THE CARLO ALBERTO NOTEBOOKS

Collegio Carlo Alberto

Revealed Ambiguity and Its Consequences: Updating

Paolo Ghirardato
Fabio Maccheroni
Massimo Marinacci

Working Paper No. 44
May 2007
www.carloalberto.org

Revealed Ambiguity and Its Consequences: Updating[1]

Paolo Ghirardato
Collegio Carlo Alberto
Università di Torino
paolo.ghirardato@unito.it

Fabio Maccheroni
IMQ and IGIER
Università Bocconi
fabio.maccheroni@unibocconi.it

Massimo Marinacci
Collegio Carlo Alberto
Università di Torino
massimo.marinacci@unito.it

March 2007

(to appear in *Advances in Decision Making under Risk and Uncertainty: Selected Papers from the FUR 2006 conference,* (M. Abdellaboui and J. D. Hey, eds.), Springer, 2007.)

[1]Most of the results in this paper were earlier circulated as part of a working paper titled "Ambiguity from the Differential Viewpoint" (Caltech Social Sciences Working Paper Number 1130, April 2002). We are grateful to Marciano Siniscalchi for many conversations on the topic of this paper.

Abstract

We study the updating of beliefs under ambiguity for invariant biseparable preferences. In particular, we show that a natural form of dynamic consistency characterizes the Bayesian updating of these beliefs.

JEL classification: D81

Keywords: ambiguity, updating

Introduction

Dynamic consistency is a fundamental property in dynamic choice models. It requires that if a decision maker plans to take some action at some juncture in the future, he should consistently take that action when finding himself at that juncture, and vice versa if he plans to take a certain action at a certain juncture, he should take that plan in mind when deciding what to do now.

However compelling *prima facie*, it is well known in the literature that there are instances in which the presence of ambiguity might lead to behavior that reasonably violates dynamic consistency, as the next Ellsberg example shows.[1]

Example 1 *Consider the classical "3-color" Ellsberg problem, in which an urn contains 90 balls, 30 of which are known to be red, while the remaining 60 are either blue or green. In period 0, the decision maker only has the information described above. Suppose that at the beginning of period 1 a ball is extracted from the urn, and the decision maker is then told whether the ball is blue or not. The decision maker has to choose between bets on the color of the drawn ball. Denoting by* $[a, b, c]$ *an act that pays* a *when a red ball is extracted,* b *when a green ball is extracted and* c *otherwise, let*

$$\begin{aligned} f &= [1,0,0] \\ g &= [0,1,0] \\ f' &= [1,0,1] \\ g' &= [0,1,1] \end{aligned}$$

Suppose that in period 0, the decision maker, like most people in Ellsberg's experiment, displays the following preference pattern

$$g' \succ f' \succ f \succ g \tag{1}$$

(the middle preference being due to monotonicity). Letting $A = \{R, G\}$*, it follows immediately from consequentialism that, conditionally on* A^c*,*

$$f' \sim_{A^c} g'.$$

On the other hand, if the decision maker's conditional preferences satisfy dynamic consistency it must be the case that if he finds an act to be optimal conditionally on A *and also conditionally on* A^c *in period 1, he must find the same act optimal in period 0. So, dynamic*

[1] We owe this example to Denis Bouyssou, who showed it to us at the RUD 1997 conference in Chantilly.

consistency implies that $g' \succ_A f'$ *(as otherwise we should have* $f' \succcurlyeq g'$*). That is, a dynamically consistent and consequentialist decision maker who is told that a blue ball has not been extracted from the Ellsberg urn (i.e., is told A) must strictly prefer to bet on a green ball having been extracted.*

Yet, it seems to us that a decision maker with the ambiguity averse preferences in Eq. (1) might *still prefer to bet on a red ball being extracted, finding that event less ambiguous than the extraction of a green ball, and that constraining him to choose otherwise is imposing a strong constraint on the dynamics of his ambiguity attitude.*

In view of this example, we claim that dynamic consistency is a compelling property only for comparisons of acts that are not affected by the possible presence of ambiguity. In other words, we think that rankings of acts unaffected by ambiguity should be dynamically consistent.

This is the starting point of this paper. We consider the preferences represented by

$$V(f) = a(f) \min_{P \in C} \int u(f(s))\, dP + (1 - a(f)) \max_{P \in C} \int u(f(s))\, dP, \tag{2}$$

where f is an act, a is a function over acts that describes the decision maker's attitudes toward ambiguity, and C is a set of priors that represents the ambiguity revealed by the decision maker's behavior. We provided an axiomatic foundation for such preferences in Ghirardato et al. (2004, henceforth GMM).[2] There, we also introduced a notion of unambiguous preference which is derived from the observable preference over acts. We argued that such derived unambiguous preference only ranks pairs of acts whose comparison is not affected by ambiguity. That is, unambiguous preference is a partial ordering, which is represented *à la* Bewley (2002) by the set of priors C (see Eq. (3) below).

Our main intuition is then naturally modelled by assuming that the derived unambiguous preference is dynamically consistent, while, in the presence of ambiguity, the primitive preference might well not be. This natural modelling idea leads to a simple and clean characterization of updating for the preferences we discuss in GMM. The main result of the present paper, Theorem 11, shows that the unambiguous preference is dynamically consistent if and only if all priors in C are updated according to Bayes' rule. This result thus characterizes prior by prior Bayesian updating, a natural updating rule for the preferences represented by Eq. (2).

[2] The multiple priors model of Gilboa and Schmeidler (1989) corresponds to the special case in which $a(f) = 1$ for all acts f. The Choquet expected utility model of Schmeidler (1989) is also seen to be a special case.

We also consider a stronger dynamic consistency restriction on preferences, which can be loosely described as imposing dynamic consistency of the decision maker's "pessimistic self." We show that such restriction (unlike the one considered earlier) leads to imposing some structure on the decision maker's ex-ante perception of ambiguity, which corresponds to the property that Epstein and Schneider (2003) have called *rectangularity*. This shows, *inter alia*, that rectangularity is not in general (i.e., for the preferences axiomatized in GMM) the characterization of dynamic consistency of the primitive preference relation, but of a different dynamic property which might even be logically unrelated to it.

We close by observing that we retain consequentialism of the primitive preference, another classic dynamic property that requires that preferences conditional on some event A only depend on the consequences inside A. This property has been weakened in Hanany and Klibanoff (2004), which also offers a survey of the literature on dynamic choice under ambiguity.

1 Preliminaries

1.1 Notation

Consider a set S of **states of the world**, an algebra Σ of subsets of S called **events**, and a set X of **consequences**. We denote by $\mathfrak{F}$ the set of all the **simple acts**: finite-valued Σ-measurable functions $f : S \to X$. Given any $x \in X$, we abuse notation by denoting $x \in \mathfrak{F}$ the constant act such that $x(s) = x$ for all $s \in S$, thus identifying X with the subset of the constant acts in $\mathfrak{F}$. Given $f, g \in \mathfrak{F}$ and $A \in \Sigma$, we denote by $f\,A\,g$ the act in $\mathfrak{F}$ which yields $f(s)$ for $s \in A$ and $g(s)$ for $s \in A^c \equiv S \setminus A$. We model the DM's preferences on $\mathfrak{F}$ by a binary relation $\succcurlyeq$. As usual, $\succ$ and $\sim$ denote respectively the asymmetric and symmetric parts of $\succcurlyeq$.

We let $B_0(\Sigma)$ denote the set of all real-valued Σ-measurable simple functions, or equivalently the vector space generated by the indicator functions 1_A of the events $A \in \Sigma$. If $f \in \mathfrak{F}$ and $u : X \to \mathbb{R}$, we denote by $u(f)$ the element of $B_0(\Sigma)$ defined by $u(f)\,(s) = u(f(s))$ for all $s \in S$. A *probability charge* on (S, Σ) is function $P : \Sigma \to [0, 1]$ that is normalized and (finitely) additive; i.e., $P(A \cup B) = P(A) + P(B)$ for any disjoint $A, B \in \Sigma$. Abusing our notation we sometimes use $P(\varphi)$ in place of $\int \varphi\, dP$, where $\varphi \in B_0(\Sigma)$.

Given a functional $I : B_0(\Sigma) \to \mathbb{R}$, we say that I is: **monotonic** if $I(\varphi) \geq I(\psi)$ for all $\varphi, \psi \in B_0(\Sigma)$ such that $\varphi(s) \geq \psi(s)$ for all $s \in S$; **constant additive** if $I(\varphi + \alpha) = I(\varphi) + \alpha$ for all $\varphi \in B_0(\Sigma)$ and $\alpha \in \mathbb{R}$; **positively homogeneous** if $I(\alpha\varphi) = \alpha I(\varphi)$ for all $\varphi \in B_0(\Sigma)$ and $\alpha \geq 0$; **constant linear** if it is constant additive and positively homogeneous.

Finally, as customary, given $f \in \mathfrak{F}$, we denote by $\Sigma(f)$ the algebra generated by f.

1.2 Invariant Biseparable Preferences

We next present the preference model used in the paper. We recall first the MEU model of Gilboa and Schmeidler (1989). In this model, a decision maker is represented by a utility function u and a set of probability charges $\mathcal{C}$, and she chooses according to the rule $\min_{P\in\mathcal{C}} \int u(\cdot)\, dP$. A generalization of this model is the so-called α**-maxmin** (α-MEU) model, in which the decision maker evaluates act $f \in \mathfrak{F}$ according to

$$\alpha \min_{P\in\mathcal{C}} \int_S u(f(s))\, dP(s) + (1-\alpha) \max_{P\in\mathcal{C}} \int_S u(f(s))\, dP(s).$$

The α-MEU model is also a generalization —to an arbitrary set of priors, rather than the set of all possible priors on Σ— of Hurwicz's α*-pessimism* decision rule, which recommends evaluating an act by taking a convex combination (with weight α) of the utility of its worst possible result and of the utility of its best possible result. In collaboration with Arrow (1972), Hurwicz later studied a generalization of his rule, which allows the "pessimism" weight α to vary according to the identity of the worst and best results that the act may yield.

As it turns out, there is a similar generalization of the α-MEU model allowing the weight α to depend on some features of the act f being evaluated. It is the model studied by GMM (see also Nehring (2001) and Ghirardato et al. (2003)), which relaxes Gilboa and Schmeidler's axiomatization of MEU by not imposing their "ambiguity aversion" axiom (and is constructed in a fully subjective setting). We present its functional characterization below, referring the reader to the cited (2004, 2003) for the axiomatic foundation and further discussion. (The axioms are simply those of Gilboa and Schmeidler (1989) minus their "uncertainty aversion" axiom.)

Definition 2 *A binary relation $\succcurlyeq$ on $\mathfrak{F}$ is called an **invariant biseparable preference** if there exist a unique monotonic and constant linear functional $I : B_0(\Sigma) \to \mathbb{R}$ and a nonconstant convex-ranged utility $u : X \to \mathbb{R}$, unique up to a positive affine transformation, such that $I(u(\cdot))$ represents $\succcurlyeq$; that is, for every $f, g \in \mathfrak{F}$,*

$$f \succcurlyeq g \Leftrightarrow I(u(f)) \geq I(u(g)).$$

It is easy to see (see GMM, page 157) that a functional $I : B_0(\Sigma) \to \mathbb{R}$ that satisfies monotonicity and constant linearity is also *Lipschitz continuous* of rank 1; i.e., $|I(\varphi) - I(\psi)| \leq \|\varphi - \psi\|$ for any $\varphi, \psi \in B_0(\Sigma)$.

In order to show how this model relates to the α-MEU model, we need to show how to derive a set of priors and consequently the decision maker's ambiguity attitude.

Suppose that act f is preferred to act g. If there is ambiguity about the state space, it is possible that such preference may not hold when we consider acts which average the payoffs of f and g with those of a common act h. Precisely, it is possible that a "mixed" act $g\,\lambda\,h$, which in each state s provides the average utility

$$u(g\,\lambda\,h)(s) = \lambda u(g(s)) + (1-\lambda)u(h(s)),$$

be preferred to a "mixed" act $f\,\lambda\,h$, which offers an analogous average of the payoffs of f and h. Such would be the case, for instance, if $g\,\lambda\,h$ has a utility profile which is almost independent of the realized state —while $f\,\lambda\,h$ does not— and the decision maker is pessimistic. On the other hand, there might be pairs of acts for which these "utility smoothing effects" are second-order. In such a case, we have "unambiguous preference." Precisely,

Definition 3 *Let $f, g \in \mathfrak{F}$. Then, f is **unambiguously preferred** to g, denoted $f \succcurlyeq^* g$, if*

$$f\,\lambda\,h \succcurlyeq g\,\lambda\,h$$

for all $\lambda \in (0,1]$ and all $h \in \mathcal{F}$.

Notice that in general $\succcurlyeq^*$ is a coarsening (hence incomplete) of $\succcurlyeq$, while on the other hand for any $x, y \in X$, $x \succcurlyeq^* y$ if and only if $x \succcurlyeq y$.

In GMM we show that given an invariant biseparable preference there exists a unique nonempty, convex and (weak*) closed set $\mathcal{C}$ of probability charges that represents the unambiguous preference relation $\succcurlyeq^*$ in the following sense

$$f \succcurlyeq^* g \Longleftrightarrow \int_S u(f(s))\,dP(s) \geq \int_S u(g(s))\,dP(s) \qquad \text{for all } P \in \mathcal{C}. \tag{3}$$

That is, unambiguous preference corresponds to preference according to every one of the possible "probabilistic scenarios" included in $\mathcal{C}$. The set $\mathcal{C}$ therefore represents the ambiguity that is revealed by the decision maker's behavior.

Given the representation $\mathcal{C}$, the decision maker's index of ambiguity aversion a is then extracted from the functional I in the following natural way:

$$I(u(f)) = a(f)\min_{P\in\mathcal{C}}\int_S u(f(s))\,dP(s) + (1-a(f))\max_{P\in\mathcal{C}}\int_S u(f(s))\,dP(s).$$

The coefficient $a : \mathfrak{F} \to [0,1]$ is uniquely identified (GMM, Thm. 11) on the set of acts whose expectation is nonconstant over $\mathcal{C}$; i.e., those f for which it is *not* the case that

$$\int_S u(f(s))\,dP(s) = \int_S u(f(s))\,dQ(s) \text{ for every } P, Q \in \mathcal{C}. \tag{4}$$

Moreover, wherever uniquely defined, a also displays a significant regularity, as it turns out that $a(f) = a(g)$ whenever f and g "order" identically the possible scenarios in $\mathcal{C}$. Formally, for all $P, Q \in \mathcal{C}$,

$$\int_S u(f(s))\,dP(s) \geq \int_S u(f(s))\,dQ(s) \Longleftrightarrow \int_S u(g(s))\,dP(s) \geq \int_S u(g(s))\,dQ(s). \tag{5}$$

(See GMM, Prop. 10 and Lemma 8 respectively for behavioral equivalents of the above conditions.) In words, the decision maker's degree of pessimism, though possibly variable, will not vary across acts which are symmetrically affected by ambiguity. Notice that in our environment the Arrow-Hurwicz rule corresponds to the case in which a decision maker's degree of pessimism only depends on the probabilities that maximize and minimize an act's evaluation. Thus, letting the degree of pessimism depend on all the ordering on $\mathcal{C}$ is a generalization of the Arrow-Hurwicz rule. Clearly, the SEU model corresponds to the special case in which $\mathcal{C}$ is a singleton. Thus, all SEU preferences whose utility is convex-ranged are invariant biseparable preferences. Less obviously, also CEU preferences with convex-ranged utility are invariant biseparable preferences. Hence, this model includes both α-MEU and CEU as special cases.

Unless otherwise noted, for the remainder of the paper preferences are always (but often tacitly) assumed to be invariant biseparable in the sense just described.

2 Some Derived Concepts

We now introduce three notions which can be derived from the primitive preference relation via the unambiguous preference relation. Besides being intrinsically interesting, such notions prove useful in presenting the main ideas of the paper.

2.1 Mixture Certainty Equivalents

For any act $f \in \mathfrak{F}$, denote by $C^*(f)$ the set of the consequences that are "indifferent" to f in the following sense:

$$C^*(f) \equiv \{x \in X : \text{ for all } y \in X,\ y \succcurlyeq^* f \text{ implies } y \succcurlyeq^* x,\ f \succcurlyeq^* y \text{ implies } x \succcurlyeq^* y\}.$$

Intuitively, these are the constants that correspond to possible certainty equivalents of f. The set $C^*(f)$ can be characterized (GMM, Prop. 18) in terms of the set of expected utilities associated with $\mathcal{C}$:

Proposition 4 *For every $f \in \mathfrak{F}$,*

$$x \in C^*(f) \Longleftrightarrow \min_{P \in \mathcal{C}} P(u(f)) \leq u(x) \leq \max_{P \in \mathcal{C}} P(u(f)).$$

Moreover, $u(C^*(f)) = [\min_{P \in \mathcal{C}} P(u(f)), \max_{P \in \mathcal{C}} P(u(f))]$.

It follows immediately from the proposition that $x \in C^*(f)$ if and only if there is a $P \in \mathcal{C}$ such that $u(x) = P(u(f))$. That is, $u(C^*(f))$ is the range of the mapping that associates each prior $P \in \mathcal{C}$ with the expected utility $P(u(f))$.

There is another sense in which the elements of $C^*(f)$ are generalized certainty equivalents of f. Consider a consequence $x \in X$ that can be substituted to f as a "payoff" in a given mixture. That is, such that for some $\lambda \in (0, 1]$ and $h \in \mathfrak{F}$,

$$x \, \lambda \, h \sim f \, \lambda \, h.$$

The following result shows that, while not all the elements of the set $C^*(f)$ can in general be expressed in this fashion, each of them is infinitesimally close (in terms of preference) to a consequence with this property.[3]

Proposition 5 *For every $f \in \mathfrak{F}$, $C^*(f)$ is the preference closure of the set*

$$\{x \in X : \exists \lambda \in (0, 1], \exists h \in \mathfrak{F} \text{ such that } x \, \lambda \, h \sim f \, \lambda \, h\}.$$

In light of this result, we abuse terminology somewhat and call $x \in C^*(f)$ a **mixture certainty equivalent** of f, and $C^*(f)$ the **mixture certainty equivalents set** of f.

2.2 Lower and Upper Envelope Preferences

Given the unambiguous preference $\succcurlyeq^*$ induced by $\succcurlyeq$, we can also define the following two relations:

Definition 6 *The **lower envelope preference** is the binary relation $\succcurlyeq^{\downarrow}$ on $\mathfrak{F}$ defined as follows: for all $f, g \in \mathfrak{F}$,*

$$f \succcurlyeq^{\downarrow} g \Longleftrightarrow \{x \in X : f \succcurlyeq^* x\} \supseteq \{x \in X : g \succcurlyeq^* x\}.$$

*The **upper envelope preference** is the binary relation $\succcurlyeq^{\uparrow}$ on $\mathfrak{F}$ defined as follows: for all $f, g \in \mathfrak{F}$,*

$$f \succcurlyeq^{\uparrow} g \Longleftrightarrow \{x \in X : x \succcurlyeq^* f\} \subseteq \{x \in X : x \succcurlyeq^* g\}.$$

[3] In the statement by "preference closure" of a subset $Y \subseteq X$, we mean $u^{-1}\left(\overline{u(Y)}\right)$.

The relation $\succcurlyeq^{\downarrow}$ describes a "pessimistic" evaluation rule, while $\succcurlyeq^{\uparrow}$ an "optimistic" evaluation rule. To see this, notice that $\succcurlyeq^{\downarrow}$ ranks acts by the size of the set of consequences that are unambiguously worse than f. In fact, it ranks f exactly as the most valuable consequence that is unambiguously worse than f. The twin relation $\succcurlyeq^{\uparrow}$ does the opposite. We denote by $\succ^{\downarrow}$ and $\sim^{\downarrow}$ (resp. $\succ^{\uparrow}$ and $\sim^{\uparrow}$) the asymmetric and symmetric components of $\succcurlyeq^{\downarrow}$ (resp. $\succcurlyeq^{\uparrow}$) respectively.

This is further clarified by the following result, which shows that the envelope relations can be represented in terms of the set $\mathcal{C}$ derived in the previous section.

Proposition 7 *For every $f, g \in \mathfrak{F}$, the following statements are equivalent:*

(i) $f \succcurlyeq^{\downarrow} g$ *(resp.* $f \succcurlyeq^{\uparrow} g$*).*

(ii) $\min_{P\in\mathcal{C}} P(u(f)) \geq \min_{P\in\mathcal{C}} P(u(g))$ *(resp.* $\max_{P\in\mathcal{C}} P(u(f)) \geq \max_{P\in\mathcal{C}} P(u(g))$*).*

It follows from this result that $\succcurlyeq^{\downarrow}$ is a 1-MEU preference, in particular an invariant biseparable preference, and that $(\succcurlyeq^{\downarrow})^*$ is represented by $\mathcal{C}$. Moreover, while for every $x, y \in X$, $x \succcurlyeq y$ if and only if $x \succcurlyeq^{\downarrow} y$, $\succcurlyeq = \succcurlyeq^{\downarrow}$ holds if and only if $\succcurlyeq$ is 1-MEU, so that $\succcurlyeq$ and $\succcurlyeq^{\downarrow}$ will be in general distinct. Symmetric observations hold for $\succcurlyeq^{\uparrow}$.

The relations between $\succcurlyeq^{\downarrow}$, $\succcurlyeq^{\uparrow}$ and $\succcurlyeq$ can be better understood by recalling the relative ambiguity aversion ranking of Ghirardato and Marinacci (2002).

Proposition 8 *The preference relation $\succcurlyeq^{\downarrow}$ is more ambiguity averse than $\succcurlyeq$, which is in turn more ambiguity averse than $\succcurlyeq^{\uparrow}$.*

Therefore, the envelope relations can be interpreted as the "ambiguity averse side" and the "ambiguity loving side" of the DM. Indeed, $\succcurlyeq^{\downarrow}$ is ambiguity averse in the absolute sense of Ghirardato and Marinacci (2002), while $\succcurlyeq^{\uparrow}$ is ambiguity loving.

3 Revealed Ambiguity and Updating

Suppose that our DM has an information structure given by some subclass Π of Σ (say, a partition or a sub-algebra), and assume that we can observe our DM's *ex ante* preference on $\mathfrak{F}$, denoted interchangeably $\succcurlyeq$ or $\succcurlyeq_S$, and his preference on $\mathfrak{F}$ after having been informed that an event $A \in \Pi$ obtained, denoted $\succcurlyeq_A$. For each $A \in \Pi' \equiv \Pi \cup S$, the preference $\succcurlyeq_A$ is assumed to be invariant biseparable, and the utility representing $\succcurlyeq_A$ is denoted by u_A. Clearly, a conditional preference $\succcurlyeq_A$ also induces an unambiguous preference relation $\succcurlyeq_A^*$, as well as mixture certainty equivalents sets $C_A^*(\cdot)$ and a lower envelope preference relation $\succcurlyeq_A^{\downarrow}$.

Because $\succcurlyeq_A$ is invariant biseparable, it is possible to represent $\succcurlyeq_A^*$ in the sense of Eq. (3) by a nonempty, weak* compact and convex set of probability measures $\mathcal{C}_A$.

We are interested in preferences conditional on events which are (*ex ante*) unambiguously non-null in the following sense:

Definition 9 *We say that $A \in \Sigma$ is* ***unambiguously non-null*** *if $x\,A\,y \succ^{\downarrow} y$ for some (all) $x \succ y$.*

That is, an event is unambiguously non-null if betting on A is unambiguosly better than getting the loss payoff y for sure (notice that this is stronger than the definition of non-null event in Ghirardato and Marinacci (2001), which just requires that $x\,A\,y \succ y$). This property is equivalently restated in terms of the possible scenarios $\mathcal{C}$ as follows: $P(A) > 0$ for all $P \in \mathcal{C}$.

We next assume that conditional on being informed of A, the DM only cares about an act's results on A, a natural assumption that we call **consequentialism**: For every $A \in \Pi$, $f \sim_A f\,A\,g$ for every $f, g \in \mathfrak{F}$. Consequentialism extends immediately to the unambiguous and lower envelope preference relations, as the following result shows:

Lemma 10 *For every $A \in \Pi$, the following statements are equivalent:*[4]

(*i*) $f \sim_A f\,A\,g$ *for every* $f, g \in \mathfrak{F}$.

(*ii*) $f \sim_A^* f\,A\,g$ *for every* $f, g \in \mathfrak{F}$.

(*iii*) $f \sim_A^{\downarrow} f\,A\,g$ *for every* $f, g \in \mathfrak{F}$.

For the remainder of this section we tacitly assume that all the preferences $\succcurlyeq_A$ are invariant biseparable and consequentialist.

An important property linking *ex ante* and *ex post* preferences is **dynamic consistency**: For all $A \in \Pi$ and all $f, g \in \mathfrak{F}$,

$$f \succcurlyeq_A g \Longleftrightarrow f\,A\,g \succcurlyeq g. \tag{6}$$

This property imposes two requirements. The first says that the DM should consistently carry out plans made *ex ante*. The second says that information is valuable to the DM, in the sense that postponing her choice to after knowing whether an event obtained does not make her worse off (see Ghirardato (2002) for a more detailed discussion).

As announced in the Introduction, we now inquire the effect of requiring dynamic consistency only in the absence of ambiguity; i.e., requiring Eq. (6) with $\succcurlyeq$ and $\succcurlyeq_A$ replaced by the

[4] In this and the remaining results of this section, we omit the equivalent statements involving the upper envelope preference relation.

unambiguous preference relations $\succsim^*$ and $\succsim^*_A$ respectively. We show that (for a preference satisfying consequentialism) this is tantamount to assuming that the DM updates all the priors in $\mathcal{C}$, a procedure that we call **generalized Bayesian updating**: For every $A \in \Pi$, the "updated" perception of ambiguity is equal to

$$\mathcal{C}|A \equiv \overline{co}^{w^*}\{P_A : P \in \mathcal{C}\},$$

where P_A denotes the posterior of P conditional on A, and $\overline{co}^{w^*}$ stands for the weak* closure of the convex hull.

Theorem 11 *Suppose that $A \in \Pi$ is unambiguously non-null. Then the following statements are equivalent:*

(i) *For every $f, g \in \mathfrak{F}$,*

$$f \succsim^*_A g \Longleftrightarrow P_A(u(f)) \geq P_A(u(g)) \text{ for all } P \in \mathcal{C}. \tag{7}$$

Equivalently, $\mathcal{C}_A = \mathcal{C}|A$ and $u_A = u$.

(ii) *The relation $\succsim^*$ is dynamically consistent with respect to A. That is, for every $f, g \in \mathfrak{F}$:*

$$f \succsim^*_A g \Longleftrightarrow f\,A\,g \succsim^* g. \tag{8}$$

(iii) *For every $x, x' \in X$, $x \succsim x' \Rightarrow x \succsim_A x'$. For every $f \in \mathfrak{F}$ and $x \in X$:*

$$x \in C^*_A(f) \Longleftrightarrow x \in C^*(f\,A\,x). \tag{9}$$

(iv) *For every $f \in \mathfrak{F}$ and $x \in X$:*

$$f \succsim^{\downarrow}_A x \Longleftrightarrow f\,A\,x \succsim^{\downarrow} x. \tag{10}$$

Alongside the promised equivalence with dynamic consistency of unambiguous preference, this results presents two other characterizations of generalized Bayesian updating. They are inspired by a result of Pires (2002), who shows that when the primitive preference relations $\succsim_A$ are 1-MEU, generalized Bayesian updating is characterized by (a condition equivalent to)

$$f \succsim_A (\sim_A)\, x \Longleftrightarrow f\,A\,x \succsim (\sim)\, x \tag{11}$$

for all $f \in \mathfrak{F}$ and $x \in X$. Statement (iii) in the proposition departs from the indifference part of Eq. (11) and applies its logic to the "indifference" notion that is generated by the incomplete preference $\succcurlyeq^*$. Statement (iv) is a direct generalization of Pires's result to preferences that are not 1-MEU. Notice that Eq. (10) is equivalent to requiring that $f \succcurlyeq_A^* x$ if and only if $fAx \succcurlyeq^* x$, a weakening of Eq. (8) that under the assumptions of the proposition is equivalent to it.

It is straightforward to show that dynamic consistency of the primitives $\{\succcurlyeq_A\}_{A\in\Pi'}$ implies condition (ii). Thus, dynamic consistency of the primitives is a sufficient condition for generalized Bayesian updating. The following example reprises the Ellsberg discussion in the Introduction to show that it is not necessary.

Example 12 *Consider the (CEU and) 1-MEU preference described by (linear utility and) the set $\mathcal{C} = \{P : P(R) = 1/3,\ P(G) \in [1/6, 1/2]\}$. It is clear that a decision maker with such $\mathcal{C}$ would display the preference pattern of Eq. (1). It follows from Theorem 11 that her preferences will satisfy consequentialism and unambiguous dynamic consistency if and only if conditionally on $A = \{R, G\}$ her updated set of priors is*

$$\mathcal{C}_A = \{P : P(R) \in [2/5, 2/3]\}.$$

Assuming that the decision maker is also 1-MEU conditionally on A, this implies that in period 1 she will still prefer betting on a red ball over betting on a green ball. As discussed in the Introduction, this cannot happen if the decision maker's conditional preferences satisfy dynamic consistency tout court*; i.e., Eq. (6).*

A different way of reinforcing the conditions of Theorem 11 is to consider imposing the full strength of dynamic consistency on the lower envelope preference relations, rather than the weaker form seen in Eq. (10). We next show that this leads to the characterization of the notion of rectangularity introduced by Epstein and Schneider (2003).

Suppose that the class Π forms a finite partition of S; i.e., $\Pi = \{A_1, \ldots, A_n\}$, with $A_i \cap A_j = \emptyset$ for every $i \neq j$ and $S = \cup_{i=1}^n A_i$. Given a set of probabilities $\mathcal{C}$ such that each A_i is unambiguously nonnull, we define

$$[\mathcal{C}] = \left\{ P : \exists Q, P_1, \ldots, P_n \in \mathcal{C} \text{ such that } \forall B \in \Sigma,\ P(B) = \sum_{i=1}^n P_i(B|A_i)\, Q(A_i) \right\}.$$

We say that $\mathcal{C}$ is Π**-rectangular** if $\mathcal{C} = [\mathcal{C}]$.[5] (We refer the reader to Epstein and Schneider (2003) for more discussion of this concept.)

Proposition 13 *Suppose that Π is a partition of S and that every $A \in \Pi$ is unambiguously non-null. Then the following statements are equivalent:*

[5]We owe this presentation of rectangularity to Marciano Siniscalchi.

(i) $\mathcal{C}$ *is* Π*-rectangular, and for every* $A \in \Pi$, $u_A = u$ *and* $\mathcal{C}_A = \mathcal{C}|A$.

(ii) *For every* $f, g \in \mathfrak{F}$ *and* $A \in \Pi$*:*

$$f \succcurlyeq^{\downarrow}_{A} g \Longleftrightarrow f\, A\, g \succcurlyeq^{\downarrow} g.$$

The rationale for this result is straightforward: Since the preference $\succcurlyeq^{\downarrow}$ is 1-MEU with set of priors $\mathcal{C}$, it follows from the analysis of Epstein and Schneider (2003) that $\mathcal{C}$ is rectangular and that for every $A \in \Pi$, $\mathcal{C}_A$ is obtained by generalized Bayesian updating. But the sets $\mathcal{C}_A$ are also those that represent the ambiguity perception of the primitive relations $\succcurlyeq_A$, as they represent the ambiguity perception of $\succcurlyeq^{\downarrow}_{A}$.

We have therefore shown that the characterization of rectangularity and generalized Bayesian updating of Epstein and Schneider can be extended to preferences which do not satisfy ambiguity hedging, having taken care to require dynamic consistency of the lower envelope (or equivalently of the upper envelope), rather than of the primitive, preference relations. The relations between dynamic consistency of the primitives $\{\succcurlyeq_A\}_{A\in\Pi'}$ and of the lower envelopes $\{\succcurlyeq^{\downarrow}_{A}\}_{A\in\Pi'}$ are not obvious and are the object of ongoing research.

A Appendix

We begin with a preliminary remark and two pieces of notation, that are used throughout this appendix. First, notice that since $u(X)$ is convex, it is w.l.o.g. to assume that $u(X) \supseteq [-1, 1]$. Second, denote by $B_0(\Sigma, u(X))$ the set of the functions in $B_0(\Sigma)$ that map into $u(X)$. Finally, given a nonempty, convex and weak* compact set $\mathcal{C}$ of probability charges on (S, Σ), we denote for every $\varphi \in B_0(\Sigma)$,

$$\underline{\mathcal{C}}(\varphi) = \min_{P \in \mathcal{C}} P(\varphi), \quad \overline{\mathcal{C}}(\varphi) = \max_{P \in \mathcal{C}} P(\varphi).$$

A.1 Proof of Proposition 5

Since the map from $B_0(\Sigma)$ to $\mathbb{R}$ defined by

$$\psi \mapsto I(u(f) + \psi) - I(\psi)$$

is continuous and $B_0(\Sigma)$ is connected, the set

$$\begin{aligned} J &= \{I(u(f) + \psi) - I(\psi) : \psi \in B_0(\Sigma)\} \\ &= \left\{ I\left(u(f) + \frac{1-\lambda}{\lambda} u(g)\right) - I\left(\frac{1-\lambda}{\lambda} u(g)\right) : g \in \mathfrak{F},\ \lambda \in (0, 1] \right\} \end{aligned}$$

is connected. That is, it is an interval. From Lemma B.4 in GMM, it follows that

$$\overline{J} = [\underline{\mathcal{C}}(u(f)), \overline{\mathcal{C}}(u(f))].$$

Let

$$M(f) = \{x \in X : \exists \lambda \in (0, 1], \exists h \in \mathfrak{F} \text{ such that } x\, \lambda\, h \sim f\, \lambda\, h\}.$$

We have $x \in M(f)$ iff

$$u(x) = I\left(u(f) + \frac{1-\lambda}{\lambda} u(h)\right) - I\left(\frac{1-\lambda}{\lambda} u(h)\right)$$

iff $x \in u^{-1}(J)$. Hence, $u(M(f)) \subseteq J$. Conversely, if $t \in J$, $t \in [\underline{\mathcal{C}}(u(f)), \overline{\mathcal{C}}(u(f))]$ and there exists $x \in X$ such that $u(x) = t$. Clearly, $x \in M(f)$, whence $u(M(f)) = J$. We conclude observing that

$$C^*(f) = u^{-1}\left([\min_{P \in \mathcal{C}} P(u(f)), \max_{P \in \mathcal{C}} P(u(f))]\right) = u^{-1}\left(\overline{J}\right) = u^{-1}\left(\overline{u(M(f))}\right).$$

A.2 Proof of Proposition 7

To prove the statement for $\succcurlyeq^{\downarrow}$ (that for $\succcurlyeq^{\uparrow}$ is proved analogously), we only need to show that

$$f \succcurlyeq^{\downarrow} g \Longleftrightarrow \underline{\mathcal{C}}(u(f)) \geq \underline{\mathcal{C}}(u(g)).$$

Applying the definition of $\succcurlyeq^{\downarrow}$ and the representation of Eq. (3), we have that $f \succcurlyeq^{\downarrow} g$ iff for every $x \in X$,

$$P(u(g)) \geq u(x) \text{ for all } P \in \mathcal{C} \Rightarrow P(u(f)) \geq u(x) \text{ for all } P \in \mathcal{C}.$$

That is, iff for every $x \in X$,

$$\underline{\mathcal{C}}(u(g)) \geq u(x) \Rightarrow \underline{\mathcal{C}}(u(f)) \geq u(x).$$

This is equivalent to $\underline{\mathcal{C}}(u(f)) \geq \underline{\mathcal{C}}(u(g))$, concluding the proof.

A.3 Proof of Proposition 8

We have proved in Proposition 7 that $\succcurlyeq^{\downarrow}$ is represented by the functional $\underline{\mathcal{C}}(u(\cdot))$, and $\succcurlyeq^{\uparrow}$ by $\overline{\mathcal{C}}(u(\cdot))$. Consider $\succcurlyeq$ and $\succcurlyeq^{\downarrow}$, and notice that (GMM, Prop. 7) for any $f \in \mathfrak{F}$, $\underline{\mathcal{C}}(u(f)) \leq I(u(f)) \leq \overline{\mathcal{C}}(u(f))$. It is clear that $\underline{\mathcal{C}}(u(f)) \leq I(u(f))$ is tantamount to saying that for every $x \in X$,

$$x \succcurlyeq f \Rightarrow x \succcurlyeq^{\downarrow} f.$$

The argument for $\succcurlyeq$ and $\succcurlyeq^{\uparrow}$ is analogous.

A.4 Proof of Lemma 10

$(i) \Leftrightarrow (ii)$: Assume that, for every $A \in \Pi$, $f \sim_A f\,A\,g$ for every $f, g \in \mathfrak{F}$, hence, for every $h \in \mathfrak{F}$ $[\lambda f + (1-\lambda)h] \sim_A [\lambda f + (1-\lambda)h]A[\lambda g + (1-\lambda)h]$, that is $\lambda f + (1-\lambda)h \sim_A \lambda f A g + (1-\lambda)h$, thus $f \sim_A^* f\,A\,g$. Conversely, if $f \sim_A^* f\,A\,g$ for every $f, g \in \mathfrak{F}$, then in particular $f \sim_A f\,A\,g$ for every $f, g \in \mathfrak{F}$.

$(ii) \Leftrightarrow (iii)$: By Eq. (3), $f \sim_A^* fAg$ iff $P(u_A(f)) = P(u_A(fAg))$ for all $P \in \mathcal{C}_A$. It immediately follows that $\underline{\mathcal{C}_A}(u_A(f)) = \underline{\mathcal{C}_A}(u_A(fAg))$. By Proposition 7, this is equivalent to $f \sim_A^{\downarrow} fAg$.

Conversely, suppose that $f \sim_A^{\downarrow} fAg$ for every $f, g \in \mathfrak{F}$. Consider $x \succ_A y$. Since $x \sim_A^{\downarrow} xAy$, it follows from Proposition 7 that

$$\begin{aligned} u_A(x) &= \min_{P \in \mathcal{C}_A} [u_A(x)\,P(A) + u_A(y)\,(1 - P(A))] \\ &= u_A(x) \min_{P \in \mathcal{C}_A} P(A) + u_A(y)\,(1 - \min_{P \in \mathcal{C}_A} P(A)). \end{aligned}$$

Since $u_A(x) > u_A(Y)$, this implies that $\min_{P\in\mathcal{C}_A} P(A) = 1$, or equivalently, that $P(A) = 1$ for all $P \in \mathcal{C}_A$. It follows that $P(u_A(f)) = P(u_A(fAg))$ for all $P \in \mathcal{C}_A$, which is equivalent to $f \sim_A^* fAg$.

A.5 Proof of Proposition 11

First, we observe that the fact that Eq. (7) implies $\mathcal{C}_A = \mathcal{C}|A$ is a consequence of Prop. A.1 in GMM. That it implies $u_A = u$ is seen by taking $f = x$ and $g = x'$ to show that $x \succcurlyeq_A x' \Leftrightarrow x \succcurlyeq x'$. The converse is trivial.

(*i*) $\Leftrightarrow$ (*ii*): $fAg \succcurlyeq^* g$ for all $P \in \mathcal{C}$ iff $\int_A u(f)\,dP + \int_{A^c} u(g)\,dP \geq \int_A u(g)\,dP + \int_{A^c} u(g)\,dP$ for all $P \in \mathcal{C}$ iff $\int_A u(f)\,dP \geq \int_A u(g)\,dP$ for all $P \in \mathcal{C}$ iff $P_A(u(f)) \geq P_A(u(g))$ for all $P \in \mathcal{C}$.

(*i*) $\Rightarrow$ (*iii*): Suppose that $u = u_A$. We first observe that it follows from Proposition 4 that for every $f \in \mathfrak{F}$ and $x \in X$, with obvious notation,

$$x \in C_A^*(f) \Longleftrightarrow \underline{\mathcal{C}_A}(u(f)) \leq u(x) \leq \overline{\mathcal{C}_A}(u(f)). \tag{12}$$

Next, we prove that for every $f \in \mathfrak{F}$ and $x \in X$, again with obvious notation,

$$x \in C^*(f\,A\,x) \Longleftrightarrow \underline{\mathcal{C}|A}(u(f)) \leq u(x) \leq \overline{\mathcal{C}|A}(u(f)). \tag{13}$$

To see this, apply again Proposition 4 to find

$$x \in C^*(f\,A\,x) \Longleftrightarrow \underline{\mathcal{C}}(u(f\,A\,x)) \leq u(x) \leq \overline{\mathcal{C}}(u(f\,A\,x)).$$

That is, $x \in C^*(fAx)$ iff both

$$\min_{P\in\mathcal{C}} \int_S u(f\,A\,x)\,dP \leq u(x) \tag{14}$$

and

$$u(x) \leq \max_{P\in\mathcal{C}} \int_S u(f\,A\,x)\,dP. \tag{15}$$

Denote resp. $\underline{P}$ and $\overline{P}$ the probabilities in $\mathcal{C}$ that attain the extrema in Eqs. (14) and (15). Then we can rewrite Eq. (14) as follows:

$$u(x) \geq \frac{1}{\underline{P}(A)} \int_A u(f)\,d\underline{P},$$

which is equivalent to saying that

$$u(x) \geq \min_{P\in\mathcal{C}} \int_S u(f)\,dP_A = \underline{\mathcal{C}|A}(u(f)).$$

Analogously, Eq. (15) can be rewritten as

$$u(x) \le \frac{1}{\overline{P}(A)} \int_A u(f)\, d\overline{P},$$

which is equivalent to

$$u(x) \le \max_{P\in\mathcal{C}} \int_S u(f)\, dP_A = \overline{\mathcal{C}|A}(u(f)).$$

This ends the proof of Eq. (13).

To prove (i), notice that $x \succcurlyeq x' \Rightarrow x \succcurlyeq_A x'$ obviously follows from the assumption $u = u_A$, and that, given Eqs. (12) and (13), Eq. (9) follows immediately from the assumption $\mathcal{C}|A = \mathcal{C}_A$.

$(iii) \Rightarrow (i)$: First, observe that the assumption that $x \succcurlyeq x' \Rightarrow x \succcurlyeq_A x'$ implies $u = u_A$ by Corollary B.3 in GMM. Hence, it follows from Eqs. (12) and (13) above that Eq. (9) is equivalent to

$$\underline{\mathcal{C}_A}(u(f)) \le u(x) \le \overline{\mathcal{C}_A}(u(f)) \Longleftrightarrow \underline{\mathcal{C}|A}(u(f)) \le u(x) \le \overline{\mathcal{C}|A}(u(f)).$$

In particular, this implies that for every $\varphi \in B_0(\Sigma, u(X))$,

$$\min_{P\in\mathcal{C}|A} P(\varphi) = \min_{Q\in\mathcal{C}_A} Q(\varphi) \tag{16}$$

The result that $\mathcal{C}|A = \mathcal{C}_A$ now follows from two applications of Prop. A.1 in GMM.

$(i) \Rightarrow (iv)$: By Proposition 7 and the assumption that $u_A = u$, for every $f \in \mathfrak{F}$ and $x \in X$ we have that $f \succcurlyeq_A^{\downarrow} x$ iff $Q(u(f)) \ge u(x)$ for all $Q \in \mathcal{C}_A$. Next, we show that

$$f\, A\, x \succcurlyeq^{\downarrow} x \Longleftrightarrow P(u(f)) \ge u(x) \quad \text{for all } P \in \mathcal{C}|A, \tag{17}$$

so that the result follows from the assumption that $\mathcal{C}|A = \mathcal{C}_A$.

To see why Eq. (17) holds, notice that by Prop. 7, $fAx \succcurlyeq^{\downarrow} x$ iff $P(u(fAx)) \ge u(x)$ for all $P \in \mathcal{C}$. Equivalently, for every $P \in \mathcal{C}$,

$$\int_A u(f)\, dP + (1 - P(A))\, u(x) \ge u(x),$$

which holds iff $P_A(u(f)) \ge u(x)$ (recall that $P(A) > 0$ for all $P \in \mathcal{C}$). In turn, the latter is equivalent to saying that $P(u(f)) \ge u(x)$ for every $P \in \mathcal{C}|A$.

$(iv) \Rightarrow (i)$: We first show (mimicking an argument of Siniscalchi (2001)) that Eq. (10) implies that $u_A = u$. To see this, notice that we have $u_A(x) \ge u_A(x')$ iff $x \succcurlyeq_A x'$ iff $x\, A\, x' \succcurlyeq^{\downarrow} x'$ iff

$$\min_{P\in\mathcal{C}} \left[u(x)P(A) + u(x')(1 - P(A))\right] \ge u(x').$$

Using the assumption that $\min_{P\in\mathcal{C}} P(A) > 0$, the latter is equivalent to $u(x) \geq u(x')$, proving that $u_A = u$.

We now show that $\mathcal{C}|A = \mathcal{C}_A$ by showing that Eq. (16) holds for every $\varphi \in B_0(\Sigma, u(X))$, so that the result follows again from Prop. A.1 in GMM. As argued above, Eq. (10) holds for f and x iff

$$P(u(f)) \geq u(x) \quad \text{for all } P \in \mathcal{C}|A \Longleftrightarrow Q(u(f)) \geq u(x) \quad \text{for all } Q \in \mathcal{C}_A.$$

Fix $\varphi \in B_0(\Sigma, u(X))$ and suppose that, in violation of Eq. (16), $\alpha \equiv \min_{P\in\mathcal{C}|A} P(\varphi) > \min_{Q\in\mathcal{C}_A} Q(\varphi) \equiv \beta$. Then there exists $\gamma \in (\beta, \alpha)$. Let x denote the consequence such that $u(x) = \gamma$. By the assumption we just made, we have that (if $f \in \mathfrak{F}$ is such that $u(f) = \varphi$),

$$\min_{P\in\mathcal{C}|A} P(u(f)) > u(x) \quad \text{and} \quad u(x) > \min_{Q\in\mathcal{C}_A} Q(u(f)),$$

which, as proved above (the proof for strict preference works *mutatis mutandis* as that for weak preference, recalling that $\mathcal{C}|A$ is weak* compact), is equivalent to $fAx \succ^{\downarrow} x$ and $f \prec^{\downarrow}_A x$, a contradiction. Suppose instead that $\alpha < \beta$ and let $\gamma \in (\alpha, \beta)$. In this case we obtain

$$\min_{P\in\mathcal{C}|A} P(u(f)) < u(x) \quad \text{and} \quad u(x) < \min_{Q\in\mathcal{C}_A} Q(u(f)),$$

which is equivalent to $f \succ^{\downarrow}_A x$ and $x \succ^{\downarrow} fAx$ (to see the latter, let $P^* \in \mathcal{C}$ be the probability whose posterior minimizes the left-hand inequality; it follows that $\min_{P\in\mathcal{C}} P(u(fAx)) \leq P^*(u(fAx)) < u(x)$), again a contradiction. This concludes the proof of Eq. (16).

References

ARROW, K. J., AND L. HURWICZ (1972): "An Optimality Criterion for Decision Making under Ignorance," in *Uncertainty and Expectations in Economics*, ed. by C. Carter, and J. Ford. Basil Blackwell, Oxford.

BEWLEY, T. (2002): "Knightian Decision Theory: Part I," *Decisions in Economics and Finance*, 25(2), 79–110, (first version 1986).

EPSTEIN, L. G., AND M. SCHNEIDER (2003): "Recursive Multiple-Priors," *Journal of Economic Theory*, 113, 1–31.

GHIRARDATO, P. (2002): "Revisiting Savage in a Conditional World," *Economic Theory*, 20, 83–92.

GHIRARDATO, P., F. MACCHERONI, AND M. MARINACCI (2004): "Differentiating Ambiguity and Ambiguity Attitude," *Journal of Economic Theory*, 118(2), 133–173.

GHIRARDATO, P., F. MACCHERONI, M. MARINACCI, AND M. SINISCALCHI (2003): "A Subjective Spin on Roulette Wheels," *Econometrica*, 17, 1897–1908.

GHIRARDATO, P., AND M. MARINACCI (2001): "Risk, Ambiguity, and the Separation of Utility and Beliefs," *Mathematics of Operations Research*, 26, 864–890.

——— (2002): "Ambiguity Made Precise: A Comparative Foundation," *Journal of Economic Theory*, 102, 251–289.

GILBOA, I., AND D. SCHMEIDLER (1989): "Maxmin Expected Utility with a Non-Unique Prior," *Journal of Mathematical Economics*, 18, 141–153.

HANANY, E., AND P. KLIBANOFF (2004): "Updating Preferences with Multiple Priors multiple priors," Mimeo, Northwestern University.

NEHRING, K. (2001): "Ambiguity in the Context of Probabilistic Beliefs," Mimeo, UC Davis.

PIRES, C. P. (2002): "A Rule for Updating Ambiguous Beliefs," *Theory and Decision*, 53, 137–152.

SCHMEIDLER, D. (1989): "Subjective Probability and Expected Utility without Additivity," *Econometrica*, 57, 571–587.

SINISCALCHI, M. (2001): "Bayesian Updating for General Maxmin Expected Utility Preferences," Mimeo, Princeton University.