

Updating preferences with multiple priors

ERAN HANANY

Faculty of Engineering, Tel Aviv University

PETER KLIBANOFF

Kellogg School of Management, Northwestern University

We propose and axiomatically characterize dynamically consistent update rules for decision making under ambiguity. These rules apply to the preferences with multiple priors of Gilboa and Schmeidler (1989), and are the first, for any model of preferences over acts, to be able to reconcile typical behavior in the face of ambiguity (as exemplified by Ellsberg's paradox) with dynamic consistency for all non-null events. Updating takes the form of applying Bayes' rule to subsets of the set of priors, where the specific subset depends on the preferences, the conditioning event, and the choice problem (i.e., a feasible set of acts together with an act chosen from that set).

KEYWORDS. Updating, dynamic consistency, ambiguity, Ellsberg, Bayesian, consequentialism.

JEL CLASSIFICATION. D81, D83, D91.

1. INTRODUCTION

A central question facing any theory of decision making under uncertainty is how preferences are updated to incorporate new information. Since updated preferences govern future choices, it is important to know how they relate to information contingent choices made *ex ante*. Dynamic consistency is the requirement that *ex ante* contingent choices are respected by updated preferences. This consistency is implicit in the standard way of thinking about a dynamic choice problem as equivalent to a single *ex ante* choice to which one is committed, and is thus ubiquitous in economic modeling.

Under subjective expected utility, updating preferences by applying Bayes' rule to the subjective probability is *the* standard way to update. Why is this so? Dynamic consistency is the primary justification for Bayesian updating. Not only does Bayesian updating imply dynamic consistency, but, if updating consists of specifying a conditional

Eran Hanany: hananye@post.tau.ac.il

Peter Klibanoff: peterk@kellogg.northwestern.edu

We thank Itzhak Gilboa, Jean-Marc Tallon, Larry Epstein, Marciano Siniscalchi, Paolo Ghirardato, Emre Ozdenoren, Sandeep Baliga, Christian Kellner, Erez Marom, Tomasz Strzalecki, Peter Wakker, Simon Grant, and especially Uzi Segal and Nabil Al-Najjar for helpful discussions and comments. The co-editor Bart Lipman and two anonymous referees offered very helpful comments. We also thank audience members who heard versions of this work at ENSAM–Paris, Erasmus University, Northwestern University, Rice University, and the RUD 2004, FUR XI, ISIPTA '05, and ORSIS '06 conferences.

Copyright © 2007 Eran Hanany and Peter Klibanoff. Licensed under the [Creative Commons Attribution-NonCommercial License 2.5](http://creativecommons.org/licenses/by-nc/2.5/). Available at <http://econtheory.org>.

probability measure for each (non-null) event, dynamic consistency implies these conditional measures must be the Bayesian updates.¹ Even under the view that Bayesian updating should be taken as given, this tells us that dynamic consistency comes “for free” under expected utility.

Since dynamic consistency leads to a well-established theory of updating under expected utility, it makes sense to ask what it implies for the updating of more general preferences. We pursue this strategy to update preferences of the max-min expected utility (MEU) form (Gilboa and Schmeidler 1989). MEU preferences are widely used in modeling ambiguity averse behavior, as exemplified by the famous Ellsberg (1961) paradoxes.² The main contribution of the paper is to provide the first update rules in the ambiguity literature, MEU or otherwise, that are dynamically consistent while maintaining the characteristic features of behavior under ambiguity as exemplified by the Ellsberg paradoxes.

In addition to our main characterization (Section 2.2) and existence (Section 2.3) results, we explore the delicate interplay between dynamic consistency and updating of MEU preferences. In Section 3.1, we show that many possible strengthenings of our dynamic consistency requirement result in the non-existence of update rules for MEU.³ When discussing related literature in Section 3.3, we show that previous definitions of dynamic consistency are strong enough to trigger this non-existence. Thus, to the extent that earlier literature investigated dynamic consistency as a requirement for updating under ambiguity, the results were negative in that one could update only a narrow subset of MEU preferences (e.g., Epstein and Le Breton 1993) or only on a restricted set of events (e.g., Epstein and Schneider 2003). The extension to all MEU preferences and all (non-null) events that we provide is critical to capturing dynamically consistent Ellsberg behavior.

As Machina (1989) has emphasized, once we move beyond expected utility and preferences are not separable across events, updating in a dynamically consistent way entails respecting these non-separabilities by allowing updated preferences to depend on more than just the conditioning event. For this reason, we will see that dynamic consistency naturally leads a decision maker (DM) concerned with ambiguity to adopt rules for updating beliefs that depend on prior choices and/or the feasible set for the problem. Results in Section 3.2 show the necessity of this dependence and explore ways in which it can and cannot be limited. It is this characteristic of our update rules that most sharply distinguishes them from the rules for updating MEU suggested in the literature (e.g., Gilboa and Schmeidler 1993).

In the remainder of the Introduction, we present a simple example to which we refer back throughout the paper. We use it to motivate and illustrate our approach and contrast it with that of the existing literature on dynamic choice models under ambiguity.

¹Although related results appear in the literature (e.g., Ghirardato 2002), we were unable to find one that says precisely this. We provide such a result in Section 3.2 (Proposition 6).

²For a survey of economic applications see Mukerji and Tallon (2004).

³We also discuss (in Section 3.1.2) a desirable strengthening that does not result in non-existence and show precisely how it further limits the set of update rules.




FIGURE 1. The two choice pairs in the dynamic three-color Ellsberg problem.

1.1 A motivating example

Consider a version of Ellsberg's three-color problem. An urn contains 120 balls, 40 of which are known to be black (B) and 80 of which are somehow divided between red (R) and yellow (Y), with no further information on the distribution. A ball is to be drawn at random from the urn, and the DM faces a choice among bets paying off depending on the color of the drawn ball. Any such bet may be written as a triple $(u_B, u_R, u_Y) \in \mathbb{R}^3$ where each ordinate represents the payoff if the respective color is drawn. Typical preferences have $(1, 0, 0) \succ (0, 1, 0)$ and $(0, 1, 1) \succ (1, 0, 1)$, reflecting a preference for the less ambiguous bets. Notice that these preferences entail a preference to bet on black over red when no bet is made on yellow and a preference to bet on red over black when a bet is made also on yellow. Now consider a simple dynamic version of this problem. In the dynamic version, there is an interim information stage, where the DM is told whether or not the drawn ball was yellow. The DM is allowed to condition her choice of betting on black or red on this information.⁴ The two relevant choice problems are depicted in Figure 1. In each tree, the nodes marked with circles represent the (partial) resolution of uncertainty about the color of the drawn ball, while the nodes marked with squares are choice nodes for the DM. In choice pair 1, the choice "Bet on B" leads to the payoff vector $(1, 0, 0)$ while the choice "Bet on R" leads to payoffs $(0, 1, 0)$. Similarly, in choice pair 2, the choice "Bet on B" leads to the payoff vector $(1, 0, 1)$ while the choice "Bet on R" leads to payoffs $(0, 1, 1)$. Observe that in choosing between the bets $(1, 0, 0)$ and $(0, 1, 0)$, the opportunity to condition one's choice on the new information does not change the problem in an essential way: if a yellow ball is drawn, either choice gives 0, so it is only conditional on the event $\{B, R\}$ that the choice changes payoffs. A similar statement applies to the choice between $(1, 0, 1)$ and $(0, 1, 1)$. Therefore, preferences should remain $(1, 0, 0) \succ (0, 1, 0)$ and $(0, 1, 1) \succ (1, 0, 1)$ as in the original three-color problem.

⁴Such dynamic extensions of Ellsberg have been suggested before in the literature. See e.g., Epstein and Schneider (2003), Section 4.1.

These preferences are inconsistent with backward induction, which requires the DM to snip the tree at the node following the event $\{B, R\}$ and to choose as if this were the entire problem. But then the choice between $(1, 0, 0)$ and $(0, 1, 0)$ must be the same as the choice between $(1, 0, 1)$ and $(0, 1, 1)$ since the snipped trees for the two choice pairs are identical, rendering the Ellsberg choices impossible. It follows that no model of dynamic choice under ambiguity implying backward induction can deliver the Ellsberg preferences in this example.

1.2 Current approaches and the example

The literature offering dynamic models of ambiguity may be roughly divided into two categories. Both categories of existing models rule out dynamic Ellsberg-type behavior.

First, there are models that rule out this behavior because they imply backward induction. The example illustrates the conflict between backward induction and dynamic Ellsberg behavior. This critique encompasses a wide variety of dynamic models in the ambiguity literature, including those of Eichberger et al. (2005), Epstein and Schneider (2003), Hayashi (2005), Klibanoff et al. (2006), Maccheroni et al. (2006), Sarin and Wakker (1998), Siniscalchi (2006), and Wang (2003).

The second category consists of models in which behavior is naively inconsistent, in the sense that an ex ante choice or plan of what to do contingent on the event $\{B, R\}$ may differ from what will actually be chosen when and if that event occurs. The existing literature on updating rules for MEU preferences (or for any other ambiguity preference model) falls into this category, in that all of the rules examined, if used to define conditional preferences, result in inconsistent behavior. This literature includes many well-known rules, such as full (or generalized) Bayesian updating,⁵ maximum likelihood updating,⁶ and Dempster–Shafer updating.⁷ Applying any of these rules to the dynamic Ellsberg problem results in a reversal of ex ante contingent choice in at least one of the two choice pairs.

For example, applying full Bayesian updating to MEU preferences delivering the Ellsberg choices ex ante results in a conditional preference for black over red, reversing the ex ante contingent choice in choice pair 2. To make this explicit, consider the following MEU preferences that deliver the Ellsberg choices ex ante. These preferences together with the state space and information structure from the dynamic Ellsberg example are useful throughout the paper, and are invoked in almost all our examples. For any MEU preference over payoff vectors in \mathbb{R}^3 , there exists a compact and convex set of probability measures, C , over the three colors and a utility function, $u : \mathbb{R} \rightarrow \mathbb{R}$, such

⁵This rule calls for updating each measure in a set of priors according to Bayes' rule. See Jaffray (1992, 1994), Fagin and Halpern (1991), Wasserman and Kadane (1990), Walley (1991), Sarin and Wakker (1998), Pires (2002), Siniscalchi (2001), Wang (2003), and Epstein and Schneider (2003) for papers suggesting, deriving, or characterizing this update rule in various settings.

⁶This rule says that, of the set of priors, only those measures assigning the highest probability to the observed event should be updated using Bayes' rule. The other priors should be discarded. See Gilboa and Schmeidler (1993).

⁷See Dempster (1968) and Shafer (1976). Gilboa and Schmeidler (1993) show that Dempster–Shafer updating is equivalent to maximum likelihood updating for sets of priors that are cores of convex capacities.

that $\forall f, g \in \mathbb{R}^3, f \succsim g \iff \min_{q \in C} \int (u \circ f) dq \geq \min_{q \in C} \int (u \circ g) dq$. Let $u(x) = x$ for all $x \in \mathbb{R}$, and let $C = \{(\frac{1}{3}, \alpha, \frac{2}{3} - \alpha) \mid \alpha \in [\frac{1}{4}, \frac{5}{12}]\}$, a set of measures symmetric with respect to the probabilities of red and yellow. Observe that, indeed, $(1, 0, 0) \succ (0, 1, 0)$ and $(0, 1, 1) \succ (1, 0, 1)$ according to these preferences. If we apply full Bayesian updating conditional on the event $E = \{B, R\}$, the updated set of measures is $C_E = \{(\alpha, 1 - \alpha, 0) \mid \alpha \in [\frac{4}{9}, \frac{4}{7}]\}$. According to these updated preferences, “Bet on B” is strictly preferred to “Bet on R” conditional on learning $E = \{B, R\}$, leading $(1, 0, 1)$ to be selected over $(0, 1, 1)$ in choice problem 2, in conflict with the unconditional preferences.

1.3 Our approach

How then may preferences in the dynamic Ellsberg problem that are consistent with the ex ante Ellsberg choices be modeled? Our approach is to look for update rules for MEU preferences that, unlike those in the existing literature, are dynamically consistent and so do not suffer from naive reversals. Key to the analysis is formulating a definition of dynamic consistency that sticks as narrowly as possible to the notion, mentioned above, that ex ante optimal contingent choices should be respected when a planned-for contingency arises. With such a definition, we show that not only do such rules exist, but they have nice characterizations—they apply Bayes’ rule to some of the probability measures used in representing the DM’s unconditional preferences. The rules are shown to necessarily have the feature that the updated preferences vary with the choice problem encountered: past choices and/or the choice sets from which they were selected matter in determining which measures are updated. This approach is in the spirit of [McClellenn \(1990\)](#) and [Machina \(1989\)](#), who recognize that consistency in choices may be achieved through having later choices be influenced by earlier ones.

When applied to the dynamic Ellsberg problem, all of the dynamically consistent rules we derive deliver the Ellsberg choices. To give a preview of how this works, consider again the MEU preferences with $u(x) = x$ for all $x \in \mathbb{R}$, and $C = \{(\frac{1}{3}, \alpha, \frac{2}{3} - \alpha) \mid \alpha \in [\frac{1}{4}, \frac{5}{12}]\}$. Dynamically consistent updating in the dynamic Ellsberg problem corresponds to updating the set of measures to be (any closed, convex subset of) $C_E^1 = \{(\alpha, 1 - \alpha, 0) \mid \alpha \in [\frac{1}{2}, \frac{4}{7}]\}$ in choice problem 1 and (any closed, convex subset of) $C_E^2 = \{(\alpha, 1 - \alpha, 0) \mid \alpha \in [\frac{4}{9}, \frac{1}{2}]\}$ in choice problem 2. The alert reader might notice that one consistent update rule for this problem simply sets $C_E = \{(\frac{1}{2}, \frac{1}{2}, 0)\}$ for both choice problems, and thus might be puzzled by our claim that consistency requires updated preferences to vary with the choice problem. This claim can be illustrated by making a small modification to the feasible sets of acts in the problems above. Specifically, in addition to the “Bet on B” and “Bet on R” options, suppose that a third option (called “Hedge”) is available in each choice problem and yields payoffs $(\frac{19}{36}, \frac{19}{36}, 0)$ in choice problem 1 and $(\frac{19}{36}, \frac{19}{36}, 1)$ in choice problem 2. One may check that the unconditionally optimal choices do not change. But now, to maintain those choices after updating, one must have (any closed, convex subset of) $C_E^1 = \{(\alpha, 1 - \alpha, 0) \mid \alpha \in [\frac{19}{36}, \frac{4}{7}]\}$ in choice problem 1 and (any closed, convex subset of) $C_E^2 = \{(\alpha, 1 - \alpha, 0) \mid \alpha \in [\frac{4}{9}, \frac{17}{36}]\}$ in choice problem 2. These updated sets of measures no longer have a non-empty intersection. This illustrates how dynamic consistency may force the measures that get updated to vary with the choice problem.

2. CHARACTERIZING UPDATE RULES

Section 2.1 sets out the formal framework, notation and some preliminary definitions. Section 2.2 defines dynamic consistency and characterizes dynamically consistent update rules for MEU preferences. Section 2.3 proves existence of such rules by identifying the dynamically consistent rule that is *ambiguity maximizing* in that it comes closest to updating each measure in the DM's set of measures.

2.1 Framework and preliminaries

Consider an Anscombe and Aumann (1963) framework, where X is the set of all simple (i.e., finite-support) lotteries over a set of consequences Z , S is a set of states of nature endowed with an algebra Σ of events, and \mathcal{A} is the set of all acts, i.e. Σ -measurable bounded functions $f : S \rightarrow X$. Abusing notation in the standard way, $z \in Z$ is used also to denote the degenerate lottery $\delta_z \in X$ assigning probability 1 to the prize z , and $x \in X$ is used also to denote the constant act for which $\forall s \in S, f(s) = x$. Let \mathcal{P}^W be the set of all weak orders on \mathcal{A} . Let $\mathcal{P}^{MEU} \subseteq \mathcal{P}^W$ denote the set of non-degenerate max-min EU preference relations over \mathcal{A} (Gilboa and Schmeidler 1989). For any preference $\succsim \in \mathcal{P}^{MEU}$, there exists a compact and convex set of (finitely-additive) probability measures, C , and a von Neumann–Morgenstern expected utility function, $u : X \rightarrow \mathbb{R}$, such that $\forall f, g \in \mathcal{A}, f \succsim g \iff \min_{q \in C} \int (u \circ f) dq \geq \min_{q \in C} \int (u \circ g) dq$. If \succsim is non-degenerate, C is unique and u is unique (among vN–M expected utility functions) up to positive affine transformations. As usual, \sim and \succ denote the symmetric and asymmetric parts of \succsim . For $E \in \Sigma$ and $f, h \in \mathcal{A}$, we use $f_E h$ to denote the act equal to f on E and h on E^c (the complement of E in S). Similarly, if a and b are real Σ -measurable bounded functions, we use $a_E b$ to denote the function equal to a on E and b on E^c . Let $\mathcal{N}(\succsim)$ denote the set of events $E \in \Sigma$ for which $\forall q \in C, q(E) > 0$. Throughout the paper we limit attention to updating on events that are non-null in this strong sense.⁸ For $E \in \Sigma$, let $\Delta(E)$ denote the set of all finitely-additive probability measures on Σ giving weight 0 to E^c . For any $q \in \Delta(S)$ with $q(E) > 0$, we denote by $q_E \in \Delta(E)$ the measure obtained through Bayesian conditioning of q on E .

Let \mathcal{B} denote the set of all non-empty subsets of acts $B \subseteq \mathcal{A}$ such that B is convex (with respect to the usual Anscombe–Aumann mixtures) and compact (according to the norm taking the supremum over states and Euclidean distance on lotteries in X). Elements of \mathcal{B} are considered feasible sets and their convexity can be justified, for example, by randomization over acts. Compactness is needed to ensure the existence of optimal acts.

⁸In common with standard Bayesian updating of subjective probabilities, we have nothing to say about updating on events that are obviously null in that $q(E) = 0$ for all $q \in C$. It can be shown that some of our main characterization results may be modified to apply to events lying in between, for which there are some $q \in C$ with $q(E) > 0$. However, the set of events cannot be expanded too much in this direction without running into problems with existence of dynamically consistent update rules. Epstein and Marinacci (forthcoming) axiomatize the restriction of MEU to the case where no such in-between events exist, and so provide one rationale for ignoring these events.

Assume a preference $\succsim \in \mathcal{P}^{MEU}$, an event $E \in \mathcal{N}(\succsim)$, and an act $g \in \mathcal{A}$ chosen according to \succsim from a feasible set $B \in \mathcal{B}$ before the realization of E (i.e., $g \succsim f$ for all $f \in B$). Denote by \mathcal{T} the set of all such quadruples (\succsim, E, g, B) . An update rule is a function $U: \mathcal{T} \rightarrow \mathcal{P}^W$, producing a conditional preference denoted by $\succsim_{E,g,B}$. Such a conditional preference is viewed as governing choice upon the realization of the conditioning event E .

Finally, it is useful to define two families of update rules that play a prominent role in the paper. The first consists of the update rules that work by changing the set of measures from the original C to a set $C_{E,g,B}$ concentrated on the conditioning event E :

DEFINITION 1. Let \mathcal{U} denote the set of all update rules producing MEU conditional preferences $\succsim_{E,g,B}$ representable using the same (up to normalization) vN–M utility function u as \succsim and non-empty, closed, and convex sets of conditional measures $C_{E,g,B}$ containing only measures in $\Delta(E)$.

The second family is the subset of \mathcal{U} satisfying the added restriction that $C_{E,g,B}$ contains only Bayesian updates of some measures in C :

DEFINITION 2. $\mathcal{U}^{Bayes} = \{U \in \mathcal{U} \mid C_{E,g,B} \subseteq \{q_E \mid q \in C\}\}$.⁹

The main positive results in the paper restrict attention to rules in \mathcal{U}^{Bayes} . Many of our negative results are stated for the larger family \mathcal{U} . We choose to start directly from these families of rules in the main text so as to focus attention on our more important results and discussion. The reader should not be misled by this into thinking that our results lack foundation. In [Appendix A](#), we provide simple preference axioms characterizing these two families of rules. These axioms may be used to make any of the results in the main text fully behavioral. We also show that many of our results proved using the restriction to \mathcal{U}^{Bayes} may be easily modified to apply to rules in \mathcal{U} . [Appendix A](#) also discusses the observability of conditional preferences. Other results and all proofs not appearing in the main text are collected in [Appendix B](#).

With these preliminaries out of the way, we turn to our main concern—dynamic consistency and its implications for updating MEU preferences.

2.2 Dynamically consistent update rules

In this section, we define dynamic consistency and describe and characterize all dynamically consistent update rules for MEU preferences that work by applying Bayes' rule to some of the measures in the representing set of measures, C . In other words, we characterize the dynamically consistent rules in \mathcal{U}^{Bayes} . Recall from the [Introduction](#) that full Bayesian updating, for instance, violates dynamic consistency.

What exactly do we mean by dynamic consistency? We mean that optimal contingent plans should be respected when a planned-for contingency arises. Formally, our

⁹The proper quantifier modifying $C_{E,g,B} \subseteq \{q_E \mid q \in C\}$ is “for each $(\succsim, E, g, B) \in \mathcal{T}$.” Furthermore, $C_{E,g,B}$ is required to be non-empty, closed, and convex. For brevity and to ease notation, we omit these qualifiers here and in all subsequent definitions of update rules and families of update rules.

axiom requires that, for each feasible set of acts, B , if act g is chosen from B unconditionally, the update rule leads it to remain optimal conditionally.

AXIOM DC (Dynamic Consistency). *For any $(\succsim, E, g, B) \in \mathcal{T}$, if $f \in B$ with $f = g$ on E^c , then $g \succsim_{E,g,B} f$.*

Observe that conditional optimality of g is checked against all feasible acts f such that $f = g$ on E^c . Why check conditional optimality only against these acts? Dynamic consistency is relevant only *ceteris paribus*, i.e., when exactly the same consequences occur on E^c . To make clear why this is reasonable, consider an environment where the DM has a fixed budget to allocate across bets on various events. It would be nonsensical (and would violate payoff dominance on the realized event) to require that the ex ante optimal allocation of bets remain better than placing all of one's bets on the realized event. This justifies the restriction of the conditional comparisons to acts that could feasibly agree on E^c .¹⁰

As identifying an appropriate definition of dynamic consistency is a crucial component to our analysis, [Section 3.1](#) contains extensive further discussion. We investigate there a number of directions in which **DC** might be strengthened. We show that many of these strengthenings lead to the non-existence of dynamically consistent update rules in our setting. We also present a strengthening that may be sensible and still admit update rules, and we characterize the resulting rules. The relationship of **DC** to previous formalizations of dynamic consistency in the literature is explored in [Section 3.3](#).

Why might dynamic consistency be desirable? Dynamic consistency of one form or another has often been put forward as a rationality criterion and thus, from a normative point of view, it is important to identify rules that satisfy some version of this property. Moreover, from the normative point of view, optimal acts are the most important acts on which dynamic consistency should be satisfied because those are the plans that are chosen. An example of a specific normative argument in favor of dynamic consistency appearing in the literature is the argument, referenced in [Section 3.1](#), that inconsistency may lead to dominated choices. Dynamic consistency is needed also to ensure that information has non-negative value (see e.g., [Wakker 1988](#)). This is easy to see in the dynamic Ellsberg example from the [Introduction](#)—if conditional choices differ from what is desired unconditionally, the DM ex ante strictly prefers to face the problem without the information as to whether E or E^c occurred as opposed to the situation in the [Introduction](#) where she is allowed to choose after learning this information. Dynamic consistency also makes it easier to describe an individual planning ahead and to make welfare statements in dynamic models. From a more psychological point of view, dynamic consistency may be viewed as a rationalization property: dynamically consistent update rules are those that support earlier choices or plans.

Now we examine the implications of **DC** for updating. We start by defining a key subset of the measures in C for each quadruple (\succsim, E, g, B) .

¹⁰A similar justification might also admit a slightly stronger version of the axiom where the qualification $f = g$ on E^c is replaced by $u(f(s)) = u(g(s))$ for all $s \in E^c$. If this change is made, then all of the results involving **DC** continue to hold when the same replacement is made in the statements of these results; furthermore, the same arguments can be used in the proofs.

DEFINITION 3. For $(\succsim, E, g, B) \in \mathcal{T}$, define the *measures in C supporting the conditional optimality of g* to be

$$Q^{E,g,B} = \left\{ q \in C \mid \int (u \circ g) dq \geq \int (u \circ f) dq \text{ for all } f \in B \text{ with } f = g \text{ on } E^c \right\}.$$

Denote by $Q_E^{E,g,B}$ the set of Bayesian conditionals on E of measures in $Q^{E,g,B}$.

There are two reasons why calling these sets “measures supporting the conditional optimality of g ” makes sense. The first is obvious: if we consider a conditional expected utility preference with measure $q_E \in Q_E^{E,g,B}$, then according to such a preference, g is conditionally optimal. The second reason is deeper: as we show next, for rules in \mathcal{U}^{Bayes} , the existence of a measure in $Q_E^{E,g,B}$ that is used to evaluate $u \circ g$ conditionally is equivalent to the conditional optimality of g .

The following result completely characterizes the set of update rules in \mathcal{U}^{Bayes} satisfying **DC**:

DEFINITION 4. $\mathcal{U}^{DC} = \{U \in \mathcal{U}^{Bayes} \mid Q_E^{E,g,B} \cap \operatorname{argmin}_{q \in C_{E,g,B}} \int (u \circ g) dq \neq \emptyset\}$.

PROPOSITION 1. \mathcal{U}^{DC} is the set of all update rules in \mathcal{U}^{Bayes} satisfying **DC**.

PROOF. We present a sketch of the proof here. The complete proof is in [Appendix B](#). From the geometry of convex optimization, g is conditionally optimal within $\{f \in B \text{ with } f = g \text{ on } E^c\}$ exactly when the utility image of that set and the conditional indifference curve (in utility space) containing $u \circ g$ have a supporting hyperplane in common on E at $u \circ g$. The work in the proof is in showing that the condition

$$Q_E^{E,g,B} \cap \operatorname{argmin}_{q \in C_{E,g,B}} \int (u \circ g) dq \neq \emptyset$$

is equivalent to the existence of such a hyperplane. Intuitively, the measures in $Q_E^{E,g,B}$ can be identified with (the normals to) hyperplanes containing $u \circ g$ that support the utility image of the set $\{f \in B \text{ with } f = g \text{ on } E^c\}$ on E , while the measures in

$$\operatorname{argmin}_{q \in C_{E,g,B}} \int (u \circ g) dq$$

are normals to the hyperplanes containing $u \circ g$ that support the conditional indifference curve containing $u \circ g$. □

Though the condition $Q_E^{E,g,B} \cap \operatorname{argmin}_{q \in C_{E,g,B}} \int (u \circ g) dq \neq \emptyset$ is easy to check for any given update rule, it is not as helpful for generating dynamically consistent update rules—to calculate

$$\operatorname{argmin}_{q \in C_{E,g,B}} \int (u \circ g) dq$$

one must already have the candidate updated sets of measures $C_{E,g,B}$ in hand (whereas $Q_E^{E,g,B}$ can be calculated directly from the unconditional preferences). For generating the dynamically consistent rules, the following alternative representation of \mathcal{U}^{DC} is useful.

$$\text{COROLLARY 1. } \mathcal{U}^{DC} = \left\{ \begin{array}{l} U \in \mathcal{U}^{Bayes} \text{ | for some } r \in Q_E^{E,g,B}, \\ r_E \in C_{E,g,B} \subseteq \{q_E \mid q \in C \text{ and } \int (u \circ g) dq_E \geq \int (u \circ g) dr_E\}. \end{array} \right\}.$$

This provides the following algorithm for constructing such update rules. Start with any $r \in Q_E^{E,g,B}$ and include its update r_E in $C_{E,g,B}$. Further members of \mathcal{U}^{DC} are found by adding to $C_{E,g,B}$ (while preserving convexity) updates of some measures in C satisfying $\int (u \circ g) dq_E \geq \int (u \circ g) dr_E$. Doing this for each $r \in Q_E^{E,g,B}$ traces out the entire \mathcal{U}^{DC} .

Recall the dynamic Ellsberg example in the **Introduction**. In choice pair 1 of that example, the feasible set $B = co\{(1, 0, 0), (0, 1, 0)\}$, $g = (1, 0, 0)$, and $Q_E^{E,g,B} = \{q \in C \mid q(\text{Black}) \geq q(\text{Red})\}$.¹¹ Given $C = \{(\frac{1}{3}, \alpha, \frac{2}{3} - \alpha) \mid \alpha \in [\frac{1}{4}, \frac{5}{12}]\}$, the algorithm indeed delivers the result $C_{E,g,B} =$ (any closed, convex subset of) $\{(\alpha, 1 - \alpha, 0) \mid \alpha \in [\frac{1}{2}, \frac{4}{7}]\}$ that was claimed in the **Introduction**. Similarly, for choice pair 2, $B' = co\{(1, 0, 1), (0, 1, 1)\}$, $g' = (0, 1, 1)$, $Q_E^{E,g',B'} = \{q \in C \mid q(\text{Red}) \geq q(\text{Black})\}$ and the algorithm delivers the previously claimed $C_{E,g',B'} =$ (any closed, convex subset of) $\{(\alpha, 1 - \alpha, 0) \mid \alpha \in [\frac{4}{9}, \frac{1}{2}]\}$.

If one considers feasible sets that are smooth at g , then the algorithm becomes even simpler, as going through it with only one r is sufficient. The reason for this is that smoothness guarantees a unique hyperplane through g supporting the part of the feasible set that agrees with g on E^c , so that there is only a single measure in $Q_E^{E,g,B}$.

COROLLARY 2. *If the feasible set B does not have a kink on E at the optimal act g , $Q_E^{E,g,B}$ is a singleton.*

An important question that has not yet been addressed is whether the set of dynamically consistent update rules defined above is non-empty. In the next section, we answer this in the affirmative by examining and proving existence of a specific update rule, denoted U^{DCmax} , that is shown to be an element of \mathcal{U}^{DC} .

2.3 Existence and ambiguity maximization

To prove existence of update rules satisfying our axioms, we identify and prove existence of rules that are most *conservative* among the rules characterized in the previous section. By *conservative* we mean maintaining the most ambiguity in the process of updating. Examining such rules is particularly illuminating because they reveal the precise extent to which dynamic consistency forces the DM to eliminate measures present in the unconditional set when updating. If, for example, one views full Bayesian updating (updating all measures in the initial set) as “the right thing to do” then examining these rules shows how far one must depart from this to maintain consistency. To begin our

¹¹We use *co* to denote the convex hull operator.

exploration, we first define what it means for one preference to display more ambiguity than another.

DEFINITION 5. Suppose $\succsim^1, \succsim^2 \in \mathcal{P}^{MEU}$. Say that \succsim^1 *displays more ambiguity than* \succsim^2 if for all $f \in \mathcal{A}$ and $x \in X$, $f \succsim^1 x \implies f \succsim^2 x$ and $f \succ^1 x \implies f \succ^2 x$.

This definition is essentially the comparative ambiguity aversion of [Ghirardato and Marinacci \(2002\)](#) (and is also closely related to [Epstein 1999](#)) in our setting.¹² The way to understand this is as follows. MEU preferences reflect an aversion to ambiguity, and so when ambiguity increases, this should be bad news for acts that may be affected by ambiguity relative to acts, such as constant acts, that are evaluated unambiguously. Thus, more ambiguity should correspond to general acts falling in the preference order relative to constant acts. In terms of the representation, the next lemma shows that enlarging the set of measures leads to a preference displaying more ambiguity.

LEMMA 1. Suppose $\succsim^1, \succsim^2 \in \mathcal{P}^{MEU}$ and $u^1 = u^2 = u$. $C^2 \subseteq C^1$ implies \succsim^1 displays more ambiguity than \succsim^2 . $C^2 \subset C^1$ implies \succsim^2 does not display more ambiguity than \succsim^1 .

We define an update rule to entail more ambiguity than another if it always produces an updated preference displaying more ambiguity than the preference updated by the other rule. This leads to the following notion of an ambiguity-maximizing update rule:

DEFINITION 6. We say that an update rule U is *ambiguity maximizing among a set* $\widehat{\mathcal{U}}$ of update rules if for all $U' \in \widehat{\mathcal{U}}$, $\succsim_{E,g,B}$ displays more ambiguity than $\succsim'_{E,g,B}$.

We start with an elementary application of this definition to the families of update rules \mathcal{U} and \mathcal{U}^{Bayes} and then proceed to identify the unique ambiguity-maximizing rule in the smaller family \mathcal{U}^{DC} . For \mathcal{U} and \mathcal{U}^{Bayes} , ambiguity maximization characterizes well known rules that always exist: conditional max-min utility, defined by $C_{E,g,B} = \Delta(E)$, and full Bayesian updating, denoted U^{FB} and defined by $C_{E,g,B} = \{q_E \mid q \in C\}$. These results follow immediately from [Lemma 1](#) and inspection of \mathcal{U} and \mathcal{U}^{Bayes} . Neither of these ambiguity-maximizing rules are dynamically consistent. After imposing **DC** we obtain the following rule that eliminates exactly those measures in C that do not conditionally evaluate the chosen act g highly enough. Recall that $Q_E^{E,g,B}$ is the set of conditionals of measures in C supporting the conditional optimality of g as in [Definition 3](#).

DEFINITION 7. U^{DCmax} is the update rule in \mathcal{U}^{Bayes} such that

$$C_{E,g,B} = \left\{ q_E \mid q \in C \text{ and } \int (u \circ g) dq_E \geq \min_{p \in Q_E^{E,g,B}} \int (u \circ g) dp \right\}.$$

¹²It is known that changes in ambiguity and in ambiguity attitude are potentially entangled in MEU models. Some formal support for the notion that, to the extent one separates them, it is ambiguity rather than ambiguity attitude that changes across the MEU class is in [Ghirardato et al. \(2004\)](#) and [Klibanoff et al. \(2005\)](#).

This rule updates the measures in C that conditionally evaluate g weakly higher than does at least one measure supporting the conditional optimality of g within $\{f \in B \text{ with } f = g \text{ on } E^c\}$. If such measures exist, it should not be surprising that this is the largest set of measures in C that may be updated while satisfying dynamic consistency. The next result proves the desired existence (using supporting hyperplane arguments) and verifies that this rule is indeed the ambiguity-maximizing one in \mathcal{U}^{DC} .

PROPOSITION 2. *U^{DCmax} exists and is the unique ambiguity maximizing update rule in \mathcal{U}^{DC} .*

Recall the dynamic Ellsberg example from the **Introduction**. The rules in \mathcal{U}^{DC} are presented for this example near the end of **Section 2.2**; U^{DCmax} corresponds to the largest set of updated measures among these rules. Thus, for choice pair 1, the updated set of measures according to U^{DCmax} is $\{(\alpha, 1 - \alpha, 0) \mid \alpha \in [\frac{1}{2}, \frac{4}{7}]\}$, so that all measures in C assigning more weight to black than red are updated. Similarly, for choice pair 2, the updated set is $\{(\alpha, 1 - \alpha, 0) \mid \alpha \in [\frac{4}{9}, \frac{1}{2}]\}$, the updates of all measures in C assigning more weight to red than black.

The definition of dynamic consistency, the fact that the rules in \mathcal{U}^{DC} characterize its implications for MEU updating, and the existence of such rules as shown through the investigation of U^{DCmax} are the main positive results of the paper. These results provide the first update rules in the ambiguity literature, MEU or otherwise, that permit consistent dynamic behavior while maintaining the characteristic features of behavior under ambiguity as exemplified by the Ellsberg example. In the remainder of the paper, we provide important discussion and extensions of these results and our dynamic consistency definition, and investigate their relation to the existing literature.

3. DISCUSSION AND EXTENSIONS

Our definition of dynamic consistency is a main ingredient in generating our results. As such, it is important to carefully examine alternative definitions and their impact on updating. In **Section 3.1**, we show that many possible strengthenings of dynamic consistency result in non-existence of update rules for MEU. We also discuss a strengthening that does not result in non-existence and show precisely how it further shapes the set of update rules. In terms of the update rules themselves, the main novelty is their dependence on the choice problem. In **Section 3.2**, we explore the necessity of this dependence and ways in which it can and cannot be limited. Related literature is discussed in **Section 3.3**.

3.1 Dynamic consistency

We now explore several possible strengthenings of our key dynamic consistency axiom and derive the implications of each for updating MEU preferences. We begin by describing three ways of strengthening **DC** motivated by suggestions in the literature. We then show the negative result that each eliminates the possibility of consistent update rules for MEU. We follow this by suggesting a strengthening that is compatible with consistent updating.

3.1.1 *Stronger consistency and impossibility* Recall that **DC** requires only conditional optimality of g among the feasible acts agreeing with g on E^c . Why not check that the ordering of *all* feasible acts agreeing with g on E^c is preserved conditionally? The following axiom does exactly this.

AXIOM DC1. For any $(\succsim, E, g, B) \in \mathcal{T}$, if $f, h \in B$ with $f = h = g$ on E^c and $f \succsim h$ then $f \succsim_{E,g,B} h$.

Requirements implying this appear in a number of places in the literature (e.g., Machina 1989, Machina and Schmeidler 1992, Epstein and Le Breton 1993, and Ghirardato 2002). Is such a stronger axiom desirable? This is debatable. At least two arguments that might be used to support **DC** do not seem to extend support to the additional requirements of **DC1**. First, the verbal essence of dynamic consistency involves reversals, which only ever have the opportunity to occur when they involve unconditionally optimal acts. As Machina (1989) writes (pp. 1636–7) “... behavior ... will be *dynamically inconsistent*, in the sense that ... actual choice upon arriving at the decision node would differ from ... planned choice for that node.” Second, many normative arguments in support of dynamic consistency, such as arguments showing how lack of consistency may lead to payoff-dominated outcomes (see e.g., Machina 1989, McClellenen 1990, Seidenfeld 2004, and Segal 1997), require only the conditional optimality of g . Most importantly for our purposes, we show below that no update rules in \mathcal{U} can satisfy **DC1**.

Epstein and Schneider 2003, when discussing differences between recursive multiple priors and the robust control model of Hansen and Sargent (2001) point out that the robust control model satisfies a version of dynamic consistency that checks only optimality of g . Aside from minor differences in the framework, the following is that condition.

AXIOM DC2. For any $(\succsim, E, g, B) \in \mathcal{T}$, if $f \in \mathcal{A}$ with $f = g$ on E^c , then $g \succsim f$ implies $g \succsim_{E,g,B} f$.

The only difference from **DC** is that comparisons with g are not restricted to acts in the feasible set. Why restrict comparisons of g to feasible acts? Again we point out that the essence of dynamic consistency involves reversals, which are only relevant if they involve ex ante feasible acts. Moreover, we show that if we impose **DC2**, consistent updating rules are impossible.

Finally, observe that **DC** allows the possibility that $g \succ f$ for some $f = g$ on E^c while conditionally $g \sim_{E,g,B} f$. In such a circumstance, it is true that the DM is willing to continue with the initially chosen act g , but this is only weakly so. Several authors (e.g., McClellenen 1990 and Segal 1997) suggest that dynamic consistency should rule out such shifts from unconditional strict preference to conditional indifference. The following axiom strengthens **DC** in precisely this way.

AXIOM DC3. For any $(\succsim, E, g, B) \in \mathcal{T}$, if $f \in B$ with $f = g$ on E^c , then $g \succ_{E,g,B} f$, and if $g \sim_{E,g,B} f$, then $g \sim f$.

Unfortunately, this requirement again renders impossible consistent updating for MEU.

We summarize and prove the three impossibility results in the next proposition. The proof works by example—for each strengthening of dynamic consistency, a slight variation on the dynamic Ellsberg example serves to prove non-existence of an update rule. The examples have the same state space, information structure, and unconditional preferences as the dynamic Ellsberg example. They differ only in the set of feasible acts available to the DM.

PROPOSITION 3. *No update rule in \mathcal{U} satisfies **DC1**, **DC2**, or **DC3**.*

PROOF. Consider the same state space, information structure, and preferences as in the dynamic Ellsberg example. So, \succsim are MEU preferences with $u(z) = z$ and $C = \{(\frac{1}{3}, \alpha, \frac{2}{3} - \alpha) \mid \alpha \in [\frac{1}{4}, \frac{5}{12}]\}$. The conditioning event is $E = \{Black, Red\}$. All update rules in \mathcal{U} yield $\succsim_{E,g,B}$ that can be represented using the same u as \succsim and using a set of measures contained in $\Delta(E)$. Thus all acts that give a weakly higher payoff if a black ball is drawn than if a red ball is drawn should be evaluated using the measure in the updated set of measures that puts the least weight on black. Denote this measure by $(\beta, 1 - \beta, 0)$, where it is understood that β may vary with g and B . The examples show that the consistency axioms each contradict the existence of such a measure. We assume all feasible acts give a payoff of 1 if a yellow ball is drawn, as in choice pair 2 of the Ellsberg example. The only thing that differs across the examples is the set of attainable payoffs for black and red.

DC1: Let the feasible set $B = co\{all\ acts\ of\ the\ form\ (a, b, 1)\ such\ that\ 4a + 3b \leq 29\ and\ a, b \geq 0\}$. One may verify that $g = (5, 3, 1)$ is unconditionally optimal within B and that $(\frac{17}{4}, 4, 1)$, $(\frac{23}{4}, 2, 1)$, $(\frac{27}{4}, 0, 1)$, and $(4, 3, 1)$ are all feasible and give higher payoffs for black than for red. Suppose $\beta < \frac{4}{7}$. Then $(\frac{17}{4}, 4, 1) \succ_{E,g,B} g$ in violation of **DC**. Suppose $\beta > \frac{4}{7}$. Then $(\frac{23}{4}, 2, 1) \succ_{E,g,B} g$ in violation of **DC**. Therefore $\beta = \frac{4}{7}$. But then $(\frac{27}{4}, 0, 1) \succ_{E,g,B} (4, 3, 1)$, contradicting the unconditional preference $(4, 3, 1) \sim (\frac{27}{4}, 0, 1)$ and thus violating **DC1**.

DC2: Let $B = \{(5, 3, 1)\}$ so that also $g = (5, 3, 1)$. Observe that $g \sim (\frac{17}{4}, 4, 1) \sim (\frac{23}{4}, 2, 1) \sim (\frac{31}{4}, 0, 1)$. Suppose $\beta < \frac{4}{7}$. Then $(\frac{17}{4}, 4, 1) \succ_{E,g,B} g$ in violation of **DC2**. Suppose $\beta > \frac{4}{7}$. Then $(\frac{23}{4}, 2, 1) \succ_{E,g,B} g$ in violation of **DC2**. Therefore $\beta = \frac{4}{7}$. But then $(\frac{31}{4}, 0, 1) \succ_{E,g,B} g$, violating **DC2**.

DC3: Let $B = co\{(1, 2, 1), (3, 0, 1)\}$. One may verify that $g = (2, 1, 1)$ is strictly optimal within B and that $(\frac{3}{2}, \frac{3}{2}, 1) \in B$. Suppose $\beta < \frac{1}{2}$. Then $(\frac{3}{2}, \frac{3}{2}, 1) \succ_{E,g,B} g$ in violation of **DC**. Suppose $\beta > \frac{1}{2}$. Then $(3, 0, 1) \succ_{E,g,B} g$ in violation of **DC**. Therefore $\beta = \frac{1}{2}$. But then $(\frac{3}{2}, \frac{3}{2}, 1) \sim_{E,g,B} g$, violating **DC3**. □

3.1.2 A more robust consistency We now introduce a property that may be used to strengthen **DC** while maintaining compatibility with updating. It says that all feasible acts that are unconditionally optimal and agree with g on E^c should remain optimal conditional on E . Under **DC**, all that is implied for such acts is that they are conditionally no better than g .

AXIOM PFI (Preservation of Feasible Optimal Indifference). *For any $(\succsim, E, g, B) \in \mathcal{T}$, if $f \in B$ with $f = g$ on E^c and $f \sim g$, then $f \sim_{E,g,B} g$.*

A motivation for this axiom is the view that dynamic consistency should permit indifference to be resolved conditionally differently than it was unconditionally. Just because an indifference was broken in favor of g unconditionally, why should it necessarily continue to be broken in favor of g conditional on E ? We hold no strong opinion on this question, but think it is important to know what imposing such a requirement means for updating.

The following example illustrates how imposing **PFI** may matter for updating. Consider again the same state space, information structure, and preferences as in the dynamic Ellsberg example. So preferences are MEU with u the identity and $C = \{(\frac{1}{3}, \alpha, \frac{2}{3} - \alpha) \mid \alpha \in [\frac{1}{4}, \frac{5}{12}]\}$. The conditioning event is $E = \{Black, Red\}$. Let the feasible acts be $B = co\{(\frac{4}{7}, \frac{4}{7}, 0), (1, 0, 0), (0, 1, 0)\}$. Observe that $(\frac{4}{7}, \frac{4}{7}, 0) \sim (1, 0, 0)$ and that both are optimal within B . Suppose the DM picks $g = (\frac{4}{7}, \frac{4}{7}, 0)$. One may calculate that $Q^{E,g,B} = C$, and thus, by **Proposition 1**, **DC** adds no restrictions to \mathcal{U}^{Bayes} , and any convex set of conditionals of measures in C satisfies **DC**. However, only $C_{E,g,B} = \{(\frac{4}{7}, \frac{3}{7}, 0)\}$ also yields $(\frac{4}{7}, \frac{4}{7}, 0) \sim_{E,g,B} (1, 0, 0)$ which is required by **PFI**. This shows how **PFI** makes dynamic consistency more robust, in the sense that it does not force commitment to a choice among various unconditionally optimal payoffs on E (like $(\frac{4}{7}, \frac{4}{7})$ vs. $(1, 0)$) until after E is realized. The “cost” of this robustness is the set of further restrictions on updating illustrated in the example. In general, these restrictions are equivalent to using an update rule contained in the following set.

DEFINITION 8. $\mathcal{U}^{DC \cap PFI} =$

$$\left\{ \begin{array}{l} U \in \mathcal{U}^{Bayes} \mid \text{for some } r \in Q^{E,g,B}, r_E \in C_{E,g,B} \subseteq \{q_E \mid q \in C \text{ and} \\ \int (u \circ f) dq_E \geq \int (u \circ g) dr_E \text{ for all } f \in B \text{ with } f = g \text{ on } E^c \text{ and } f \sim g. \end{array} \right\}$$

This set of rules differs from the rules in \mathcal{U}^{DC} (**Definition 4**) in only one respect. The condition $\int (u \circ f) dq_E \geq \int (u \circ g) dr_E$ for measures q_E in the updated set of measures is required to hold only for $f = g$ in \mathcal{U}^{DC} while it is required for all $f \in B$ with $f = g$ on E^c and $f \sim g$ in $\mathcal{U}^{DC \cap PFI}$. In other words, measures in the updated set must not simply value g highly enough, they must also do the same for all feasible acts that are unconditionally optimal and agree with g on E^c . Our next result shows that this fully characterizes the additional restrictions imposed by **PFI**.

PROPOSITION 4. $\mathcal{U}^{DC \cap PFI}$ is the set of all update rules in \mathcal{U}^{Bayes} satisfying **DC** and **PFI**.

We implied that strengthening dynamic consistency by imposing **PFI** is compatible with updating. Identifying and proving the existence of the unique ambiguity-maximizing update rule in $\mathcal{U}^{DC \cap PFI}$ suffices to back up this claim. As with our earlier result (**Proposition 2**) carrying out a similar exercise for \mathcal{U}^{DC} , this not only serves as an existence proof, but also shows the precise extent to which the axioms force the DM to eliminate measures present in the unconditional set when updating.

To define the ambiguity-maximizing rule in $\mathcal{U}^{DC \cap PFI}$, it is useful to define a subset, denoted $K^{E,g,B}$, of $Q^{E,g,B}$, the measures supporting the conditional optimality of g . Measures in $K^{E,g,B}$ are required to satisfy one additional condition beyond those

determining $Q^{E,g,B}$: $\int(u \circ g) dq_E = \int(u \circ f) dq_E$ for all $f \in B$ with $f = g$ on E^c and $f \sim g$. Thus, measures in $K^{E,g,B}$ may be thought of as supporting not only the conditional optimality of g , but also the conditional optimality of all unconditionally optimal acts agreeing with g on E^c .

DEFINITION 9. For $(\succsim, E, g, B) \in \mathcal{T}$,

$$K^{E,g,B} = \left\{ q \in C \mid \int(u \circ g) dq \geq \int(u \circ f) dq \text{ for all } f \in B \text{ with } f = g \text{ on } E^c \text{ and } \int(u \circ g) dq = \int(u \circ f) dq \text{ if, in addition, } f \sim g \right\}.$$

Denote by $K_E^{E,g,B}$ the set of Bayesian conditionals on E of measures in $K^{E,g,B}$.

The next definition specifies an update rule by modifying the requirements for the updated set of measures, $C_{E,g,B}$, in $\mathcal{U}^{DC \cap PFI}$ (Definition 8) by replacing \subseteq with $=$ and requiring the updated measure used to evaluate g , r_E , to be an element of

$$\operatorname{argmin}_{p \in K_E^{E,g,B}} \int (u \circ g) dp.$$

DEFINITION 10. $U^{DC \cap PFI \max}$ is the update rule in \mathcal{U}^{Bayes} such that

$$C_{E,g,B} = \left\{ q_E \mid q \in C \text{ and } \int(u \circ f) dq_E \geq \min_{p \in K_E^{E,g,B}} \int(u \circ g) dp \text{ for all } f \in B \text{ with } f = g \text{ on } E^c \text{ and } f \sim g \right\}.$$

This update rule exists and is the rule with the largest updated sets of measures satisfying the axioms:

PROPOSITION 5. $U^{DC \cap PFI \max}$ exists and is the unique ambiguity-maximizing update rule in $\mathcal{U}^{DC \cap PFI}$.

Finally, we note that if we were to strengthen **PFI** in the ways that **DC1**, **DC2**, or **DC3** strengthen **DC**, no update rule in \mathcal{U} would satisfy the stronger condition. The formal statement of the strengthenings and the result is in **Appendix B**. The proof is omitted as it is the same as for the impossibility results in **Proposition 3**.

Collectively, the results in this section have shown that **PFI** is a way to strengthen and “robustify” **DC** while maintaining the ability to update MEU preferences, and that many directions of strengthening **DC** and **PFI** are unsuitable for our purposes.

Although **DC** and **PFI** allow for updating to depend on the choice problem through g and B , the extent to which updating of MEU preferences *must* depend on these in order to satisfy **DC** and **PFI** has not yet been fully investigated. It is to this question we now turn.

3.2 Dependence of conditional preference on the past choice and feasible set

The most striking feature of the update rules we have characterized in this paper is that conditional preferences may vary depending on the choice of g and the feasible set B . In

contrast, more restrictive definitions of dynamic consistency that rule out such dependence in updating are not uncommon in the literature (see e.g., Ghirardato 2002). It is well-known that this more restrictive dynamic consistency makes it impossible to capture non-expected utility behavior under the assumption that the DM cares only about the acts induced by their strategy (see Ghirardato 2002 among many others). For the purpose of modeling behavior departing from expected utility, others in the literature on subjective uncertainty have allowed, as we do, the dependence of conditional preferences on g (e.g., Machina and Schmeidler 1992, Epstein and Le Breton 1993).¹³ One may wonder if it is really necessary for dynamically consistent updating of MEU preferences to allow dependence of the conditional preference on the feasible set B . To address this issue, the axiom below formally states what it means not to depend on B .

AXIOM IFS (Independence from Feasible Sets). *For every $(\succsim, E, g, B_1), (\succsim, E, g, B_2) \in \mathcal{T}$, $\succsim_{E,g,B_1} = \succsim_{E,g,B_2}$.*

Are there update rules that satisfy our axioms and do not depend on B ? Note that if one restricts attention to expected utility preferences, the answer is yes, and this uniquely identifies Bayes' rule. This is the content of the next proposition, which formalizes our claim in the Introduction about the relationship between dynamic consistency and Bayes' rule under expected utility.

PROPOSITION 6. *Bayes' rule is the unique update rule in \mathcal{U} satisfying DC and IFS when \succsim and $\succsim_{E,g,B}$ are restricted to be non-degenerate expected utility preferences.*

However, once we allow MEU preferences, our next result shows there are no longer any rules satisfying these conditions. The proof (in Appendix B) works by considering two examples with the same state space, information structure, and unconditional preferences as the dynamic Ellsberg example. The examples have a common g and differ only in the set of feasible acts. There is a kink at g in the unconditional preferences. All update rules in \mathcal{U} deliver preferences conditional on $E = \{Black, Red\}$ that are smooth at g . If the feasible sets have different slopes at g (permitted by the kink in \succsim), different updates are required to obtain conditional optimality. For one feasible set, updating measures that put too little weight on black leads to dynamic inconsistency. For the other feasible set, *not* updating measures that put too little weight on black leads to dynamic inconsistency. This shows that no common updating can work across these feasible sets.

PROPOSITION 7. *There does not exist an update rule in \mathcal{U} satisfying DC and IFS.*

Allowing dependence on the feasible set, over and above dependence on the conditioning event and the unconditionally optimal choice g , is therefore crucial to dynamically consistent updating for MEU preferences. Our next result makes this more precise

¹³Wakker (1997) presents an updating rule for non-expected utility with objective probabilities that depends on the feasible set beyond g . This is the only other update rule we know of in the literature that shares this feature with our rules.

by showing that dependence on only part of the feasible set is needed. Specifically, the crucial acts are those that agree with g on E^c and that are not strictly dominated on E by other feasible acts agreeing with g on E^c . Intuitively, checking that dynamic consistency is satisfied with respect to these undominated acts is sufficient to ensure that it is satisfied overall. Formally, the feasible acts undominated on E and agreeing with g on E^c are the following:

DEFINITION 11. Given $(\succsim, E, g, B) \in \mathcal{T}$, the *conditionally undominated feasible set* is

$$B^{Un,E,g} = \left\{ \begin{array}{l} f \in B \text{ with } f = g \text{ on } E^c \mid \text{for no } h \in B \text{ with } h = g \text{ on } E^c, \\ h(s) \succsim f(s) \text{ for all } s \in E, \text{ with } h(s) \succ f(s) \text{ for some } s \in E \end{array} \right\}.$$

What does it mean for updating to depend only on the conditionally undominated feasible set? It means that if everything else is held fixed, and two feasible sets have the same conditionally undominated parts, then the two updated preferences must be the same. This is exactly the content of the next axiom.

AXIOM **IDFS** (Independence from the Dominated Part of Feasible Sets). *For every $(\succsim, E, g, B_1), (\succsim, E, g, B_2) \in \mathcal{T}$, if $B_1^{Un,E,g} = B_2^{Un,E,g}$, then $\succsim_{E,g,B_1} = \succsim_{E,g,B_2}$.*

The next proposition shows that this limited independence is compatible with **DC** (and **PFI**), unlike the complete independence required in **IFS**.

PROPOSITION 8. *The rules U^{DCmax} and $U^{DC \cap PFI max}$ satisfy **IDFS**.*

Generally speaking, if one wanted to build the whole paper while considering only the conditionally undominated parts of feasible sets, very little would change. Specifically, the only change would be that the *option* to vary update rules with other aspects of the feasible set would be eliminated. This option is never used to prove dynamic consistency, existence, ambiguity maximality, or any other important properties of our rules.

Given that it is necessary to allow conditional preference to depend on the feasible set, one may wonder if it is also necessary to have further dependence on g . We show that the answer is no. To do this, we define what it means to be independent of g and identify the ambiguity-maximizing rule satisfying this condition plus our earlier axioms.

AXIOM **ICA** (Independence from the Chosen Act). *For every $(\succsim, E, g_1, B), (\succsim, E, g_2, B) \in \mathcal{T}$, $\succsim_{E,g_1,B} = \succsim_{E,g_2,B}$.*

DEFINITION 12. $U^{DC \cap IC max}$ is the update rule in \mathcal{U}^{Bayes} such that

$$C_{E,B} = \bigcap_g \left\{ \begin{array}{l} q_E \mid q \in C \text{ and } \int (u \circ f) dq_E \geq \min_{p \in K_E^{E,g,B}} \int (u \circ g) dp \\ \text{for all } f \in B \text{ with } f = g \text{ on } E^c \text{ and } f \sim g \end{array} \right\}.$$

Recall that $K_E^{E,g,B}$ is defined in **Definition 9** of **Section 3.1.2** and notice that the updated set of measures, $C_{E,B}$, is the intersection of the updated sets of measures produced by the rule $U^{DC \cap PFI max}$ as g varies.

PROPOSITION 9. $U^{DC \cap IC \text{max}}$ exists and is the unique ambiguity-maximizing update rule in $\mathcal{U}^{DC \cap PFI}$ satisfying **ICA**.

We note that **DC** and **ICA** together imply **PFI**, thus $U^{DC \cap IC \text{max}}$ is also the unique ambiguity-maximizing update rule in \mathcal{U}^{DC} satisfying **ICA**. Furthermore $U^{DC \cap IC \text{max}}$ does not depend on any conditionally dominated feasible acts and so satisfies **IDFS** as well. The following example shows that **ICA** can be restrictive. In the example, while the other axioms allow a substantial amount of ambiguity after conditioning no matter which optimal act is chosen as g , adding **ICA** removes the ambiguity entirely.

EXAMPLE 1. Suppose there are three states. Consider $\succsim \in \mathcal{P}^{MEU}$ with

$$C = co\{(0.3, 0.2, 0.5), (0.1, 0.7, 0.2), (0.7, 0.2, 0.1)\},$$

$Z = \mathbb{R}$, and $u(z) = z$ for $z \in Z$. When faced with the choice set

$$B = co\{(1, 16.75, 1), (2, 15.25, 1), (8.4, 5.9, 0.9), (7.4, 7.4, 0.9)\},$$

the DM is indifferent among all acts in B and all are evaluated by $(0.3, 0.2, 0.5)$. Suppose that the conditioning event is $E = \{1, 2\}$ (i.e., the DM learns that the true state is one of the first two) and that $g = (1, 16.75, 1)$. According to the update rule $U^{DC \cap PFI \text{max}}$, the updated set of measures is

$$\{q \in \Delta(E) \mid 0.6 \geq q_E(1) \geq 0.125\}.$$

If, instead, $g = (8.4, 5.9, 0.9)$ then the updated set of measures according to $U^{DC \cap PFI \text{max}}$ is

$$\{q \in \Delta(E) \mid \frac{7}{9} \geq q_E(1) \geq 0.6\}.$$

From the definition of $U^{DC \cap IC \text{max}}$, the updated set of measures according to any rule in \mathcal{U}^{DC} satisfying **ICA** is the singleton set

$$\{(0.6, 0.4, 0)\}. \quad \diamond$$

In sum then, consistency forces the selection of measures that get updated by Bayes' rule to depend on the choice problem through the conditionally undominated part of the feasible set, but does not require further dependence on g . The rule $U^{DC \cap IC \text{max}}$ is notable in exhibiting only the required dependence.

3.3 Related literature

Many papers, some of which we mention in the **Introduction**, have examined update rules for MEU preferences. **Gilboa and Schmeidler (1993)** describe, but do not characterize in terms of preferences, a large class of rules for updating sets of priors that they call *classical* update rules. This class includes both maximum likelihood and full Bayesian updating, among others, and turns out to be exactly the set of rules in \mathcal{U}^{Bayes}

that also are independent of g and B . Thus our axioms characterizing \mathcal{U}^{Bayes} (see [Appendix A](#)) plus **ICA** and **IFS** provide a preference characterization of the entire set of classical update rules. Update rules we propose and characterize in this paper, for example U^{DCmax} and $U^{DC \cap PFImax}$, are distinct from any previously mentioned in the literature. In particular, *all* of the previous rules in \mathcal{U} proposed for updating MEU preferences also are independent of the feasible set B . [Proposition 7](#) then implies that these rules must violate **DC**.

3.3.1 Dynamic consistency requirements Given the importance of **DC** to our paper, it is helpful to compare it to existing definitions of dynamic consistency in the literature. This comparison shows that our requirements are weak, and beneficially so. In particular, the existing definitions are too strong to provide a basis for updating MEU preferences and incorporate more than the idea that optimal contingent plans remain optimal contingently.

[McClennen \(1990\)](#) provides an excellent and deep analysis of the problem of rational dynamic choice and as a key part of this defines dynamic consistency (p. 120) within a formal framework of decision trees. In our terms, this definition seems closest to the following: For $k \in B$, $k \succsim h$ for all $h \in B$ if and only if for some $f \in B$ with $f = k$ on E , $f \succsim_{E,g,B} h$ for all $h \in B$ such that $f = h$ on E^c . In words, an act k is unconditionally optimal within the feasible set if and only if it agrees on the conditioning event E with an act f that is conditionally optimal among all feasible acts that agree with f on E^c . The ‘only if’ direction (running from unconditional to conditional preference) is stronger than **DC** and **PFI** together. Specifically, **DC** and **PFI** together imply that direction only for those f such that $f = g$ on E^c , where g was chosen unconditionally. The ‘only if’ direction of McClennen’s condition is what we would get if we were to impose, in addition to **DC** and **PFI**, that conditional preferences depended only on E and B and not on the act g (i.e., axiom **ICA** defined in the previous section). More importantly, the ‘if’ direction of McClennen’s definition is not implied by any combination of our axioms, and seems too strong a requirement. For example, consider an environment where the DM has a fixed budget to allocate across bets on various events. Allocating the entire budget to bets on events within E is clearly conditionally optimal given E , but this in no way implies it is unconditionally optimal to do so. If we strengthen **DC** to imply even a weaker version of this condition (weaker in that we require $f = g$ on E^c), we get an impossibility result (whose proof is the same as for the part of [Proposition 3](#) relating to **DC3**).

[Machina \(1989\)](#) discusses dynamic consistency in the context of preferences over lotteries in decision trees with known probabilities. As reflected in the quote we used in [Section 3.1](#), Machina’s verbal explanation goes quite well with our definition. However, when later in the same paper Machina proposes a rule for conditional preferences, his rule imposes a much stronger consistency property. [Machina and Schmeidler \(1992\)](#) adopt the subjective analogue of Machina’s rule, and thus provide the appropriate basis for comparison with our work. They require that if g is the initially chosen act and the event E occurs, then conditional preferences are defined by $f \succsim_{E,g} h$ if and only if $f_E g \succsim h_E g$. [Epstein and Le Breton \(1993\)](#) show that Machina and Schmeidler’s update rule together with [Savage’s \(1954\)](#) axioms minus P2 (the sure-thing principle) applied

to both unconditional and conditional preferences imply preferences must be probabilistically sophisticated and thus incompatible with Ellsberg behavior. This implies that using Machina and Schmeidler's update rule to define dynamic consistency leads to nonexistence in our setting. Our **Proposition 3** above shows that even **DC1**, a weaker version of Machina and Schmeidler's condition, is sufficient to generate impossibility in our problem.

Segal (1997) also discusses dynamic consistency in the context of preferences over lotteries in decision trees. The weakest dynamic consistency property he states (his axiom 1) says that any conditionally optimal choice must also be part of an unconditionally optimal plan. Since this axiom requires strict unconditional preference for an optimal act to remain strict conditionally, it leads to nonexistence in our setting (this follows from the part of **Proposition 3** relating to **DC3**). **Segal (1997)** goes on to propose an even stronger consistency axiom (his axiom 3) which weakens Machina's update rule to apply only to the indifference curve containing the initially chosen lottery. All of the analysis in **Segal (1997)** is carried out using an extension of this stronger axiom.

Gul and Lantto (1990) propose an axiom (called dynamic programming solvability) on preferences over lotteries in two-stage decision trees with known probabilities. They prove this axiom is equivalent to a consistency condition. The axiom can be understood as saying roughly the following: If one can bike or walk to work and both are optimal plans contingent on the weather, then it must also be optimal to bike if it is sunny and walk if it is rainy or to walk if it is sunny and bike if it is rainy. This requires that the optimization contingent on sun is in some sense separable from the optimization contingent on rain. Translated into acts, dynamic programming solvability implies that if the feasible set contains $\{f, h, f_E h, h_E f\}$, and f and h are both unconditionally optimal, then $f_E h$ and $h_E f$ must also be unconditionally optimal. It is easy to see that this condition is not generally satisfied by MEU preferences. For example, suppose there are four states of the world, $S = \{1, 2, 3, 4\}$, $E = \{1, 2\}$, the set of measures $C = \{(\frac{1}{2}p, \frac{1}{2}(1-p), \frac{1}{2}p, \frac{1}{2}(1-p)) \mid p \in [0, 1]\}$, $Z = \mathbb{R}$, and $u(z) = z$. Letting $f = (1, 0, 0, 1)$ and $h = (\frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2})$ one may check that f and h are both optimal within $B = \text{co}\{f, h, f_E h, h_E f\}$, yet $f_E h = (1, 0, \frac{1}{2}, \frac{1}{2})$ is suboptimal, violating dynamic programming solvability. Thus, this condition leads to nonexistence of an update rule for MEU preferences. The same conclusion applies to a closely related axiom proposed by **Grant et al. (2000)**.

Compared to all of this literature, then, **DC** and **PFI** are weak requirements, and strengthening them in the direction of any of these proposals quickly leads to nonexistence of an update rule.

3.3.2 Consequentialism and reduction In addition to dynamic consistency, two other conditions important to justifications of updating in the decision theory literature are consequentialism and reduction of compound acts/decision trees. In our setting, consequentialism means that preference conditional on an event E depends only on the unconditional preference, the event E , and treats E^c as a null event.¹⁴ Reduction means

¹⁴The null event requirement is exactly our axiom **NC** defined in **Appendix A**.

that preferences are defined over acts, and thus the DM is assumed to care only about the mapping from states to (lotteries over) outcomes induced by his actions. It has been shown in various contexts that consequentialism, reduction, and some version of dynamic consistency (together with standard assumptions) imply expected utility and Bayesian updating (see e.g., Karni and Schmeidler 1991 and Ghirardato 2002). As a result, at least one of these properties must be relaxed if updated preferences are not to be of the expected utility form. Our approach maintains reduction but relaxes consequentialism. Two papers on updating multiple priors that relax reduction while maintaining consequentialism are Wang (2003) and Hayashi (2005).¹⁵

Epstein and Schneider (2003) have shown that if dynamic consistency is required only for a single filtration of events, then a subclass of MEU preferences is compatible with both consequentialism and reduction (and recursion) through the use of full Bayesian updating. This subclass requires the representing set of measures C to be *rectangular*, meaning it satisfies a type of stochastic independence with respect to events in the filtration. Formally, this property is defined as follows.

DEFINITION 13. Suppose $E, E^c \in \mathcal{N}(\mathcal{Z})$. The set C is *rectangular* with respect to $\{E, E^c\}$ if $\forall q^1, q^2, q^3 \in C, \exists q^4 \in C$ such that $\forall F, q^4(F) = q^3(E)q^1(F \cap E)/q^1(E) + q^3(E^c)q^2(F \cap E^c)/q^2(E^c)$.

For these preferences, when conditioning on an event in the filtration, each act is evaluated by minimizing separately the conditional expectation on E and on E^c and then minimizing the expectation of these conditional expectations. This leads to the following result.

PROPOSITION 10. Any of the ambiguity-maximizing update rules defined above that belong to \mathcal{U}^{Bayes} coincide with full Bayesian updating for $(\mathcal{Z}, E, g, B) \in \mathcal{T}$ such that $E, E^c \in \mathcal{N}(\mathcal{Z})$ and C is rectangular with respect to $\{E, E^c\}$.

This shows that these update rules agree with that of Epstein and Schneider's recursive multiple priors model on the domain of that model.¹⁶ Thus, when updating on events with respect to which preferences are rectangular, the ambiguity-maximizing rules identified by our approach satisfy the same properties (including consequentialism and recursion) as recursive multiple priors. At the same time, our rules are able to allow for updating on all other non-null events in a manner satisfying desirable axioms including **DC**, while the recursive multiple priors approach rules out and therefore does not address such events (e.g., the event that the ball drawn is not yellow in the dynamic Ellsberg example of the **Introduction**).

¹⁵Both of these papers relax reduction in a manner influenced by the seminal work of Kreps and Porteus (1978), which was done in the context of decision trees with objective probabilities.

¹⁶Note that Epstein and Schneider's framework is not identical to ours in that it places more temporal structure on the acts and outcomes than we do. For example, their ultimate outcomes are vectors of consumption over time while we have unstructured ultimate outcomes. Thus, literally speaking, with a rectangular set of measures, the two models are identical in their sets of measures and how they are updated, but not in all the structural details.

Given dynamic consistency for all non-null events, our focus on relaxing consequentialism rather than reduction has at least three justifications. First and most importantly, maintaining consequentialism (with or without reduction) rules out the Ellsberg pattern of choices over acts in the dynamic Ellsberg example of the **Introduction**. Second, relaxing reduction means that uncertainty is treated differently at different decision nodes, a property that may not be attractive in terms of a pure updating interpretation where the role of time is not, in and of itself, an important consideration. Third, viewing the current state of the literature, the only previous theory of updating MEU preferences while at the same time maintaining dynamic consistency with respect to updating on all non-null events is in **Wang (2003)**. As mentioned above, Wang's theory relaxes reduction. In light of this, one might have thought that relaxing reduction was the only way to maintain this scope of dynamic consistency. Thus, it makes sense for us to focus on relaxing consequentialism because there is more to be learned in that direction. Furthermore, there is an important sense in which relaxing reduction breaks completely the connection between dynamic consistency and the form of updating. In contrast, our approach of relaxing consequentialism maintains important and interesting links between the two, as we have shown.

3.3.3 Dynamic inconsistency There are approaches to dynamic decision making that take dynamic inconsistency of preferences over acts as given. Dynamic behavior is derived by coupling the underlying preferences with an assumption about how the different "selves" deal with the inconsistency. The best known approach here is due to **Strotz (1955–6)**, and assumes that behavior is *sophisticated* in the sense that the DM correctly anticipates her future conditional desires and chooses taking the results of those future decisions as a constraint. Observe that this is quite a distinct treatment of dynamic choice from the one put forward here, where conditional preferences agree with the optimal choices of the unconditional preference. **Siniscalchi (2006)** investigates the preferences over decision trees that are compatible with the Strotzian type of backward induction approach. As we pointed out in the **Introduction**, no such preferences can deliver behavior in the dynamic modification of Ellsberg's example that agrees with the typical preferences in Ellsberg's original problem (namely, $(1, 0, 0) \succ (0, 1, 0)$ and $(0, 1, 1) \succ (1, 0, 1)$). For a critique of the sophistication approach based on value of information considerations see e.g., **Epstein and Le Breton (1993, pp. 11–12)**. There are examples in the literature arguing in favor of (**Eichberger et al. 2007**) and against (**Grünwald and Halpern 2004** and **Seidenfeld 2004**) such preferences in the context of ambiguity.

Three alternative approaches that, like sophistication, accept that preferences over acts may be dynamically inconsistent, are **McClennen's (1990)** resolute choice, the cooperative decision processes of **Jaffray (1999)**, and the game-theoretic approach of **Peleg and Yaari (1973)**. In resolute choice, it is assumed that conditional choices are in agreement with an unconditionally optimal plan, even when those conditional choices may conflict with some underlying conditional preference. **McClennen (1990)** does not specify how this agreement is to be obtained. The resulting behavior is as if the conditional choices were governed by dynamically consistent underlying conditional preferences

that depend on earlier choices and the choice problem, such as those in our paper.¹⁷ In this sense, adopting one of our update rules satisfying **DC** is a way of implementing resolute choice for MEU preferences, while also preserving the property that conditional choices are based solely on conditional preferences. Jaffray (1999) suggests that the inconsistency between unconditional and conditional preferences might be resolved in a way that is more of a compromise between the different preferences. He examines a selection criterion that chooses a plan that is “not too bad” in a utility sense according to any of these preferences and is not dominated in that no feasible plan is better according to all the preferences. Peleg and Yaari (1973) propose treating the conflict generated by dynamic inconsistency as a non-cooperative game between the different “selves” of the DM and solving for an equilibrium.

Ozdenoren and Peck (forthcoming) use extensive form games of conflict with nature to illustrate how varying the game the DM thinks she is playing is an alternative modeling strategy for generating variation in conditional choices in Ellsberg-like problems. Tallon and Vergnaud (2006) show that it may be important in terms of dynamic consistency in the Ellsberg problem whether one updates on information about the color of the ball drawn or the color composition of the urn.

4. CONCLUSION

We have proposed and characterized novel update rules that apply to MEU preferences over acts. These preferences are important in analyzing dynamic behavior in the presence of ambiguity. This paper provides the first theories of ambiguity sensitive preferences that maintain, in dynamic extensions (as in our introductory example), typical behavior under ambiguity as exemplified by the classic Ellsberg problem. In particular, no theories based on recursion or backward induction may deliver this behavior, and neither can any theory that naively updates in a dynamically inconsistent manner.

The major new feature of the rules we characterize is that they are dynamically consistent when updating on any non-null event, and as a result, depend on the choice problem. In fact, we have shown that dependence on the feasible set of acts (and not simply the unconditionally optimal act) is necessary. Nonetheless, all of the rules take the simple form of applying Bayes’ rule to a subset of measures used in representing the DM’s unconditional preferences. In addition to deriving the entire set of rules compatible with our assumptions, we identify the unique ambiguity-maximizing rules, thus delimiting the extent to which dynamic consistency and other conditions we examine *force* the DM to reduce the set of measures she considers in the face of new information.

The dynamic consistency axioms we use, **DC** and **PFI**, are weaker than those appearing elsewhere, yet effectively capture the fundamental idea that optimal contingent plans should be respected when a planned-for contingency arises. As we discussed in Section 3.1, strengthening these axioms in various ways leads to non-existence.

While MEU is a well-known and important class of preferences for modeling behavior under ambiguity, it is far from the only such model. In ongoing work, we are

¹⁷Though McClennen of course had his own, stronger, notion of dynamic consistency, which we discussed above.

investigating the implications of similar axioms for other preference models, such as the smooth ambiguity model in [Klibanoff et al. \(2005\)](#).

A. APPENDIX: FURTHER PREFERENCE FOUNDATIONS

A.1 *A behavioral characterization of applying Bayes' rule to part of the set of priors*

In this appendix, we present four axioms that, collectively, are shown to characterize the restriction to update rules that work by applying Bayes' rule to some of the measures in the representing set of measures, C , i.e., the restriction to rules in \mathcal{U}^{Bayes} . Thus, these axioms, when combined with the dynamic consistency conditions introduced above, complete the behavioral characterization of the update rules proposed in this paper.

The first axiom requires that since the unconditional preferences that we consider are MEU, conditional preferences also should satisfy the axioms characterizing MEU ([Gilboa and Schmeidler 1989](#)). There is little reason to imagine different families of preferences at different stages of a decision and good reason to think that properties reasonable for a preference relation are reasonable also for its updates. In particular, the attractiveness of a given set of axioms characterizing a class of preferences is not usually thought to depend on whether the preferences at hand are conditional or unconditional preferences.

AXIOM CL (Closure with respect to \mathcal{P}^{MEU}). *For any $(\succsim, E, g, B) \in \mathcal{T}$, $\succsim_{E,g,B} \in \mathcal{P}^{MEU}$.*

The next axiom states that a preference conditional on an event E should not depend on the consequences outside E , a basic requirement on conditional preferences.

AXIOM NC (Null Complement). *For any $(\succsim, E, g, B) \in \mathcal{T}$ and $f, h \in \mathcal{A}$, $f \sim_{E,g,B} f_E h$.*

The next axiom requires updating to preserve the ordering of constant acts (acts yielding the same lottery in each state of the world). In particular this yields a separation between attitudes towards risk, which are held constant by this unchanged ordering, and the remainder of the preference, which may be affected by updating.

AXIOM UT (Unchanged Tastes). *For any $(\succsim, E, g, B) \in \mathcal{T}$ and $x, y \in X$, $x \succsim y \Leftrightarrow x \succsim_{E,g,B} y$.*

The straightforward proposition below states that these first three axioms are equivalent to the conditional preferences having an MEU representation using the same u as the unconditional preferences and all measures in the updated set $C_{E,g,B}$ placing zero weight on E^c .

PROPOSITION 11. *The set \mathcal{U} ([Definition 1](#)) is the set of all update rules satisfying **CL**, **NC**, and **UT**.*

Notice that the axioms so far require no connection between the unconditional measures in C and the conditional measures in $C_{E,g,B}$. Our next axiom generates such a connection. A stronger version of this axiom appears in [Jaffray \(1994\)](#) and is used by [Pires \(2002\)](#) to characterize full Bayesian updating for MEU.

AXIOM **IIW** (Information Improves the Worst-case). For any $(\succsim, E, g, B) \in \mathcal{T}$, $f \in \mathcal{A}$, and $x \in X$, if $f_E x \sim x$ then $f \succsim_{E,g,B} x$.

To understand this axiom, note that $\min_{q \in C} \int (u \circ f) dq_E$ is the worst-case evaluation of the act f on the event E according to the measures in C . If $f_E x \sim x$, then it must be that this worst-case evaluation equals $u(x)$. Similarly, $f \succsim_{E,g,B} x$ requires $\min_{q \in C_{E,g,B}} \int (u \circ f) dq \geq u(x)$. Together these imply

$$\min_{q \in C_{E,g,B}} \int (u \circ f) dq \geq \min_{q \in C} \int (u \circ f) dq_E.$$

Thus, the axiom says that learning that E has occurred may not lower the worst-case evaluation of any act on E . Why does this make sense? We usually think of information as reducing ambiguity. If we think of the ex ante ambiguity in the evaluation of an act f on E as given by the interval $[\min_{q \in C} \int (u \circ f) dq_E, \max_{q \in C} \int (u \circ f) dq_E]$ and ask that this ambiguity weakly shrink with updating, in the sense that

$$\left[\min_{q \in C_{E,g,B}} \int (u \circ f) dq, \max_{q \in C_{E,g,B}} \int (u \circ f) dq \right] \subseteq \left[\min_{q \in C} \int (u \circ f) dq_E, \max_{q \in C} \int (u \circ f) dq_E \right],$$

then this implies $\min_{q \in C_{E,g,B}} \int (u \circ f) dq \geq \min_{q \in C} \int (u \circ f) dq_E$, which in turn yields **IIW**.

We now show that, together with the earlier axioms, **IIW** is equivalent to the updated sets of measures $C_{E,g,B}$ being generated only from Bayesian conditionals of measures in C . In other words, no new sets of relative weights on E may suddenly appear in the conditional set of measures.

PROPOSITION 12. \mathcal{U}^{Bayes} is the set of all update rules satisfying **CL**, **NC**, **UT**, and **IIW**.

PROOF. Observe that if $f_E x \sim x$, then $\min_{q \in C} \int (u \circ f) dq_E = u(x)$. If

$$C_{E,g,B} \subseteq \{p_E \mid p \in C\}$$

then

$$\min_{q \in C_{E,g,B}} \int (u \circ f) dq \geq \min_{q \in \{p_E \mid p \in C\}} \int (u \circ f) dq = \min_{q \in C} \int (u \circ f) dq_E = u(x)$$

and **IIW** is satisfied. If $C_{E,g,B}$ contains a measure not in $\{q_E \mid q \in C\}$ then by separating hyperplane arguments there exists a $q^* \in C_{E,g,B}$ and an act f such that

$$\int (u \circ f) dq^* < \min_{q \in \{p_E \mid p \in C\}} \int (u \circ f) dq = u(x)$$

and thus $f \prec_{E,g,B} x$, in violation of **IIW**. This proves the result. □

This result justifies restricting attention to update rules that apply Bayes' rule to some of the measures in C . However, it is worth noting that, with the exception of

Proposition 10, all of the characterization results in the paper may be easily modified to ones that hold when **IIW** is dropped. Specifically, wherever the restriction $q \in C$ appears, replacing this with $q \in \Delta(S)$ such that $q(E) > 0$ and replacing $U \in \mathcal{U}^{Bayes}$ with $U \in \mathcal{U}$ gives analogous results that hold without **IIW**. We choose to present the results with **IIW** rather than without because we find the restriction to measures in C natural for updating. Furthermore, **IIW** helps in relating our rules to those in the literature (e.g., **Proposition 10**) and is satisfied by all previously proposed update rules for MEU preferences. For example, strengthening **IIW** by replacing $f \succsim_{E,g,B} x$ with $f \sim_{E,g,B} x$ (and thus requiring that information leaves the worst-case unchanged) is equivalent to requiring full Bayesian updating. The lemma below is a slight variation of **Pires'** (2002) axiomatic characterization of full Bayesian updating. Whereas Pires uses an axiom (due to **Jaffray 1994**) that runs from the conditional to the unconditional, i.e., $f \sim_{E,g,B} x$ implies $f_E x \sim x$, ours runs in the other direction, as in **IIW**.

LEMMA 2. *Full Bayesian updating (U^{FB}) is the unique update rule satisfying **CL**, **NC**, **UT**, and **IIW** with $f \succsim_{E,g,B} x$ replaced by $f \sim_{E,g,B} x$.*

PROOF. **CL**, **NC** and **UT** are clearly satisfied by U^{FB} . Observe that $f_E x \sim x$ if and only if $\min_{q \in C} \int (u \circ f) dq_E = u(x)$. Under U^{FB} , f is evaluated conditionally according to $\min_{q \in C} \int (u \circ f) dq_E$, and therefore $f_E x \sim x$ if and only if $f \sim_{E,g,B} x$. For the other direction, assume **CL**, **NC**, **UT**, and **IIW** with $f \succsim_{E,g,B} x$ replaced by $f \sim_{E,g,B} x$ hold. This implies that each f must be evaluated conditionally as $\min_{q \in C} \int (u \circ f) dq_E$. This can happen only if $C_{E,g,B}$ is the set of all conditionals that can be formed from C (see **Pires 2002**). This is U^{FB} . \square

From the **Introduction** we know that full Bayesian updating is dynamically inconsistent, and thus the lemma shows that this stronger version of **IIW** cannot be imposed along with dynamic consistency.

A.2 Observability of conditional preferences

In any theory including both conditional and unconditional preferences, observability of these preferences may be an issue. Specifically, observing conditional preference requires, at a minimum, observability of the conditioning event relevant to the DM. In a theory such as ours, where conditional preference may depend not only on the conditioning event, but also on the unconditionally chosen act, g , and the feasible set, B , from which it was chosen, one needs to assume that these features are observable as well. In the context of the dynamic Ellsberg example, for instance, one needs to observe the tree within which the choice of betting on black versus red is made to properly interpret the conditional preference to which this choice belongs.

Unconditional preferences, \succsim , can in principle be observed, as usual, by having a DM make choices among acts and paying her according to the acts she chooses and the realized states of the world. Having elicited \succsim in this manner, one may then face the DM with a given feasible set B and observe her choice, g , from this set. Then, to elicit conditional preference, $\succsim_{E,g,B}$, the same problem is presented but now the DM is

told that if the state of the world lies in E she will be given an opportunity to revise her initial choice, and then, assuming the state of the world is in E , asked to make choices among acts at that point. Notice that eliciting the conditional preference between f and h where neither is conditionally optimal within B requires offering a different choice set at E than was specified in B . Just as the preference \succsim does not depend on the choice set to which it is applied, it seems quite natural (and in the spirit of imposing the same requirements on conditional preferences as on unconditional) to assume the same for $\succsim_{E,g,B}$. In the absence of this assumption, one could only hope to elicit conditionally optimal choices from B , rather than the whole conditional preference relation.

B. APPENDIX: SOME RESULTS AND SELECTED PROOFS

B.1 Proofs of results from Section 2

We begin with two key lemmata. The first is used in several of the proofs to come and shows that there is at least one measure in C that both supports the conditional optimality of g and is used to evaluate g unconditionally. The second lemma is used in proving the first, and relates measures used to evaluate g unconditionally to measures separating g from acts weakly better than g . The convexity of the feasible set, B , and of the set of acts weakly better than g is crucial in these arguments.

LEMMA 3. For $(\succsim, E, g, B) \in \mathcal{T}$, $Q^{E,g,B} \cap \operatorname{argmin}_{p \in C} \int (u \circ g) dp \neq \emptyset$.

PROOF. Let I be the set of Σ -measurable, bounded real-valued functions on S . Let \succsim be a complete, transitive binary relation on I defined by

$$a \succsim b \text{ if and only if } \min_{p \in C} \int a dp \geq \min_{p \in C} \int b dp.$$

Let \succ and \approx be the asymmetric and symmetric parts of \succsim . Consider the convex sets $D_1 \equiv \{a \mid a \in I \text{ with } a \succ u \circ g\}$ and $D_2 \equiv \{u \circ f \mid f \in B \text{ with } f = g \text{ on } E^c\}$. Unconditional optimality of g implies $D_1 \cap D_2 = \emptyset$. The set D_1 is non-empty by inspection and also has a non-empty interior. The set D_2 is non-empty since it contains $u \circ g$. By a separating hyperplane theorem (e.g., Aliprantis and Border 1999, Thm. 5.50, p. 190), there exists a hyperplane separating D_1 and D_2 . Without loss of generality, such a hyperplane may be defined by $\{a \in I \mid \int a dr = \alpha\}$ for a finitely additive measure $r \in \Delta(S)$ and real α such that $\int a dr \geq \alpha \geq \int b dr$ for all $a \in D_1$ and $b \in D_2$. Since $u \circ g \in D_2$, $\alpha \geq \int (u \circ g) dr$. Suppose $\alpha > \int (u \circ g) dr$. Then by the event-wise continuity of $\int (\cdot) dr$ and monotonicity of \succsim , there exists an $a \in D_1$ such that $\int a dr < \alpha$, a contradiction. Thus, $\alpha = \int (u \circ g) dr$. Therefore, $\int (u \circ g) dr \geq \int b dr$ for all $b \in D_2$. The same argument by contradiction shows $a \approx u \circ g$ implies $\int a dr \geq \alpha$ and thus $[a \approx u \circ g \implies \int a dr \geq \int (u \circ g) dr]$ for all $a \in I$. Now, by the argument in the “only if” direction of the proof of Lemma 4 applied to \succsim , $r \in \operatorname{argmin}_{p \in C} \int (u \circ g) dp$. From the definition of $Q^{E,g,B}$ and the fact that $\int (u \circ g) dr \geq \int b dr$ for all $b \in D_2$, $r \in Q^{E,g,B}$. Thus, $r \in Q^{E,g,B} \cap \operatorname{argmin}_{p \in C} \int (u \circ g) dp \neq \emptyset$. \square

LEMMA 4. *Suppose there is no best or worst consequence in Z . Then $q \in \Delta(S)$ satisfies $[f \sim h \implies \int (u \circ f) dq \geq \int (u \circ h) dq]$ for all $f \in \mathcal{A}$ if and only if $q \in \operatorname{argmin}_{p \in C} \int (u \circ h) dp$.*

PROOF. (If) Suppose $q \in \operatorname{argmin}_{p \in C} \int (u \circ h) dp$. Then $f \sim h \implies \int (u \circ f) dq \geq \min_{p \in C} \int (u \circ f) dp = \min_{p \in C} \int (u \circ h) dp = \int (u \circ h) dq$.

(Only if) Let q be such that $[f \sim h \implies \int (u \circ f) dq \geq \int (u \circ h) dq]$ for all $f \in \mathcal{A}$. Let $u(x) = \min_{p \in C} \int (u \circ h) dp$. Since $x \sim h$, $u(x) \geq \int (u \circ h) dq$. Suppose $u(x) > \int (u \circ h) dq$. Consider $h' \in \mathcal{A}$ such that $h = \alpha h' + (1 - \alpha)x$ for some $\alpha \in (0, 1)$. Such an h' exists because the assumption of no best or worst consequence in Z ensures the existence of an open neighborhood of h (using the norm defined by the supremum over state utility differences). Observe that $h' \sim x \sim h$ but $\int (u \circ h') dq < \int (u \circ h) dq$, contradicting $[h' \sim h \implies \int (u \circ h') dq \geq \int (u \circ h) dq]$. Therefore, $u(x) = \int (u \circ h) dq = \min_{p \in C} \int (u \circ h) dp$. It remains to show that $q \in C$. For all $f \sim h$, $\int (u \circ f) dq \geq \int (u \circ h) dq = u(x) = \min_{p \in C} \int (u \circ h) dp = \min_{p \in C} \int (u \circ f) dp$. Since the C -linearity of MEU preferences implies that all indifference curves have the same shape, the same holds for any $f \in \mathcal{A}$, i.e., $\int (u \circ f) dq \geq \min_{p \in C} \int (u \circ f) dp$. Since $C = \{q \in \Delta(S) \mid \int (u \circ f) dq \geq \min_{p \in C} \int (u \circ f) dp \text{ for all } f \in \mathcal{A}\}$ (see, e.g., the proof of the Fundamental lemma in [Chateauneuf 1991](#)), we have shown $q \in C$ and thus $q \in \operatorname{argmin}_{p \in C} \int (u \circ h) dp$. \square

We can now prove the results stated in [Section 2](#).

PROOF OF [PROPOSITION 1](#). First we show that all rules in \mathcal{W}^{DC} satisfy **DC**. Let q_E^g be an element of

$$Q_E^{E,g,B} \cap \operatorname{argmin}_{q \in C_{E,g,B}} \int (u \circ g) dq.$$

Since

$$\min_{q \in C_{E,g,B}} \int (u \circ g) dq = \int (u \circ g) dq_{q_E^g} \geq \int (u \circ f) dq_{q_E^g} \geq \min_{q \in C_{E,g,B}} \int (u \circ f) dq$$

for all $f \in B$ with $f = g$ on E^c , g is conditionally optimal and thus **DC** holds.

Now we show that no rules in \mathcal{W}^{Bayes} that are not in \mathcal{W}^{DC} can satisfy **DC**. Consider the convex sets $D_1 \equiv \{f \mid f \succ_{E,g,B} g \text{ with } f = g \text{ on } E^c\}$ and $D_2 \equiv \{f \mid f \in B \text{ with } f = g \text{ on } E^c\}$. **DC** is equivalent to $D_1 \cap D_2 = \emptyset$. An argument very similar to the one in the proof of [Lemma 3](#), with $\succ_{E,g,B}$ in place of \succ and the measure r corresponding to the separating hyperplane in $\Delta(E)$ rather than $\Delta(S)$, may be used to show that $D_1 \cap D_2 = \emptyset$ (and thus **DC**) implies

$$Q_E^{E,g,B} \cap \operatorname{argmin}_{q \in C_{E,g,B}} \int (u \circ g) dq \neq \emptyset.$$

Therefore, any rules in \mathcal{W}^{Bayes} satisfying **DC** must be in \mathcal{W}^{DC} . \square

PROOF OF [COROLLARY 2](#). By the argument in the proof of [Lemma 3](#), since there is no kink on E in the feasible set at g , there exists a unique tangent hyperplane to B on

E at g . Thus, $Q_E^{E,g,B}$ consists of the unique measure $q_E \in \Delta(E)$ associated with that hyperplane. □

PROOF OF LEMMA 1. We have $f \succsim^i x \iff \min_{p \in C^i} \int (u \circ f) dp \geq u(x)$. Now, $C^2 \subseteq C^1$ implies $\min_{p \in C^2} \int (u \circ f) dp \geq \min_{p \in C^1} \int (u \circ f) dp$ and thus $f \succsim^1 x$ implies $f \succsim^2 x$ and $f \succ^1 x \implies f \succ^2 x$. Suppose $C^2 \subset C^1$. Then there exist an f and an x such that $\min_{p \in C^2} \int (u \circ f) dp \geq u(x) > \min_{p \in C^1} \int (u \circ f) dp$ and thus $f \succsim^2 x \not\Rightarrow f \succsim^1 x$. □

PROOF OF PROPOSITION 2. To show U^{DCmax} exists, it suffices to show that

$$\min_{q \in Q_E^{E,g,B}} \int (u \circ g) dq$$

always exists, as $C_{E,g,B}$ is then non-empty. If $Q_E^{E,g,B}$ is non-empty and (weak*-)compact then $\min_{q \in Q_E^{E,g,B}} \int (u \circ g) dq$ exists.

That $Q_E^{E,g,B}$ (and thus $C_{E,g,B}$) is non-empty follows from Lemma 3. We now show that $Q_E^{E,g,B}$ is compact. Let

$$M_E = \left\{ q_E \in \Delta(E) \mid \int (u \circ g) dq_E \geq \int (u \circ f) dq_E \text{ for all } f \in B \text{ with } f = g \text{ on } E^c \right\}.$$

By inspection M_E is closed as weak*-convergence ensures convergence of $\int (u \circ g) dq_E - \int (u \circ f) dq_E$ and the set of all conditionals on E is closed. We have $Q_E^{E,g,B} = \{q_E \mid q \in C\} \cap M_E$, the intersection of two closed sets, and so this set is closed. It is bounded because it is a set of measures. Thus it is compact. So, $\min_{q \in Q_E^{E,g,B}} \int (u \circ g) dq$ exists and thus so does U^{DCmax} .

Since

$$\emptyset \neq \operatorname{argmin}_{q \in Q_E^{E,g,B}} \int (u \circ g) dq \subseteq \operatorname{argmin}_{q \in C_{E,g,B}} \int (u \circ g) dq,$$

$U^{DCmax} \in \mathcal{U}^{DC}$. Since

$$Q_E^{E,g,B} \cap \operatorname{argmin}_{q \in C_{E,g,B}} \int (u \circ g) dq \neq \emptyset$$

implies that only $q \in C$ such that

$$\int (u \circ g) dq_E \geq \min_{q \in Q_E^{E,g,B}} \int (u \circ g) dq$$

may have their conditionals included in any $\hat{C}_{E,g,B}$ associated with rules in \mathcal{U}^{DC} , Lemma 1 implies U^{DCmax} is the unique ambiguity-maximizing rule in \mathcal{U}^{DC} . □

B.2 Proofs and results from Section 3

PROOF OF PROPOSITION 4. By inspection, $\mathcal{U}^{DC \cap PFI} \subseteq \mathcal{U}^{DC}$ and so these rules satisfy **DC**. We now show they satisfy **PFI**. Fix $f \in B$ with $f = g$ on E^c and $f \sim g$. Since $r \in Q^{E,g,B}$, $\int (u \circ f) dr_E \leq \int (u \circ g) dr_E$. By the definition of $\mathcal{U}^{DC \cap PFI}$ (Definition 8), $\int (u \circ f) dr_E \geq \int (u \circ g) dr_E$. Thus, $\int (u \circ f) dr_E = \int (u \circ g) dr_E$ and

$$\int (u \circ f) dr_E = \min_{p \in C_{E,g,B}} \int (u \circ f) dp.$$

Setting $f = g$ in the definition of $C_{E,g,B}$ gives

$$\int (u \circ g) dr_E = \min_{p \in C_{E,g,B}} \int (u \circ g) dp.$$

Therefore, $f \sim_{E,g,B} g$ for all $f \in B$ with $f = g$ on E^c and $f \sim g$ and **PFI** is satisfied.

Consider an update rule $U \in \mathcal{U}^{Bayes}$ satisfying the axioms, with associated sets of measures $C_{E,g,B}^U$. Since U satisfies **DC**, Corollary 1 implies that there exists an $r \in Q^{E,g,B}$ such that $r_E \in C_{E,g,B}^U \subseteq \{q_E \mid q \in C \text{ and } \int (u \circ g) dq_E \geq \int (u \circ g) dr_E\}$. For any $f \in B$ with $f = g$ on E^c and $f \sim g$, **PFI** implies

$$\min_{p \in C_{E,g,B}^U} \int (u \circ f) dp = \min_{p \in C_{E,g,B}^U} \int (u \circ g) dp.$$

Therefore, it implies also that for any $q_E \in C_{E,g,B}^U$,

$$\int (u \circ f) dq_E \geq \min_{p \in C_{E,g,B}^U} \int (u \circ g) dp$$

for all $f \in B$ with $f = g$ on E^c and $f \sim g$. This shows that $C_{E,g,B}$ has the form stated in $\mathcal{U}^{DC \cap PFI}$. \square

LEMMA 5. For $(\succsim, E, g, B) \in \mathcal{T}$, $Q^{E,g,B} \cap \arg \min_{p \in C} \int (u \circ g) dp \subseteq K^{E,g,B} \subseteq Q^{E,g,B}$.

PROOF. Suppose $q \in Q^{E,g,B} \cap \arg \min_{p \in C} \int (u \circ g) dp$. Consider any $f \in B$ with $f = g$ on E^c . By definition, $q \in Q^{E,g,B}$ implies $\int (u \circ g) dq \geq \int (u \circ f) dq$. If $f \sim g$, then $\int (u \circ f) dq \geq \min_{p \in C} \int (u \circ f) dp = \min_{p \in C} \int (u \circ g) dp = \int (u \circ g) dq$ and so $\int (u \circ g) dq = \int (u \circ f) dq$. Therefore, $q \in K^{E,g,B}$. Since the definition of $K^{E,g,B}$ includes more restrictions than that of $Q^{E,g,B}$, $K^{E,g,B} \subseteq Q^{E,g,B}$. \square

PROOF OF PROPOSITION 5. First we show that $U^{DC \cap PFI \max}$ exists and $U^{DC \cap PFI \max} \in \mathcal{U}^{DC \cap PFI}$. Then we show it is uniquely ambiguity maximizing within $\mathcal{U}^{DC \cap PFI}$. We begin proving existence by noting that Lemmata 3 and 5 imply $K^{E,g,B}$ (and thus $K_E^{E,g,B}$) is non-empty. Since

$$K_E^{E,g,B} = Q_E^{E,g,B} \cap \left\{ q_E \in \Delta(E) \mid \int (u \circ f) dq_E = \int (u \circ g) dq_E \text{ for all } f \in B \text{ with } f = g \text{ on } E^c \text{ and } f \sim g \right\}$$

and $Q_E^{E,g,B}$ is shown to be closed in the proof of **Proposition 2**, $K_E^{E,g,B}$ is the intersection of two closed sets, and so is closed. It is bounded because it is a set of measures. Thus it is compact. So,

$$\min_{q \in K_E^{E,g,B}} \int (u \circ g) dq$$

exists. By **Lemma 5**, $K_E^{E,g,B} \subseteq Q_E^{E,g,B}$ and so

$$\operatorname{argmin}_{p \in K_E^{E,g,B}} \int (u \circ g) dp \subseteq Q_E^{E,g,B}.$$

By inspection of $U^{DC \cap PFI \max}$,

$$\operatorname{argmin}_{p \in K_E^{E,g,B}} \int (u \circ g) dp \subseteq C_{E,g,B}.$$

Thus, taking

$$r_E \in \operatorname{argmin}_{p \in K_E^{E,g,B}} \int (u \circ g) dp$$

and an associated unconditional $r \in K_E^{E,g,B}$, we see that $C_{E,g,B}$ is non-empty and $U^{DC \cap PFI \max} \in \mathcal{U}^{DC \cap PFI}$. Given r_E , it is obvious that any $\hat{C}_{E,g,B}$ associated with an update rule in $\mathcal{U}^{DC \cap PFI}$ is made largest by replacing \subseteq with $=$ in the definition of $\mathcal{U}^{DC \cap PFI}$. It remains to show only that such a choice of r_E leads to the smallest value of

$$\min_{p \in \hat{C}_{E,g,B}} \int (u \circ g) dp$$

for rules in $\mathcal{U}^{DC \cap PFI}$. Suppose there exists a $\hat{r} \in Q_E^{E,g,B}$ such that

$$\int (u \circ g) d\hat{r}_E < \min_{p \in K_E^{E,g,B}} \int (u \circ g) dp.$$

Without loss of generality, assume

$$\min_{p \in \hat{C}_{E,g,B}} \int (u \circ g) dp = \int (u \circ g) d\hat{r}_E.$$

If and only if such an $\hat{r}_E \in \hat{C}_{E,g,B}$ is there a rule in $\mathcal{U}^{DC \cap PFI}$ displaying more ambiguity than $U^{DC \cap PFI \max}$. However, if $\hat{r}_E \in \hat{C}_{E,g,B}$, then

$$\int (u \circ f) d\hat{r}_E \geq \int (u \circ g) d\hat{r}_E = \min_{p \in \hat{C}_{E,g,B}} \int (u \circ g) dp$$

for all $f \in B$ with $f = g$ on E^c and $f \sim g$. This, together with $\hat{r} \in Q_E^{E,g,B}$, implies $\hat{r} \in K_E^{E,g,B}$, a contradiction. This shows $C_{E,g,B}$ is the maximal set of measures among all sets associated with rules in $\mathcal{U}^{DC \cap PFI}$, and so, by **Lemma 1**, $U^{DC \cap PFI \max}$ is uniquely ambiguity maximizing. \square

AXIOM **PFI1**. For any $(\succsim, E, g, B) \in \mathcal{T}$, if $f, h \in B$ with $f = h = g$ on E^c and $f \sim h$, then $f \sim_{E,g,B} h$.

AXIOM **PFI2**. For any $(\succsim, E, g, B) \in \mathcal{T}$, if $f \in \mathcal{A}$ with $f = g$ on E^c and $f \sim g$, then $f \sim_{E,g,B} g$.

AXIOM **PFI3**. For any $(\succsim, E, g, B) \in \mathcal{T}$, if $f \in B$ with $f = g$ on E^c , $f \sim g$ if and only if $f \sim_{E,g,B} g$.

PROPOSITION 13. No update rule in \mathcal{U} satisfies **PFI1**, **PFI2**, or **PFI3**.

PROOF. Identical to the proof of Proposition 3 with **PFI** n replacing **DC** n , $n = 1, 2, 3$. \square

PROOF OF PROPOSITION 6. Applying Bayes' rule to the DM's subjective probability, p , satisfies the axioms. We prove it does so uniquely. By Proposition 1 and the argument in Appendix A describing how to modify our characterization results when **IHW** is not imposed, **DC** requires the updated probability, $p_{E,g,B}$, to belong to $\tilde{Q}_E^{E,g,B}$ where

$$\tilde{Q}_E^{E,g,B} \equiv \left\{ q_E \mid q \in \Delta(S), q(E) > 0, \int (u \circ g) dq \geq \int (u \circ f) dq \text{ for all } f \in B \text{ with } f = g \text{ on } E^c \right\}.$$

By inspection, the Bayesian update, p_E , is in $\tilde{Q}_E^{E,g,B}$. By **IFS**, without loss of generality we can assume the feasible set is smooth at g . This implies $\tilde{Q}_E^{E,g,B}$ is a singleton by arguments as in the proof of Corollary 2. Thus $p_{E,g,B} = p_E$ and updating is Bayesian. \square

PROOF OF PROPOSITION 7. Consider the same state space, information structure, and preferences as in the dynamic Ellsberg example. So, \succsim are MEU preferences with $u(z) = z$ and $C = \{(\frac{1}{3}, \alpha, \frac{2}{3} - \alpha) \mid \alpha \in [\frac{1}{4}, \frac{5}{12}]\}$. The conditioning event is $E = \{Black, Red\}$.

For any update rule in \mathcal{U} , $\succsim_{E,g,B}$ can be represented using the same u and using a set of measures contained in $\Delta(E)$. Thus all acts that give a weakly higher payoff if a black ball is drawn than if a red ball is drawn should be evaluated using the measure in the updated set of measures that puts the least weight on black. Denote this measure by $(\beta, 1 - \beta, 0)$. **IFS** implies that, fixing g , β is the same for all B . We show that the existence of such a β contradicts **DC**.

Fix feasible sets $B_1 = co\{(\frac{3}{2}, \frac{3}{2}, 1), (3, 0, 1)\}$ and $B_2 = co\{(\frac{14}{9}, \frac{14}{9}, 1), (\frac{22}{9}, \frac{4}{9}, 1)\}$. One may verify that $g = (2, 1, 1)$ is the unique unconditionally optimal choice in B_1 and in B_2 . Note that $\beta < \frac{5}{9}$ implies $(\frac{14}{9}, \frac{14}{9}, 1) \succ_{E,g,B_2} g$ and $\beta > \frac{1}{2}$ implies $(3, 0, 1) \succ_{E,g,B_1} g$, both violations of **DC**. Therefore no update rule satisfying the axioms exists. \square

PROOF OF PROPOSITION 8. To see that **IDFS** is satisfied for U^{DCmax} , recall that the only place the feasible set B enters the definition of the rule (Definition 7) is through the

set $Q^{E,g,B}$ (Definition 3), the measures in C supporting the conditional optimality of g . Only acts agreeing with g on E^c matter in determining $Q^{E,g,B}$, and adding or removing acts dominated on E among these does not affect which measures make g conditionally optimal. For $U^{DC \cap PFI \max}$, the set $K_E^{E,g,B}$ (Definition 9) is similarly unaffected, as is the set of feasible acts such that $f \sim g$ and $f = g$ on E^c . \square

The proof of Proposition 9 makes use of the following lemma.

LEMMA 6. Fix any (\succsim, E, B) . There exists a $\nu \in C$ such that, for all acts f with $(\succsim, E, f, B) \in \mathcal{T}$, $\nu \in \operatorname{argmin}_{p \in C} \int (u \circ f) dp$.

PROOF. Let $(\succsim, E, g, B) \in \mathcal{T}$. By an argument similar to the proof of Lemma 3, there exists $\nu \in \operatorname{argmin}_{p \in C} \int (u \circ g) dp$ such that $\int (u \circ g) d\nu \geq \int (u \circ f) d\nu$ for all $f \in B$. Therefore, $f \sim g$ implies $\int (u \circ g) d\nu = \int (u \circ f) d\nu = \min_{p \in C} \int (u \circ f) dp$. \square

PROOF OF PROPOSITION 9. Consider $\nu \in C$ as identified in Lemma 6. From the proof of that lemma and the definition of $K_E^{E,g,B}$, it follows that $\nu_E \in C_{E,B}$. Therefore $C_{E,B}$ is non-empty. It is closed and convex since it is the intersection of closed convex sets. Thus $U^{DC \cap IC \max}$ exists. We have $U^{DC \cap IC \max} \in \mathcal{U}^{DC \cap PFI}$ because for each $g, \nu \in K_E^{E,g,B} \subseteq Q^{E,g,B}$ and $C_{E,B} \subseteq \{q_E \mid q \in C \text{ and } \int (u \circ f) dq_E \geq \int (u \circ g) d\nu_E \text{ for all } f \in B \text{ with } f = g \text{ on } E^c \text{ and } f \sim g\}$. Now, suppose that for some E and B a rule in $\mathcal{U}^{DC \cap PFI}$ uses a set $\hat{C}_{E,B}$, independent of g , that contains measures not in $C_{E,B}$. Then, by the definition of $C_{E,B}$, there is some $\hat{q} \in \hat{C}_{E,B}$ and some g such that $\hat{q} \notin C_{E,g,B}^{U^{DC \cap PFI \max}}$. But this implies that **DC** or **PFI** would be violated, because $U^{DC \cap PFI \max}$ is the ambiguity-maximizing rule in $\mathcal{U}^{DC \cap PFI}$. This proves that $U^{DC \cap IC \max}$ is the unique ambiguity-maximizing rule in $\mathcal{U}^{DC \cap PFI}$ satisfying **ICA**. \square

PROOF OF PROPOSITION 10. For $E, E^c \in \mathcal{N}(\succsim)$, rectangularity with respect to $\{E, E^c\}$ implies that if $f = h$ on E^c then

$$\min_{p \in C} \int (u \circ h) dp \geq \min_{p \in C} \int (u \circ f) dp$$

if and only if

$$\min_{p \in \{q_E \mid q \in C\}} \int (u \circ h) dp \geq \min_{p \in \{q_E \mid q \in C\}} \int (u \circ f) dp.$$

Therefore U^{FB} satisfies both **DC** and **PFI** under rectangularity. This together with the fact that U^{FB} is the unique ambiguity-maximizing rule in \mathcal{U}^{Bayes} proves the result. \square

REFERENCES

Aliprantis, Charalambos D. and Kim C. Border (1999), *Infinite Dimensional Analysis: A Hitchhiker's Guide*, second edition. Springer, Berlin. [288]

Anscombe, Frank J. and Robert J. Aumann (1963), "A definition of subjective probability." *Annals of Mathematical Statistics*, 34, 199–205. [266]

Chateauneuf, Alain (1991), "On the use of capacities in modeling uncertainty aversion and risk aversion." *Journal of Mathematical Economics*, 20, 343–369. [289]

Dempster, Arthur P. (1968), "A generalization of bayesian inference." *Journal of the Royal Statistical Society, Series B (Methodological)*, 30, 205–247. [264]

Eichberger, Jürgen, Simon Grant, and David Kelsey (2005), "CEU preferences and dynamic consistency." *Mathematical Social Sciences*, 49, 143–151. [264]

Eichberger, Jürgen, Simon Grant, and David Kelsey (2007), "Updating Choquet beliefs." *Journal of Mathematical Economics*, 43, 888–899. [283]

Ellsberg, Daniel (1961), "Risk, ambiguity, and the Savage axioms." *Quarterly Journal of Economics*, 75, 643–669. [262]

Epstein, Larry G. (1999), "A definition of uncertainty aversion." *Review of Economic Studies*, 66, 579–608. [271]

Epstein, Larry G. and Michel Le Breton (1993), "Dynamically consistent beliefs must be Bayesian." *Journal of Economic Theory*, 61, 1–22. [262, 273, 277, 280, 283]

Epstein, Larry G. and Massimo Marinacci (forthcoming), "Mutual absolute continuity of multiple priors." *Journal of Economic Theory*. [266]

Epstein, Larry G. and Martin Schneider (2003), "Recursive multiple-priors." *Journal of Economic Theory*, 113, 1–31. [262, 263, 264, 273, 282]

Fagin, Ronald and Joseph Y. Halpern (1991), "A new approach to updating beliefs." In *Uncertainty in Artificial Intelligence 6* (Piero P. Bonissone, Max Henrion, Laveen N. Kanal, and John F. Lemmer, eds.), 347–374, North-Holland, Amsterdam. [264]

Ghirardato, Paolo (2002), "Revisiting Savage in a conditional world." *Economic Theory*, 20, 83–92. [262, 273, 277, 282]

Ghirardato, Paolo, Fabio Maccheroni, and Massimo Marinacci (2004), "Differentiating ambiguity and ambiguity attitude." *Journal of Economic Theory*, 118, 133–173. [271]

Ghirardato, Paolo and Massimo Marinacci (2002), "Ambiguity made precise: A comparative foundation." *Journal of Economic Theory*, 102, 251–289. [271]

Gilboa, Itzhak and David Schmeidler (1989), "Maxmin expected utility with non-unique prior." *Journal of Mathematical Economics*, 18, 141–153. [262, 266, 285]

Gilboa, Itzhak and David Schmeidler (1993), "Updating ambiguous beliefs." *Journal of Economic Theory*, 59, 33–49. [262, 264, 279]

Grant, Simon, Atsushi Kajii, and Ben Polak (2000), "Decomposable choice under uncertainty." *Journal of Economic Theory*, 92, 169–197. [281]

Grünwald, Peter D. and Joseph Y. Halpern (2004), "When ignorance is bliss." In *Uncertainty in Artificial Intelligence: Proceedings of the Twentieth Conference* (Max Chickering and Joseph Y. Halpern, eds.), 226–234, AUAI Press, Arlington, Virginia. [283]

Gul, Faruk and Outi Lantto (1990), "Betweenness satisfying preferences and dynamic choice." *Journal of Economic Theory*, 52, 162–177. [281]

Hansen, Lars Peter and Thomas J. Sargent (2001), "Robust control and model uncertainty." *American Economic Review*, 91, 60–66. [273]

Hayashi, Takashi (2005), "Intertemporal substitution, risk aversion and ambiguity aversion." *Economic Theory*, 25, 933–956. [264, 282]

Jaffray, Jean-Yves (1992), "Bayesian updating and belief functions." *IEEE Transactions on Systems, Man, and Cybernetics*, 22, 1144–1152. [264]

Jaffray, Jean-Yves (1994), "Dynamic decision making with belief functions." In *Advances in the Dempster–Shafer Theory of Evidence* (Ronald R. Yager, Janusz Kacprzyk, and Mario Fedrizzi, eds.), 331–352, Wiley, New York. [264, 285, 287]

Jaffray, Jean-Yves (1999), "Rational decision making with imprecise probabilities." In *Proceedings of the First International Symposium on Imprecise Probabilities and Their Applications* (Gert De Cooman, Fabio Gagliardi Cozman, Serafin Moral, and Peter Walley, eds.), 183–188, Morgan Kaufmann, San Francisco. Available at <http://decsai.ugr.es/~smc/isipta99/proc/027.html>. [283, 284]

Karni, Edi and David Schmeidler (1991), "Atemporal dynamic consistency and expected utility theory." *Journal of Economic Theory*, 54, 401–408. [282]

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji (2005), "A smooth model of decision making under ambiguity." *Econometrica*, 73, 1849–1892. [271, 285]

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji (2006), "Recursive smooth ambiguity preferences." Notebook 17, Fondazione Collegio Carlo Alberto. [264]

Kreps, David M. and Evan L. Porteus (1978), "Temporal resolution of uncertainty and dynamic choice theory." *Econometrica*, 46, 185–200. [282]

Maccheroni, Fabio, Massimo Marinacci, and Aldo Rustichini (2006), "Dynamic variational preferences." *Journal of Economic Theory*, 128, 4–44. [264]

Machina, Mark J. (1989), "Dynamic consistency and non-expected utility models of choice under uncertainty." *Journal of Economic Literature*, 27, 1622–1668. [262, 265, 273, 280]

Machina, Mark J. and David Schmeidler (1992), "A more robust definition of subjective probability." *Econometrica*, 60, 745–780. [273, 277, 280]

- McClennen, Edward F. (1990), *Rationality and dynamic choice: Foundational explorations*. Cambridge University Press, Cambridge. [265, 273, 280, 283]
- Mukerji, Sujoy and Jean-Marc Tallon (2004), “An overview of economic applications of david schmeidler’s models of decision making under uncertainty.” In *Uncertainty in Economic Theory* (Itzhak Gilboa, ed.), 283–302, Routledge, London. [262]
- Ozdenoren, Emre and James Peck (forthcoming), “Ambiguity aversion, games against nature, and dynamic consistency.” *Games and Economic Behavior*. [284]
- Peleg, Bezalel and Menahem E. Yaari (1973), “On the existence of a consistent course of action when tastes are changing.” *Review of Economic Studies*, 40, 391–401. [283, 284]
- Pires, Cesaltina Pacheco (2002), “A rule for updating ambiguous beliefs.” *Theory and Decision*, 53, 137–152. [264, 285, 287]
- Sarin, Rakesh and Peter P. Wakker (1998), “Dynamic choice and nonexpected utility.” *Journal of Risk and Uncertainty*, 17, 87–119. [264]
- Savage, Leonard J. (1954), *The Foundations of Statistics*. Wiley, New York. [280]
- Segal, Uzi (1997), “Dynamic consistency and reference points.” *Journal of Economic Theory*, 72, 208–219. [273, 281]
- Seidenfeld, T. (2004), “A contrast between two decision rules for use with (convex) sets of probabilities: γ -maximin versus e -admissibility.” *Synthese*, 140, 69–88. [273, 283]
- Shafer, Glenn (1976), *A Mathematical Theory of Evidence*. Princeton University Press, Princeton. [264]
- Siniscalchi, Marciano (2001), “Bayesian updating for general maxmin expected utility preferences.” Discussion Paper 1366, Center for Mathematical Studies in Economics and Management Science, Northwestern University. [264]
- Siniscalchi, Marciano (2006), “Dynamic choice under ambiguity.” Unpublished paper, Northwestern University. [264, 283]
- Strotz, R. H. (1955–6), “Myopia and inconsistency in dynamic utility maximization.” *Review of Economic Studies*, 23, 165–180. [283]
- Tallon, Jean-Marc and Jean-Christophe Vergnaud (2006), “Beliefs and dynamic consistency.” In *Knowledge, Beliefs and Economics* (Richard Arena and Agnès Festré, eds.), 137–154, Edward Elgar, Northampton. [284]
- Wakker, Peter P. (1988), “Nonexpected utility as aversion of information.” *Journal of Behavioral Decision Making*, 1, 169–175. [268]
- Wakker, Peter P. (1997), “Are counterfactual decisions relevant for dynamically consistent updating under nonexpected utility?” Unpublished paper, CentER, Tilburg University. [277]

Walley, Peter (1991), *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, London. [264]

Wang, Tan (2003), "Conditional preferences and updating." *Journal of Economic Theory*, 108, 286–321. [264, 282, 283]

Wasserman, Larry A. and Joseph B. Kadane (1990), "Bayes' theorem for Choquet capacities." *Annals of Statistics*, 18, 1328–1339. [264]

Submitted 2006-3-13. Final version accepted 2007-8-2. Available online 2007-8-3.