

Siniscalchi, Marciano

Working Paper

Bayesian Updating for General Maxmin Expected Utility Preferences

Discussion Paper, No. 1366

Provided in Cooperation with:

Kellogg School of Management - Center for Mathematical Studies in Economics and Management Science, Northwestern University

Suggested Citation: Siniscalchi, Marciano (2001) : Bayesian Updating for General Maxmin Expected Utility Preferences, Discussion Paper, No. 1366, Northwestern University, Kellogg School of Management, Center for Mathematical Studies in Economics and Management Science, Evanston, IL

This Version is available at:

<https://hdl.handle.net/10419/221715>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Bayesian Updating for General Maxmin Expected Utility Preferences

Marciano Siniscalchi*

September 14, 2001

First draft—Comments welcome!

Abstract

A characterization of “generalized Bayesian updating” in a maxmin expected utility setting is provided. The key axioms are consequentialism and *constant-act dynamic consistency*. The latter requires that, if an arbitrary act f is preferred (inferior) to a constant act y conditional upon E , and if f dominates (is dominated by) y pointwise on the complementary event E^c , then f is unconditionally preferred (inferior) to y .

The result provides a basis for a model of dynamic choice that accommodates arbitrary unconditional maxmin EU preferences, and allows for deviations from full dynamic consistency related to ambiguity.

Standard Expected Utility (EU) preferences are *separable across events*. In a static setting, the notion of separability is formalized by Savage’s Postulate P2 (the “Sure-Thing Principle”). In a dynamic framework, separability corresponds to *dynamic consistency*: if the decision maker would prefer some course of action to another if she learned that some event has obtained, and also if she learned that the same event has not obtained, then she should prefer it even prior to learning whether or not the event in question has obtained. As is well-known, P2 and dynamic consistency are closely related (see e.g. Ghirardato, 2001).

In a static setting, Ellsberg (1961) demonstrates that separability may fail if the decision maker perceives some ambiguity in the relative likelihood of events. Thus, in a dynamic setting, it is at least plausible to expect some tension between ambiguity and dynamic consistency. Recent experimental evidence (Cohen et al., 2000) based on a dynamic version of the single-urn Ellsberg example seems to indicate that ambiguity may indeed lead to violations of dynamic consistency.

*Economics Department, Princeton University, Princeton, NJ 08544-1021. Email: marciano@princeton.edu. I thank Larry Epstein, Paolo Ghirardato and Fabio Maccheroni for extensive discussions. All errors are my own.

This note takes as primitive an arbitrary, unconditional preference relation belonging to the maxmin-expected utility (MEU) family (Gilboa and Schmeidler, 1989), and characterizes conditional preferences derived from it via *generalized Bayesian updating*. This provides the axiomatic basis for a model of dynamic choice that allows for deviations from dynamic consistency related to the decision maker’s perception of ambiguity.

In the MEU framework, beliefs are represented by a set of (prior) probability measures; according to generalized Bayesian updating, conditional beliefs are represented by the collection of Bayesian updates of all elements of the set of priors.

The characterizing axioms are:

- *Consequentialism*. When evaluating an act f conditional upon the realization of the event E , outcomes assigned by f at states outside of E do not matter.
- *Constant-Act Dynamic Consistency*. If an arbitrary act f is preferred (inferior) to a constant act y conditional upon E , and if f dominates (is dominated by) y pointwise on the complementary event E^c , then f is unconditionally preferred (inferior) to y .

I show that the decision maker’s preferences conditional upon an event E satisfy the above axioms if and only if they are derived from her unconditional MEU preferences via generalized Bayesian updating.

The setup and result are presented in Section 1. Further discussion and comparison with related literature can be found in Section 2.

1 Setup and Results

1.1 Basic Framework and Definitions

The Anscombe-Aumann (1963) setup is adopted here for simplicity; see the discussion section for details on how to extend the analysis to a Savage-style framework.

Consider a set S of states, endowed with an algebra Σ , a set X of prizes, the set Y of finite lotteries over X , and the set L of all acts, i.e. all bounded, Σ -measurable Y -valued functions on S .

The main objects of interest are binary relations $\{\succsim_E\}_{E \in \Sigma}$ on L , to be interpreted as conditional or unconditional preferences among acts. That is, for any event $E \in \Sigma$, the relation $\succsim_E$ represents the decision maker’s preference ordering of acts, conditional upon learning that the event E has occurred. It is convenient to denote $\succsim_S$ simply by $\succsim$.

Interpretation. The collection $\{\succsim_E\}_{E \in \Sigma}$ provides a standard (Gilboa and Schmeidler, 1993; Myerson, 1986) ‘reduced-form’ representation of the decision maker’s preferences in dynamic choice problems for which the state space is (isomorphic to) S , and outcomes are drawn from the set Y .

Some readers may find it helpful to briefly relate this representation to *decision trees*. In that setting, a *strategy* is a specification of actions at each decision node; an *act* encodes the payoff consequences of a strategy. Moreover, every specification of moves at each chance node defines a state; thus, it is natural to associate with each decision node an *event* that corresponds to a specification of moves only at *preceding* chance nodes. Then, the relation $\succsim_E$ reflects the decision maker's preferences among strategies at the decision node associated with E . The mapping from decision trees to systems of conditional preferences is further illustrated in the leading example of the next Section.

Since other formulations of dynamic decision problems are possible (and may be more convenient in specific settings), I prefer to adopt the more encompassing, if more abstract, 'reduced-form' specification above.

Throughout this note, $\succsim$ will be assumed to be non-trivial.

Assumption 1.1 (NT) *There exist $f, g \in L$ such that $f \succ g$.*

For any event $E \in \Sigma$ and pair of acts $f, g \in L$,

$$f E g(s) = \begin{cases} f(s) & s \in E \\ g(s) & s \in E^c \end{cases}$$

The non-trivial, unconditional preference relation $\succsim$ will be assumed to admit a maxmin EU representation.

Assumption 1.2 (MP) *There exists a weak*-closed, convex set of probability measures $C \subset \Delta(S, \Sigma)$ and an affine function $u : Y \rightarrow \mathbb{R}$ such that*

$$\forall f, g \in L, \quad f \succsim g \quad \Leftrightarrow \quad \min_{q \in C} \int_S u(f(s))q(ds) \geq \min_{q \in C} \int_S u(g(s))q(ds).$$

Note that, by Assumption NT and the results of Gilboa and Schmeidler (1989), the utility function u is unique up to a positive affine transformation, and the set C of measures is unique.

An event $E \in \Sigma$ is $\succsim$ -null iff, for all $y, y' \in Y$ such that $y \succ y'$, $y E y' \sim y'$. Note that, under MP, an event E is $\succsim$ -null iff $\min_{q \in C} q(E) = 0$. Thus, in particular, $\emptyset$ is null.

1.2 Axioms and Main Results

Consider the following three behavioral restrictions, stated as assumptions regarding an arbitrary conditional preference $\succsim_E$; the event E will be assumed to be non-null in the theorems below.

First, each conditional preference of interest is a weak order:

Axiom 1.1 (WO) *The conditional preference $\succsim_E$ is complete and transitive.*

Second, preferences conditional upon the event E are not affected by outcomes at states outside E . This is a version of *consequentialism*.

Axiom 1.2 (Cons) *For every pair of acts $f, h \in L$: $f \sim_E fEh$.*

Third, a weakening of the standard *dynamic consistency* axiom is imposed. Its motivation is discussed in the next section.

Axiom 1.3 (c-DC) *For every act $f \in L$ and outcome $y \in Y$:*

$$\begin{aligned} f \succsim_E y, \quad f(s) \succsim y \quad \forall s \in E^c &\Rightarrow f \succsim y; \\ f \precsim_E y, \quad f(s) \precsim y \quad \forall s \in E^c &\Rightarrow f \precsim y. \end{aligned}$$

Moreover, if the preference conditional upon E is strict, then so is the unconditional preference.

Observe that the dominance conditions $f(s) \succsim y$ and $f(s) \precsim y$ are stated in terms of the unconditional preference. It is clear that one could separately assume that conditional and unconditional preferences agree on Y , and state the dominance conditions in terms of the conditional preference $\succsim_{E^c}$. Note also that strict preference conditional on the event E is required to imply strict unconditional preference.

The main result follows; see the Appendix for the proof.

Theorem 1 *Assume that $\succsim$ satisfies MP and NT. Then, for any non-null event $E \in \Sigma$, the following are equivalent:*

- (1) *$\succsim$ satisfies WO, Cons and c-DC;*
- (2) *For all acts $f, g \in L$,*

$$f \succsim_E g \quad \Leftrightarrow \quad \min_{q \in C} \int_E u \circ f \, q(ds|E) \geq \min_{q \in C} \int_E u \circ g \, q(ds|E).$$

2 Discussion

While not entirely uncontroversial, Consequentialism is assumed here in order to focus on the issue of dynamic consistency. This section is thus devoted to the analysis of Axiom c-DC, its implications, and extensions. It also discusses alternative updating rules, and related literature.

2.1 Weakening Dynamic Consistency

As a starting point, full dynamic consistency may be defined as follows in the present setting.

Axiom 2.1 (DC) *For every pair of acts $f, g \in L$ and non-null event $E \in \Sigma$:*

$$f \succsim_E (\succ_E) g, \quad f \succsim_{E^c} g \quad \Rightarrow \quad f \succsim (\succ) g$$

Thus, c-DC weakens DC in two ways:

- It only involves comparisons of acts with constants.
- It involves dominance, not just conditional preference, on one subevent.

To clarify the rationale for the above weakenings, it is helpful to focus on a concrete example. I adopt an explicitly dynamic formulation and illustrate how it relates to the ‘reduced-form’ analysis of the previous section.

The following dynamic decision problem is adapted from Seidenfeld and Wasserman (1993). An urn contains an *unspecified*, but equal number n of black and white balls. Two balls will be drawn without replacement; note that the prior probability of drawing a black (or white) ball on the i -th draw is $\frac{1}{2}$ independently of $i = 1, 2$ and n . However, given the outcome of the first draw, the posterior probability of drawing a black ball in the second extraction is $\frac{n-1}{2n-1}$ if the first ball drawn was black, and $\frac{n}{2n-1}$ if it was white.

A risk-neutral decision maker is offered a bet on the outcome of the *second* draw, according to the following protocol. At time 0, the decision maker can choose to opt out (o_1) of the gamble, which yields a payoff of 0, or accept it (b_2). If she accepts it, then, at time 1, she learns the outcome of the first draw (B_1 or W_1). She can then decide to opt out (o_2), this time with a payoff of -1 , or continue with the bet (c). Finally, at time 2, all uncertainty is resolved; the decision maker receives a payoff of 3 if the second ball drawn is black (B_2), and -2 otherwise (W_2). The situation is depicted in Figure 1.

Let the state space be $S = \{B_1B_2, B_1W_2, W_1B_2, W_1W_2\}$. To reflect the assumption that the number of balls is left unspecified, consider the set of probability measures

$$C = \left\{ q \in \Delta(S) : q(B_1B_2) = q(W_1W_2) = \frac{1}{2}p, \quad q(B_1W_2) = q(W_1B_2) = \frac{1}{2}(1-p); \quad p \in \left[0, \frac{1}{2}\right] \right\},$$

where $p = 0$ corresponds to the case of an urn containing exactly 2 balls, and $p = \frac{1}{2}$ corresponds to the limiting case of an urn containing infinitely many balls.

A *strategy* for the decision maker is a list of choices the decision maker plans to make at each decision node: e.g. $\sigma = (b_2, c, o_2)$ indicates an initial choice of b_2 , followed by c upon learning




Figure 1: A dynamic decision problem.

σ	$f_\sigma(B_1B_2)$	$f_\sigma(B_1W_2)$	$f_\sigma(W_1B_2)$	$f_\sigma(W_1W_2)$	min EU	min EU $_{B_1}$	min EU $_{W_1}$
(b_2, c, c)	3	-2	3	-2	0.5	-2	0.5
(b_2, o_2, c)	-1	-1	3	-2	-0.25	-1	0.5
(b_2, c, o_2)	3	-2	-1	-1	-1.5	-2	-1
(b_2, o_2, o_2)	-1	-1	-1	-1	-1	-1	-1
(o_1)	0	0	0	0	0	0	0

Table 1: Strategies and Acts

B_1 , and o_2 upon learning W_1 ; $\sigma = (o_1)$ denotes the strategy whereby the decision maker opts out immediately. Thus, a strategy is a function from decision nodes to choices. To each strategy σ corresponds a (‘reduced-form’) *act* f_σ , i.e. a function from states to outcomes. Table 1 maps strategies to acts, and indicates the corresponding minimum expected payoffs at each decision node. For the nodes following B_1 and W_1 , the set of posteriors is constructed via Generalized Bayesian updating.

I begin with two remarks on the formal properties of generalized Bayesian updating. The example demonstrates that this updating procedure may violate DC, and that its characterization via c-DC is “tight”.

To see that DC may fail, consider the non-null event B_1 ; note that $(b_2, o_2, c) \succ_{B_1} (b_2, c, c)$ and $(b_2, o_2, c) \sim_{W_1} (b_2, c, c)$, yet $(b_2, o_2, c) \prec (b_2, c, c)$.

To see that c-DC is tight, begin by reinterpreting the rankings above. One has $(b_2, o_2, c) \succ_{B_1} (b_2, c, c)$ and $(b_2, o_2, c)(s) \sim_{W_1} (b_2, c, c)(s)$ for all $s \in W_1$, yet $(b_2, o_2, c) \prec (b_2, c, c)$. Thus, it is not possible to allow for comparisons between arbitrary pairs of acts, even if dominance on E^c is retained. Next, consider the alternative set of priors

$$C' = \left\{ q : \in \Delta(S) : q(B_1 B_2) = q(W_1 W_2) = \frac{1}{2}p, q(B_1 W_2) = q(W_1 B_2) = \frac{1}{2}(1-p); p \in [0, 1] \right\}.$$

Let $\succ'_E$ denote the generic MEU preference determined by these beliefs. It is easy to see that $(b_2, c, c) \prec'_{B_1} (o_1)$ and $(b_2, c, c) \prec'_{W_1} (o_1)$ (the left-hand strategy has a minimum conditional expected payoff of -2), but $(b_2, c, c) \succ' (o_1)$. Thus, c-DC cannot be strengthened by allowing for conditional preference on E^c , even if the conditionally preferred act is constant.¹

Turn now to the interpretation of the above preferences. As noted above, the number of balls in the urn affects the conditional probability of drawing a black ball in the second extraction. In particular, if the first ball drawn is black (resp. white), the probability of drawing another black ball is increasing (resp. decreasing) in the total number of balls in the urn.

Then, loosely speaking, if a black ball is drawn first, an ambiguity-averse decision maker will evaluate the option of continuing with the bet assuming that the urn contains very few balls, and will therefore choose to opt out—even if this entails a cost.

On the other hand, ex-ante, the probability of drawing a black ball in the second extraction is $\frac{1}{2}$ regardless of the initial number of balls; differences in conditional probabilities *exactly average out*. Thus, ambiguity does not play a role in the ex-ante evaluation of the bet.

The example also suggests a possible interpretation of deviation from dynamic consistency (Axiom DC). Note first that, at time 0, a strategy comprises an immediate *choice* (i.e. b_2 or o_1) and a continuation *plan*. Regardless of their strength, dynamic consistency axioms impose direct restrictions on the latter only—although, of course, they do have implications for choice as well.

In this example, Axiom DC requires that the decision maker's time-0 plan respect her time-1 conditional ranking of choices. For instance, the conditional preference $(b_2, o_2, c) \succ_{B_1} (b_2, c, c)$ reveals (by Consequentialism) a conditional preference for choosing o_2 over c ; under DC, this would imply the ex-ante preference $(b_2, o_2, c) \succ (b_2, c, c)$.

Dynamic consistency is appealing if, in the decision maker's perception, the time-0 optimization problem is *separable* into sub-problems, which can be solved independently. Loosely speaking,

¹ It is also possible to show that, for a maxmin EU decision maker, $f \succ_E y$ and $f \succ_{E^c} y$ does imply $f \succ y$. In other words, dominance (vs. conditional preference) on E^c is only required if y is the dominating option. However, for a maxmax decision maker, the opposite would be true. The weaker, symmetric form used in the text thus accommodates both cases, and is sufficient to imply Generalized Bayesian updating.

time-1 conditional preferences may help ‘construct’ time-0 preferences in a recursive fashion. However, it is a widely accepted view that ambiguity manifests itself precisely via *non-separability* of preferences. Ellsberg (1961) relates ambiguity to violations of Savage’s Sure-Thing Principle; Epstein and Zhang (2001) *define* it in terms of a specific form of non-separability. Thus, one can at least expect some tension between dynamic consistency and ambiguity.

From a slightly different perspective, dynamic consistency is appealing if one can argue that the motivations supporting future choices are also applicable and relevant at time 0, when the decision maker evaluates her continuation plan. But, in the presence of ambiguity, this is not necessarily the case. Consider for example the choice of o_2 vs. c at the node following B_1 . The ambiguity-averse decision maker under consideration is worried that the urn might have initially contained *very few* balls (perhaps only one of each color), so that it is now unlikely (if not impossible) that another black ball will be drawn. She is, of course, aware that the urn might contain a large number of balls instead, but, since she is ambiguity-averse, this consideration is not relevant to her ex-post comparison of o_2 vs. c . On the other hand, when evaluating o_2 vs. c following B_1 as part of the strategy (b_2, o_2, c) , the decision maker is worried that the urn may contain a *large number* of balls, which is ‘bad news’ if W_1 obtains in the first draw. Again, she is aware of the possibility that the urn might contain very few balls, but, again, ambiguity aversion renders this consideration irrelevant.

To summarize, ambiguity may thus influence the decision maker’s evaluation of choices and plans in different ways at different decision nodes. This may lead to both non-separabilities and violations of dynamic consistency.

Of course, these observations are not meant to suggest that Axiom DC *should* be violated, or will *necessarily* be violated, in the presence of ambiguity. They merely indicate that, if one is willing to depart from separability in order to accommodate ambiguity, then one might also entertain weakening dynamic consistency, for closely related considerations.

Finally, Axiom c-DC does impose some restrictions on the relationship between present plans and future choices. I suggest that these restrictions involve comparisons that are robust to changes in the decision maker’s perception of ambiguity upon the arrival of new information.

To clarify, recall that, as noted above, ambiguity manifests itself via non-separabilities; in particular, the decision maker’s preferences, and hence her beliefs, may reveal *complementarities* between certain events (Epstein and Zhang, 2001, §3.1). For instance, in the example under consideration, the events B_1B_2 and W_1B_2 are complementary at time 0: depending on the number of balls in the urn, their probabilities can be as low as 0 and $\frac{1}{4}$, respectively; however, regardless of the number of balls in the urn, the probability of $B_2 = \{B_1B_2, W_1B_2\}$ is $\frac{1}{2}$. Loosely speaking, the dependency on the initial number of balls ‘averages out’.

The first part of Axiom c-DC reflects the assumption that, for an ambiguity-averse decision maker, *complementarities are favorable*. Consider for instance the strategy (b_2, c, c) . Conditional upon B_1 , the choice c prescribed by this strategy is unattractive because it leads to a favorable outcome only if B_1B_2 obtains; as noted above, an ambiguity-averse decision maker is worried that B_1B_2 might be very unlikely, or perhaps impossible. However, ex-ante, the complementarity between B_1B_2 and W_1B_2 (equivalently, the ‘averaging out’ of dependency on the number of balls) eliminates this concern. Thus, it makes sense to require that, if (b_2, c, c) is preferred to some constant act y conditional upon W_1 and pointwise² on B_1 , then it should also be preferred to y unconditionally.

The second part of Axiom c-DC reflects the fact that complementarities may lead to a more favorable evaluation of an act only by increasing the weight assigned to ‘good’ outcomes. For instance, conditional upon B_1 , the strategy (b_2, c, c) is unattractive because it delivers the favorable outcome 3 only if the event B_1B_2 obtains, and the latter receives zero weight as a consequence of ambiguity aversion. Ex-ante, the complementarity between B_1B_2 and W_1W_2 induces the decision maker to assign a higher weight to the favorable outcome 3, but this cannot lead to a preference for (b_2, c, c) over the certain outcome 3. Axiom c-DC encodes this requirement.

2.2 Alternative Updating Rules

Gilboa and Schmeidler (1993) consider a class of updating rules for (CEU and) MEU preferences. Given an unconditional preference relation $\succsim$, an act $h \in L$, and an event $E \in \Sigma$, the conditional preference $\succsim_E$ is the *h-Bayesian update* of $\succsim$ iff, for all acts $f, g \in L$, $f \succsim_E g$ if and only if $fEh \succsim gEh$.

h-Bayesian updates satisfy Consequentialism, but not necessarily Axiom c-DC. This suggests that the latter reflects dynamic consistency properties that are specific to generalized Bayesian updating.

In fact, generalized Bayesian updating may not correspond (for arbitrary unconditional MEU preferences) to any *h*-Bayesian update. However, generalized Bayesian updating is also characterized by a fixpoint condition that is somewhat reminiscent of *h*-Bayesian updating. Specifically, the conditional preference relation $\succsim_E$ is derived from $\succsim$ via generalized Bayesian updating if and only if, for every non-null event E , act f and constant lottery y ,

$$\forall f \in L, y \in Y : \quad f \sim_E y \quad \Leftrightarrow \quad fEy \sim y \quad (1)$$

(cf. Claim 1 in the Appendix).

²Cf. Footnote 1.

Finally, Gilboa and Schmeidler (1993) show that, for preferences that admit both a CEU and a MEU representation, the h -Bayesian update rule whereby $h = x^*$ is the most preferred prize³ in X characterizes “maximum-likelihood updating” in the MEU framework, and the Dempster-Shafer rule in the CEU setting. Define the conditional preference $\succsim_E^*$ by postulating that $f \succsim_E^* g$ if and only if $fEx^* \succsim gEx^*$; then, in the MEU representation,

$$f \succsim_F^* g \quad \Leftrightarrow \quad \min_{\substack{q \in C: \\ q(F) = \max_{q' \in C} q'(F)}} \int_F u(f(s))q(ds|F) \geq \min_{\substack{q \in C: \\ q(F) = \max_{q' \in C} q'(F)}} \int_F u(g(s))q(ds|F). \quad (2)$$

In other words, one only considers Bayesian updates of those priors $q \in C$ that maximize the probability of the conditioning event F . Intuitively, while generalized Bayesian updating tends to “preserve ambiguity”, maximum-likelihood updating tends to reduce it as much as possible. For this reason, maximum-likelihood and generalized Bayesian updating may be regarded as “polar opposites”.

2.3 Alternative Approaches to Generalized Bayesian Updating

In this note, DC is weakened essentially by restricting the set of *acts*, but not the collection of *events*, to which it applies. Epstein and Schneider (2001) and Wang (2000) take the opposite route, and obtain characterizations of recursive versions of the MEU model.

It should be noted that an explicit objective of both papers is to obtain a model of *dynamically consistent* choice for MEU decision makers. By way of contrast, the model of dynamic choice analyzed here is designed to allow for deviations from dynamic consistency.

Here, I discuss only some aspects of their results, (substantially) simplifying and adapting their definitions and notation to mine.

Fix a MEU preference $\succsim$ and a partition $\mathcal{F} = \{F_1, \dots, F_n\}$ of S ; and assume that each cell F_i is non-null. Along with the appropriate form of consequentialism, both Epstein-Schneider and Wang impose the following version of dynamic consistency:

Axiom 2.2 ($\mathcal{F}$ -DC) *For all acts $f, g \in L$: if $f \succsim_{F_i} g$ for all $i = 1, \dots, n$, then $f \succsim g$. If, additionally, $f \succ_{F_{i^*}} g$ for some i^* , then $f \succ g$.*

Axiom $\mathcal{F}$ -DC is necessary in order to obtain a recursive representation; conversely, recursivity is necessary in order to ensure that choices are dynamically consistent. Furthermore, the axiom is sufficiently weak to allow for a rich class of ambiguous beliefs.

However, Axiom $\mathcal{F}$ -DC does rule out certain preference patterns, including the modal preferences in the dynamic Ellsberg-style experiment of Cohen et al. (2000). In the example of

³Gilboa and Schmeidler (1993) assume that the set of prizes is bounded for $\succsim$.

Subsection 2.1, $\mathcal{F}$ -DC is violated for $\mathcal{F} = \{W_1, B_1\}$: for instance, $(b_2, o_2, c) \succ_{W_1} (b_2, c, c)$ and $(b_2, o_2, c) \sim_{B_1} (b_2, c, c)$, but $(b_2, c, c) \succ (b_2, o_2, c)$.

The comparison with the result of this note highlights the tension between ambiguity and dynamic consistency. Epstein and Schneider (2001) and Wang (2000) essentially resolve this tension by favoring dynamic consistency, at the expense of the ability to represent certain preference patterns. This note resolves it in favor of the latter, by relaxing dynamic consistency.

I conclude with two technical remarks. First, it can be shown that $\mathcal{F}$ -DC implies that all h -Bayesian updating rules discussed in Gilboa and Schmeidler (1993) yield the same conditional preferences as generalized Bayesian updating. Thus, unlike c-DC, this axiom does not reflect consistency properties that are unique to generalized Bayesian updating. Also, observe that Wang (2000) characterizes preferences that admit both a CEU and a MEU representation. This note deals with general MEU preferences.

Finally, Halpern (1998) provides an alternative, non-decision-theoretic approach to generalized Bayesian updating. The author defines an *update function* as a map from sets of prior probability measures to sets of conditional probability measures; he then exhibits a list of postulates on update functions that uniquely characterize generalized Bayesian updating.

2.4 CEU preferences

Theorem 1 clearly implies that, given *any* unconditional MEU preference relation, conditional preferences that satisfy WO, Cons and c-DC are also MEU preferences. In other words, the class of MEU preferences is closed under conditioning, if the latter conforms to these axioms.

It turns out that this is not the case for unconditional CEU preferences. More precisely, given an unconditional CEU preference $\succsim$ and an event E , the conditional preference relation $\succsim_E$ identified⁴ by Axioms WO, Cons and c-DC admits a CEU representation only if the capacity representing prior beliefs satisfies certain ‘separability’ restrictions. The Appendix provides an example and further discussion.

Clearly, since axioms in the spirit of $\mathcal{F}$ -DC imply⁵ c-DC, the same considerations hold for stronger notions of dynamic consistency.

2.5 Independence

Klibanoff (2001) introduces a behavioral definition of “stochastically independent randomizing de-

⁴By Claim 1, these axioms do uniquely identify a conditional preference relation.

⁵This is the case under the natural assumptions that all conditional preferences satisfy monotonicity and agree with the unconditional preference on the set of constant acts.

vice,” or SIRD, in a setting à la Savage. The proposed notion of stochastic independence is easy to interpret in view of the axioms adopted here.

Assume that $S = S_1 \times S_2$, and denote by L_1 the subset of acts $f \in L$ such that $f(s) = f(s')$ whenever $s = (s_1, s_2)$ and $s' = (s_1, s'_2)$ for some $s_1 \in S_1$ and $s_2, s'_2 \in S_2$. That is, acts in L_1 only depend on the realization of the first component of the Cartesian state space.

In Klibanoff (2001), S_1 is deemed stochastically independent of S_2 if, for any act $f \in L_1$ and outcome x , $f \sim x$ implies that $f(S_1 \times \{s_2\})x \sim x$ for all $s_2 \in S_2$.

If conditional preferences satisfy WO, Cons and c-DC, then Eq. (1) holds, so the condition proposed by Klibanoff can be restated as follows:

$$\forall s_2 \in S_2, \quad f \sim x \quad \Rightarrow \quad f \sim_{S_1 \times \{s_2\}} x.$$

In words, if the decision maker is unconditionally indifferent between f and x , then she must remain so upon learning that $s_2 \in S_2$ has obtained. Hence, the axioms proposed here allow for a simple and direct interpretation of Klibanoff’s definition.

2.6 Other Comments

The Anscombe-Aumann setup is adopted here merely as a matter of expositional convenience. The proof of Theorem 1 only relies on the fact that the utility function u is convex-ranged: in particular, this ensures that, for any act f and non-null event E , there exists a lottery y_f such that $u(y_f) = \min_{q \in C} \int_E u(f(s))q(ds|E)$.

Characterizations of MEU preferences in a Savage-like settings have been provided in Casadesus-Masanell et al. (1998a,b) and Ghirardato et al (2001). In all three papers, the utility function u is convex-ranged, so Theorem 1 holds verbatim in the setup adopted by these authors.

A Proof of Theorem 1

Recall that an event E is non-null iff $\min_{q \in C} q(E) > 0$; this implies that, for all $q \in C$, conditional probabilities $q(\cdot|E)$ are well-defined whenever E is non-null.

Throughout the proof, the notation $q_{f,E}$ denotes the generic minimizer of $\int_E u(f(s))q(ds|E)$ on C ; if $E = S$, I simply write q_f . Thus, in particular,

$$\int u(f(s))q_f(ds) \leq \int u(f(s))q_{f,E}(ds) \quad \text{and} \quad \int_E u(f(s))q_f(ds|E) \geq \int_E u(f(s))q_{f,E}(ds|E). \quad (3)$$

(2) $\Rightarrow$ (1): Fix a non-null $E \in \Sigma$. WO and Cons are obvious, as is c-DC in the trivial cases with $\min_{q \in C} q(E) = 1$; to see that c-DC holds with E such that $\min_{q \in C} q(E) \in (0, 1)$, consider first

the case $f \succsim_E y$, $f(s) \succsim y$ for all $s \in E^c$. Then

$$\begin{aligned} \int u(f(s))q_f(ds) &= q_f(E) \int_E u(f(s))q_f(ds|E) + \int_{E^c} u(f(s))q_f(ds) \geq \\ &\geq q_f(E) \int_E u(f(s))q_{f,E}(ds|E) + \int_{E^c} u(f(s))q_f(ds) \geq \\ &\geq q_f(E)u(y) + q_f(E^c)u(y) = u(y); \end{aligned}$$

the first inequality follows from Eq. (3). Hence, $f \succsim y$; moreover, if $f \succ_E y$, the last inequality is strict, and thus $f \succ y$. Next, consider the case $f \precsim_E y$, $f(s) \precsim y$ for all $s \in E^c$. Then

$$\begin{aligned} \int u(f(s))q_f(ds) &\leq \int u(f(s))q_{f,E}(ds) = \\ &= q_{f,E}(E) \int_E u(f(s))q_{f,E}(ds|E) + \int_{E^c} u(f(s))q_{f,E}(ds) \leq \\ &\leq q_{f,E}(E)u(y) + q_{f,E}(E^c)u(y) = u(y), \end{aligned}$$

where the first inequality follows from Eq. (3). Thus, $f \precsim y$, again with strict preference if $f \prec_E y$.

(1) $\Rightarrow$ (2): fix a non-null $E \in \Sigma$. Begin with two preliminary claims.

Claim 1: For all acts f and outcomes y , $f \succsim_E y \Leftrightarrow fEy \succsim y$ and $f \precsim_E y \Leftrightarrow fEy \precsim y$.

To see this, suppose $f \succsim_E y$. By Cons, $fEy \sim_E f \succsim_E y$. Clearly, $fEy(s) \sim y$ for all $s \in E^c$. Thus, by c-DC, we get $fEy \succsim y$. If instead $f \prec_E y$, the same argument shows that $fEy \prec y$, which proves the first part of the claim. The second is proved similarly.

Claim 2: For all outcomes y, y' , $y \succsim_E y' \Leftrightarrow y \succsim y'$.

The preceding claim implies that $y \succsim_E y'$ iff $yEy' \succsim y'$; by MP, this is equivalent to $\pi u(y) + (1 - \pi)u(y') \geq u(y')$, where π equals either $\min_{q \in C} q(E)$ or $\max_{q \in C} q(E)$. In either case, $\pi > 0$ because E is non-null, so the preceding expression reduces to $u(y) \geq u(y')$. This implies the claim.

To complete the proof, by Claims 1 and 2 it is sufficient to show that

$$fEy \sim y \quad \Leftrightarrow \quad u(y) = \min_{q \in C} \int_E u(f(s))q(ds|E). \quad (4)$$

To see that this is sufficient, assume Eq.(4) holds; then, it is always possible to find outcomes y_f and y_g such that $fEy_f \sim y_f$ and $gEy_g \sim y_g$, respectively. By Claim 1, $f \sim_E y_f$ and $g \sim_E y_g$, so $f \succsim_E g$ iff $y_f \succsim_E y_g$; by Claim 2, this is equivalent to $y_f \succsim y_g$, i.e. to $\min_{q \in C} \int_E u(f(s))q(ds|E) \geq \min_{q \in C} \int_E u(g(s))q(ds|E)$, as needed.

Now turn to the proof of Eq. (4). Assume $fEy \sim y$; then

$$q_{fEy}(E^c)u(y) + q_{fEy}(E) \int_E u(f(s))q_{fEy}(ds|E) = \int u(fEy(s))q_{fEy}(ds) = u(y).$$

Collecting terms involving $u(y)$ in the r.h.s. and dividing both sides by $1 - q_{fEy}(E^c) = q_{fEy}(E) > 0$ yields $u(y) = \int_E u(f(s))q_{fEy}(ds|E) \geq \int_E u(f(s))q_{f,E}(ds|E)$; the inequality follows from the definition of $q_{f,E}$. On the other hand, by the definition of q_{fEy} ,

$$\begin{aligned} q_{f,E}(E^c)u(y) + q_{f,E}(E) \int_E u(f(s))q_{f,E}(ds|E) &= \int u(fEy(s))q_{f,E}(ds) \geq \\ &\geq \int u(fEy(s))q_{fEy}(ds) = \\ &= u(y), \end{aligned}$$

and hence, proceeding as above, $u(y) \leq \int_E u(f(s))q_{f,E}(ds|E)$. Thus, $u(y) = \int_E u(f(s))q_{f,E}(ds|E) = \min_{q \in C} \int_E u(f(s))q(ds|E)$.

Conversely, assume $u(y) = \min_{q \in C} \int_E u(f(s))q(ds|E) = \int_E u(f(s))q_{f,E}(ds|E)$. Then

$$\int u(fEy(s))q_{fEy}(ds) = q_{fEy}(E^c)u(y) + q_{fEy}(E) \int_E u(f(s))q_{fEy}(ds|E) \geq u(y),$$

because $\int_E u(f(s))q_{fEy}(ds|E) \geq u(y) = \int_E u(f(s))q_{f,E}(ds|E)$ by the definition of $q_{f,E}$. On the other hand,

$$\int u(fEy(s))q_{fEy}(ds) \leq \int u(fEy(s))q_{f,E}(ds) = q_{f,E}(E^c)u(y) + q_{f,E}(E) \int_E u(f(s))q_{f,E}(ds|E) = u(y),$$

where the inequality follows from the definition of q_{fEy} . Thus, $fEy \sim y$.

B Remarks on CEU preferences

Let $S = \{s_1, s_2, s_3, s_4\}$ and $\Sigma = 2^S$. Also let $X = \mathbb{R}$ and $u(x) = x$ for all $x \in X$. Let p be the uniform probability distribution on S , and define a capacity ν by letting $\nu(E) = p^2(E)$ for all $E \in \Sigma$. Thus, ν and u jointly define a preference relation that lies in the intersection of the MEU and CEU families.

Now let $E = \{s_1, s_2, s_3\}$. I claim that, if $\succsim_E$ satisfies WO, Cons and c-DC, then $\succsim_E$ does not admit a CEU representation.

Define two acts f and g by letting $f = [22, 2, 1, 0]$ and $g = [11, 5, 1, 0]$, where $[h_1, h_2, h_3, h_4]$ represents the act h such that $h(s_i) = h_i$ for $i = 1, \dots, 4$. It may be verified that

$$fE3 \sim gE3 \sim 3;$$

by Claim 1, this implies that $f \sim_E g \sim_E 3$: that is, 3 is the conditional certainty equivalent of both f and g . Note also that f and g are comonotonic. Hence, if $\succsim_E$ is a CEU preference relation,

$\frac{1}{2}f + \frac{1}{2}g = [16.5, 3.5, 1, 0] \sim_E g$. But a simple calculation shows that

$$\left(\frac{1}{2}f + \frac{1}{2}g\right) E3 \succ 3,$$

which, by Claim 1 and transitivity, yields $\frac{1}{2}f + \frac{1}{2}g \succ_E g$. Thus, $\succ_E$ cannot be a CEU preference.

The example generalizes immediately; the following condition is clearly necessary for the existence of a CEU representation of $\succ_E$:

$$\forall f, g \in L, y \in Y : \quad f, g \text{ comonotonic, } fEy \sim gEy \sim y \quad \Rightarrow \quad \left(\frac{1}{2}f + \frac{1}{2}g\right) Ey \sim y.$$

Note that, for the unconditional CEU preference $\succ$, comonotonicity implies that, since f and g are comonotonic, $f \sim g \Rightarrow \frac{1}{2}f + \frac{1}{2}g \sim g$; however, as the example demonstrates, fEy and gEy may fail to be comonotonic, even if f and g are. Thus, the above condition may be viewed as a strengthening of comonotonicity that imposes a form of separability of preferences across the events E, E^c . I conjecture that the condition may also be sufficient for the existence of a conditional CEU preference.

References

- Anscombe, F. and R. Aumann (1963): “A Definition of Subjective Probability,” *Annals of Mathematical Statistics*, 34, 199-205.
- Ramon Casadesus-Masanell, Peter Klibanoff, and Emre Ozdenoren (1998). “Maxmin expected utility over savage acts with a set of priors,” *Journal of Economic Theory*, 92:33–65, 2000.
- Ramon Casadesus-Masanell, Peter Klibanoff, and Emre Ozdenoren (1998). “Maxmin expected utility through statewise combinations,” *Economics Letters*, 66:49–54, 2000.
- Cohen, M., Gilboa, I., Jaffray, J.Y., and D. Schmeidler (2000): “An Experimental Study of Updating Ambiguous Beliefs,” *Risk, Decision, and Policy*, 5, 123-133.
- Ellsberg, D. (1961): “Risk, Ambiguity and the Savage Axioms,” *Quarterly Journal of Economics* 75, 643-669.
- Epstein, L. and M. Schneider (2001): “Recursive Multiple Priors,” mimeo, University of Rochester.
- Epstein, L. and J. Zhang (2001): “Subjective Probabilities on Subjectively Unambiguous Events,” *Econometrica* 69, 265-306.

- Ghirardato, P. (2001): "Revisiting Savage in a Conditional World," *Economic Theory*, forthcoming.
- Ghirardato, P., Maccheroni, F., Marinacci, M., and M. Siniscalchi (2001): "A Subjective Spin on Roulette Wheels," Caltech WP #1127.
- Gilboa, I. and D. Schmeidler (1989): "Maxmin Expected Utility with Non-Unique Prior," *Journal of Mathematical Economics*, 18, 141-153.
- Gilboa, I. and D. Schmeidler (1993): "Updating Ambiguous Beliefs," *Journal of Economic Theory* 55, 33-49.
- Halpern, J. and A. Grove (1998): "Updating Sets of Probabilities," *Proceedings of the Fourteenth Conference on Uncertainty in AI*, pp. 173-182.
- Klibanoff, P. (2001): "Stochastically Independent Randomization and Uncertainty Aversion," *Economic Theory* 18, 605-620.
- Myerson, R. (1986): "Multistage Games with Communication," *Econometrica*, 54, 323-358.
- Savage, L.J. (1954), *The Foundations of Statistics*. Wiley, New York.
- Schmeidler, D. (1989): "Subjective Probability and Expected Utility without Additivity," *Econometrica*, 57, 571-587.
- Seidenfeld, T. and Wasserman, L. (1993): "Dilation for Sets of Probabilities," *The Annals of Statistics*, 21 n.3, 1139-1154.
- Wang, T. (2000): "Conditional Preferences and Updating," mimeo, University of British Columbia.