the capacity for the conditional preference is obtained from the unconditional capacity by the Full Bayesian Updating (FBU) rule of Jaffray (1992).[3]

Theorem 1. *Fix $\langle \succsim_E \rangle_{E \in \mathcal{E}}$ a collection of CEU preferences. Let (u, ν) be the utility index and capacity associated with $\succsim$. For each $E \in \mathcal{E}$, let ν_E be the capacity associated with $\succsim_E$. The following two statements are equivalent*:

(i) $\langle \succsim_E \rangle_{E \in (\mathcal{E})}$ *satisfies Consequentialism, State Independence and Conditional Certainty Equivalent Consistency.*

(ii) *For each event $E \in \mathcal{E}$,*

$$f \succsim_E g \Leftrightarrow \int u \circ f \, d\nu_E \geq \int u \circ g \, d\nu_E, \text{ where } \nu_E(E^c) = 0.$$

Moreover, if E is non-null, then for every event $A \in \mathcal{E}$

$$\nu_E(A) \equiv \frac{\nu(A \cap E)}{\nu(A \cap E) + \bar{\nu}(A^c \cap E)}.$$

Remark 1. The theorem clearly extends Pires's updating result to the class of CEU preferences. But notice that we have done this in a setting of purely subjective uncertainty. Hence our State Independence assumption is much weaker than hers, as it only entails that the ordinal ranking of outcomes remains unchanged no matter on which event preferences are being conditioned. In the Anscombe-Aumann setting of mixed subjective and objective uncertainty employed by Pires (2002), State Independence says that the certainty equivalent of any pure roulette lottery is the same no matter on which state it obtains and no matter what obtains on other states. This

[3] This is also referred to as the Generalized Updating Rule by Walley (1991).

immediately gives her the invariance of the 'risk preferences over pure roulette lotteries' and hence that the conditional utility functions are the same up to a positive affine transformation. The novel aspect of our proof is that provided not all subsets of the conditioning event are null or conditionally universal we are able to show the 'cardinal invariance' of the state utility functions over outcomes arises from Conditional Certainty Equivalent Consistency. That is, it is implied by the updating rule embodied in this axiom.

5. Conclusion

It is one of the attractive features of expected utility theory that Bayesian Updating provides a natural method for considering information which becomes sequentially available. In contrast, for Choquet Expected Utility (CEU) theories of decision making, there are many "sensible" updating rules in the literature, all of which share the property that the updating rule converges to Bayesian Updating as capacities become additive. Which updating rule one chooses in an application will most likely depend on the intended application as much as on a priori criteria from decision theory and statistics.

In this paper we find that two axioms connecting conditional and unconditional CEU preferences characterize the Full Bayesian Updating rule (FBU). As we show by an example which generalizes the Ellsberg paradox, new information may either generate ambiguity or remove it altogether. Hence, a reasonable consistency condition should allow decision makers to maintain their attitudes towards ambiguity in the face of information that changes the ambiguity of an act. The two axioms, *Conditional Certainty Equivalence Consistency* and *Consequentialism,* yield the FBU rule for capacities and achieve this desideratum.

Acknowledgements

We thank participants at the RUD 2005 meetings and in particular, Paulo Ghirardato, Eran Hanany, Robert Kast, Peter Klibanoff, Mark Machina and Marciano Siniscalchi for their comments and suggestions. This research is in part supported by ESRC grant RES-000-22-0650.

Appendix A. Proof of Theorem 1

Fix an event $E \in \mathcal{E}$. Without loss of generality, normalize $u(M){:=}1$ and $u(0){:=}0$.

Part A. $(i) \Rightarrow (ii)$

Step 1. Consequentialism implies that for CEU representation of $\succsim_E$, we have $\nu_E(E^c) = 0$.

Step 2. We next show the theorem holds for constant acts. Note that for any constant act x, $\int u \circ x \,\mathrm{d}\nu = u(x) = \int u \circ x \,\mathrm{d}\nu_E$. Hence for any pair of constant acts x and y, we have from the CEU representations and State Independence: $u(x) \geq u(y)$ if and only if $x \succsim y$ if and only if $x \succsim_E y$ if and only if $\int u \circ x \,\mathrm{d}\nu_E \geq \int u \circ y \,\mathrm{d}\nu_E$, as required.

Step 3. If E is null, then as Conditional Certainty Equivalent Consistency entails no relationship between $\succsim_E$ and $\succsim$, the result extends trivially to all non-constant acts as well.

For the remainder of the proof of Part A, we shall take E to be non-null.

Step 4. We next show that for every event A, either $\nu(A \cap E) > 0$ or $\bar{\nu}(A^c \cap E) > 0$ (or both). Suppose to the contrary that $\nu(A \cap E) = 0$ and $\bar{\nu}(A^c \cap E) = 0$, (and so, $\nu(A \cup E^c) = 1$). Then $\int u \circ [(M_A 0)_E 0] \,\mathrm{d}\nu = \nu(A \cap E) u(0) = 0$ and $\int u \circ [(M_A 0)_E M] \mathrm{d}\nu = \nu(A \cup E^c) u(M) = 1$. That is, $(M_A 0)_E 0 \sim 0$ and $(M_A 0)_E M \sim M$. But then by Conditional Certainty Equivalent Consistency, $M_A 0 \sim_E M$ and $M_A 0 \sim_E 0$, a contradiction.

Furthermore it follows from the argument in the previous paragraph and Conditional Certainty Equivalent Consistency, that $\nu_E(A) = 0$ whenever $\nu(A \cap E) = 0$ and that $\nu_E(A) = 1$ whenever $\bar{\nu}(A^c \cap E) = 0$.

Step 5. If for every event A, either $\nu(A \cap E) = 0$ or $\bar{\nu}(A^c \cap E) = 0$ (but by Step 4 we know this cannot be true for both), then with respect to the conditional capacity $\nu_E(\cdot)$, every event either has capacity zero (and its complement has capacity one) or it has capacity one (and its complement has capacity zero). Hence, for every act f, $\int u \circ f \, \mathrm{d}\nu_E = u(x^f)$, where

$$x^f \in \mathrm{argmax}_{z \in X}\{u(z) : f^{-1}(z) \neq \emptyset \quad \text{and} : \nu_E(f^{-1}(z : y \succcurlyeq z)) = 1\}.$$

That is, the conditional CEU is the utility of the best outcome in the range of f, for which the capacity of getting something at least as good as that outcome is one, and so (ii) holds as required.

For the rest of the proof of Part A, suppose there exists an event B, such that $\nu(B \cap E) > 0$ and $\bar{\nu}(B^c \cap E) > 0$.

Step 6. We shall show that u_E must be an affine transformation of u. From State Independence it immediately follows that u_E is a monotonic transformation of u. Now, since X is a connected topological space, both u and u_E are continuous and u_E is a monotonic transformation of u, it is sufficient to show that for any x, y, z, if $u_E(x) > u_E(y)$ and $u_E(z) = [u_E(x) + u_E(y)]/2$ then $u(z) = [u(x) + u(y)]/2$.

Fix $x \succ y$. By State Independence, $u_E(x) > u_E(y)$.

Lemma 1. $\nu_E(B) \in (0, 1)$.

Proof. Suppose to the contrary that $\nu_E(B) = 0$, that is, $M_B 0 \sim_E 0$. Since by Consequentialism, $E^c \in \mathcal{N}_E$, $M_{B\cap E} 0 \sim_E 0$. And by Conditional Certainty Equivalent Consistency we also have $M_{B\cap E} 0 \sim 0$. That is, $\nu(B \cap E) = 0$, which contradicts the hypothesis that $\nu(B \cap E) > 0$. So suppose instead that $\nu_E(B) = 1$, that is, $M_B 0 \sim_E M$. By set monotonicity, it follows $M_{B\cup E^c} 0 \sim_E M$ and Conditional Certainty Equivalent Consistency entails $M_{B\cup E^c} 0 \sim M$. That is, $\nu(B \cup E^c) = 1$. But this is a contradiction since we have by hypothesis $\bar{\nu}(B^c \cap E) = 1 - \nu(B \cup E^c) > 0$. □

Since $\nu_E(B) \in (0, 1)$, X is a connected topological space and u_E is continuous, outcomes with the following properties exist.

$$z : u_E(z) = \frac{1}{2} u_E(x) + \frac{1}{2} u_E(y) \tag{2}$$

$$z' : u_E(z') = \nu_E(B) u_E(x) + [1 - \nu_E(B)] u_E(z) \tag{3}$$

$$y' : u_E(y') = \nu_E(B) u_E(z) + [1 - \nu_E(B)] u_E(y) \tag{4}$$

$$z'' : u_E(z'') = \nu_E(B) u_E(x) + [1 - \nu_E(B)] u_E(y) \tag{5}$$

Eqs. (2)–(5) together imply

$$x_{B\cap E} z_{B^c\cap E} z' \sim_E z' \tag{6}$$

$$z_{B\cap E} y_{B^c\cap E} y' \sim_E y' \tag{7}$$

$$x_{B\cap E} y_{B^c\cap E} z'' \sim_E z'' \tag{8}$$

$$x_{B\cap E}y_{B^c\cap E}z''\sim_E z'_{B\cap E}y'_{B^c\cap E}z'' \tag{9}$$

Each of the indifferences (6)–(8) follow directly from Eqs. (3)–(5), respectively. To see that the last indifference also follows, notice that the conditional (on E obtaining) Choquet Expected utility of the act $z'_{B\cap E}y'_{B^c\cap E}z''$ may be expressed

$$\begin{aligned}
&\nu_E(B)u_E(z') + [1-\nu_E(B)]u_E(y')\\
&\quad= \nu_E(B)[\nu_E(B)u_E(x) + [1-\nu_E(B)]u_E(z)]\\
&\qquad+[1-\nu_E(B)][\nu_E(B)u_E(z) + [1-\nu_E(B)]u_E(y)]\\
&\quad= [\nu_E(B)]^2u_E(x) + 2\nu_E(B)[1-\nu_E(B)]\left[\frac{1}{2}u_E(x)+\frac{1}{2}u_E(y)\right] + [1-\nu_E(B)]^2u_E(y)\\
&\quad= \nu_E(B)u_E(x) + [1-\nu_E(B)]u_E(y).
\end{aligned}$$

But $\nu_E(B)u_E(x) + [1-\nu_E(B)]u_E(y)$ is the conditional (on E obtaining) Choquet Expected Utility of the act $x_{B\cap E}y_{B^c\cap E}z''$ and so the indifference (9) holds.

By applying Conditional Certainty Equivalent Consistency to (6)–(9) we obtain

$$x_{B\cap E}z_{B^c\cap E}z' \sim z' \tag{10}$$

$$z_{B\cap E}y_{B^c\cap E}y' \sim y' \tag{11}$$

$$x_{B\cap E}y_{B^c\cap E}z'' \sim z'' \tag{12}$$

$$x_{B\cap E}y_{B^c\cap E}z'' \sim z'_{B\cap E}y'_{B^c\cap E}z''. \tag{13}$$

These four indifference relations imply the following four equations

$$u(z') = \nu(B\cap E)u(x) + [\nu(B\cup E^c) - \nu(B\cap E)]u(z') + [1-\nu(B\cup E^c)]u(z) \tag{14}$$

$$u(y') = \nu(B\cap E)u(z) + [\nu(B\cup E^c) - \nu(B\cap E)]u(y') + [1-\nu(B\cup E^c)]u(y) \tag{15}$$

$$u(z'') = \nu(B\cap E)u(x) + [\nu(B\cup E^c) - \nu(B\cap E)]u(z'') + [1-\nu(B\cup E^c)]u(y) \tag{16}$$

$$u(z'') = \nu(B\cap E)u(z') + [\nu(B\cup E^c) - \nu(B\cap E)]u(z'') + [1-\nu(B\cup E^c)]u(y') \tag{17}$$

Substituting from (14) for $u(z')$ and from (15) for $u(y')$ into (17), and equating (17) with (16) to eliminate $u(z'')$, we obtain

$$\begin{aligned}
&\nu(B\cap E)u(x) + [1-\nu(B\cup E^c)]u(y)\\
&\quad= \nu(B\cap E)\left[\frac{\nu(B\cap E)u(x) + [1-\nu(B\cup E^c)]u(z)}{\nu(B\cap E)+1-\nu(B\cup E^c)}\right]\\
&\qquad+[1-\nu(B\cup E^c)]\left[\frac{\nu(B\cap E)u(z) + [1-\nu(B\cup E^c)]u(y)}{\nu(B\cap E)+1-\nu(B\cup E^c)}\right].
\end{aligned}$$

Collecting terms,

$$\begin{aligned}
&([\nu(B\cap E)+1-\nu(B\cup E^c)]\nu(B\cap E) - [\nu(B\cap E)]^2)u(x) + ([\nu(B\cap E)\\
&\quad+1-\nu(B\cup E^c)][1-\nu(B\cup E^c)] - [1-\nu(B\cup E^c)]^2)u(y)\\
&\qquad= 2\nu(B\cap E)(1-\nu(B\cup E^c))u(z).
\end{aligned}$$

Simplifying yields

$$\nu(B\cap E)\bar{\nu}(B^c\cap E)u(x)+\nu(B\cap E)\bar{\nu}(B^c\cap E)u(y)=2\nu(B\cap E)\bar{\nu}(B^c\cap E)u(z).$$

That is,

$$u(z)=\frac{1}{2}u(x)+\frac{1}{2}u(y).$$

as required.

Step 7. We know from Step 6, that u_E is a positive affine transformation of u. So we can normalize by setting $u_E(0)=u(0)=0$, and $u_E(M)=u(M)=1$. So fix $A\in\mathcal{E}$. By Step 3 we have $\nu(A\cap E)+\bar{\nu}(A^c\cap E)>0$. Define z to be the outcome for which $M_{A\cap E}0_{A^c\cap E}z\sim_E z$. That is, $u(z)=\nu_E(A\cap E)$. Since by Consequentialism $E^c\in\mathcal{N}_E$, we also have $M_A 0\sim_E M_{A\cap E}0_{A^c\cap E}z$, hence $\nu_E(A)=u(z)$. By Conditional Certainty Equivalent Consistency, $M_{A\cap E}0_{A^c\cap E}z\sim z$. There are three cases to consider,

1. $z=0$, that is, $\nu_E(A\cap E)=0$. Thus, $M_{A\cap E}0_{A^c\cap E}z\sim z$ implies $\nu(A\cap E)=0$. And since $\nu(A\cap E)+\bar{\nu}(A^c\cap E)>0$, it must be the case, $\bar{\nu}(A^c\cap E)>0$. Thus we have

$$\nu_E(A)=\frac{\nu(A\cap E)}{\nu(A\cap E)+\bar{\nu}(A^c\cap E)}=\frac{0}{0+\bar{\nu}(A^c\cap E)}=0,\quad \text{as required.}$$

2. $z=M$, that is, $\nu_E(A\cap E)=1$. Thus, $M_{A\cap E}0_{A^c\cap E}z\sim z$ implies $\nu(A\cup E^c)=1$ or equivalently, $\bar{\nu}(A^c\cap E)=1-\nu(A\cup E^c)=0$. And so, $\nu(A\cap E)>0$ and we have

$$\nu_E(A)=\frac{\nu(A\cap E)}{\nu(A\cap E)+\bar{\nu}(A^c\cap E)}=\frac{\nu(A\cap E)}{\nu(A\cap E)+0}=1,\quad \text{as required.}$$

3. $u(z)\in(0,1)$, that is, $\nu_E(A\cap E)\in(0,1)$. Thus, $M_{A\cap E}0_{A^c\cap E}z\sim z$ implies $u(z)=\nu(A\cap E)+[\nu(A\cup E^c)-\nu(A\cap E)]u(z)$, and since $\nu(A\cap E)+\bar{\nu}(A^c\cap E)>0$, then we have

$$u(z)=\frac{\nu(A\cap E)}{\nu(A\cap E)+\bar{\nu}(A^c\cap E)}=\nu_E(A),\quad \text{as required.}$$

Part B. $(ii)\Rightarrow(i)$

Consequentialism follows immediately from the CEU representation of $\succsim_E$ and the fact that $\nu_E(E^c)=0$.

State Independence follows since for every constant act x, $\int u\circ x\,\mathrm{d}\nu=u(x)=\int u\circ x\,\mathrm{d}\nu_E$.

Conditional Certainty Equivalent Consistency follows from the CEU representations of $\succsim$ and $\succsim_E$ (the representations are additive for comonotonic acts) and the definition of conditional capacity.

References

Baron, J., Frisch, D., 1988. Ambiguity and rationality. Journal of Behavioral Decision Making 1, 149–157.

Denneberg, D., 1994. Conditioning (updating) non-additive measures. Annals of Operations Research 52, 21–42.

Denneberg, D., 2002. Conditional expectation for monotone measures, the discrete case. Journal of Mathematical Economics 37, 105–121.

Dow, J., Werlang, S., 1992. Uncertainty aversion, risk aversion, and the optimal choice of portfolio. Econometrica 60, 197–204.

Eichberger, J., Kelsey, D., 2000. Non-additive beliefs and strategic equilibria. Games and Economic Behavior 30, 183–215.

Eichberger, J., Kelsey, D., 2002. Strategic complements, substitutes, and ambiguity: the implications for public goods. Journal of Economic Theory 106, 436–466.

Eichberger, J., Grant, S., Kelsey, D., 2005. CEU preferences and dynamic consistency. Mathematical Social Sciences 49, 143–151.

Epstein, E., Schneider, M., 2003. Recursive multiple-priors. Journal of Economic Theory 113, 1–31.

Gilboa, I., Schmeidler, D., 1989. Maxmin expected utility with a non-unique prior. Journal of Mathematical Economics 18, 141–153.

Gilboa, I., Schmeidler, D., 1993. Updating ambiguous beliefs. Journal of Economic Theory 59, 33–49.

Hanany, E., Klibanoff, P., 2006. Updating Preferences with Multiple Priors. Mimeo Kellog Graduate School of Management, Northwestern University.

Jaffray, J.-Y., 1992. Bayesian updating and belief functions. IEEE Transactions on Systems, Man and Cybernetics 22, 1144–1152.

Keynes, J.M., 1921. A Treatise on Probability. Macmillan, London.

Knight, F.H., 1921. Risk, Uncertainty and Profit. Houghton Mifflin, New York.

Lapied, A., Kast, R., 2005. Updating Choquet Valuation and Discounting Information Arrivals. Mimeo CNRS, Lameta et Idep.

Lehrer, E., 2005. Updating non-additive probabilities—a geometric approach. Games and Economic Behavior 50, 42–57.

Mukerji, S., 1998. Ambiguity aversion and incompleteness of contractual form. American Economic Review 88, 1207–1231.

Mukerji, S., Tallon, J.-M., 2001. Ambiguity aversion and incompleteness of financial markets. Review of Economic Studies 68, 883–904.

Mukerji, S., Tallon, J.-M., 2004. Ambiguity aversion and the absence of wage indexation. Journal of Monetary Economics 51, 653–670.

Pires, C.P., 2002. A rule for updating ambiguous beliefs. Theory and Decision 53, 137–152.

Sarin, R., Wakker, P., 1998. Dynamic choice and non-expected utility. Journal of Risk and Uncertainty 17, 87–119.

Schmeidler, D., 1989. Subjective probability and expected utility without additivity. Econometrica 57, 571–587.

Shafer, G., 1976. A Mathematical Theory of Evidence. Princeton University Press, Princeton New Jersey.

Siniscalchi, M., 2006. Dynamic Choice Under Ambiguity. Mimeo Department of Economics, Northwestern University.

Skiadas, C., 1997. Conditioning and aggregation of preferences. Econometrica 65, 347–367.

Walley, P., 1991. Statistical Reasoning with Imprecise Probabilities. Chapman and Hall, London.

Wakker, P., 2001. Testing and characterizing properties of nonadditive measures through violations of the Sure-Thing Principle. Econometrica 69, 1039–1059.

Wang, T., 2003. Conditional preferences and updating. Journal of Economic Theory 108, 286–321.