Dynamically Consistent Updating of Multiple Prior Beliefs – an Algorithmic Approach

Eran Hanany[*] Peter Klibanoff[†] Erez Marom[‡]

January 2011 Revised: June 2011

Abstract

This paper develops algorithms for dynamically consistent updating of ambiguous beliefs in the maxmin expected utility model of decision making under ambiguity. Dynamic consistency is the requirement that ex-ante contingent choices are respected by updated preferences. Such updating, in this context, implies dependence on the feasible set of payoff vectors available in the problem and/or on an ex-ante optimal act for the problem. Despite this complication, the algorithms are formulated concisely and are easy to implement, thus making dynamically consistent updating operational in the presence of ambiguity.

Key Words: Utility Theory, Uncertainty Modelling, Risk Analysis, Updating Beliefs, Ambiguity

1 Introduction

A central task facing any theory of decision making under uncertainty is updating preferences in response to new information (see e.g., Winkler 1972). 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 (see Hanany and Klibanoff 2007 for a formal definition of dynamic consistency and a detailed discussion of its relation to other dynamic consistency

[*]Department of Industrial Engineering, Tel Aviv University, Tel Aviv 69978, Israel. Email: hananye@post.tau.ac.il

[†]Department of Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, Evanston, IL 60208, USA. E-mail: peterk@kellogg.northwestern.edu

[‡]Department of Industrial Engineering, Tel Aviv University, Tel Aviv 69978, Israel. Email: erezmarom@hotmail.com

concepts in the decision theory literature). 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 decision analysis.

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

The study of dynamic consistency is in a well defined sense the study of optimal updating, as dynamically consistent update rules result in maximal (ex-ante) welfare. Moreover, 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. In a recent paper, Hanany and Klibanoff 2007 pursued this strategy to update preferences of the max-min expected utility (MEU) form (Gilboa and Schmeidler 1989). For these preferences, beliefs are ambiguous in the sense that they are represented by a (compact and convex) set of probability measures, rather than the usual single measure of EU, and acts are evaluated by the minimum expected utility generated by these measures. MEU preferences are widely used in modeling ambiguity averse behavior, as exemplified by the famous Ellsberg 1961 paradoxes (for a survey of economic applications of MEU see Mukerji and Tallon 2004). The objective of this paper is to offer ways to explicitly compute the decision maker's (DM's) updated beliefs according to the update rules suggested by Hanany and Klibanoff 2007, thus making these rules operational. For each rule, the updated ambiguous beliefs are computed using constructive algorithms. The algorithms we describe allow explicit implementation of these rules, for example via computer programs. The rules allow the dynamically consistent updating of any set of MEU beliefs upon observing any non-null event. These rules all are generalizations of Bayesian updating in the sense that they specialize to Bayes' rule when the set of measures is a singleton. In common with Bayesian updating and many other models, our approach does not address preferences conditional on completely unanticipated (i.e., null) events. Similarly, it does not consider costs, cognitive or otherwise, of describing events. For a discussion of these and other aspects of dynamic decisions see, e.g., Hacking 1967, Modica 2008, Schipper 2010.

To better understand the issues involved in dynamically consistent updating under ambiguity, consider a version of Ellsberg's three-color problem. There is an urn containing 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 $(1,0,0)$ preferred to $(0,1,0)$ and $(0,1,1)$ preferred to $(1,0,1)$, 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 also made on yellow (thus violating the sure-thing principle of EU). 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. In the first choice pair, the bet on black or red is not paired with a bet on yellow, so the choice "Bet on B" leads to the payoff vector $(1,0,0)$ while the choice "Bet on R" leads to payoffs $(0,1,0)$. In the second choice pair, a bet on yellow is included, so the choice "Bet on B" leads to the payoff vector $(1,0,1)$ while the choice "Bet on R" leads to payoffs $(0,1,1)$. Since the DM can condition on whether or not a yellow ball is drawn, the complete set of pure strategies available in each choice problem is: ("Bet on B" if not Y; "Bet on B" if Y), ("Bet on B" if not Y; "Bet on R" if Y), ("Bet on R" if not Y; "Bet on B" if Y), ("Bet on R" if not Y; "Bet on R" if Y). Notice that when translated into payoff vectors, these strategies yield only $\{(1,0,0),(0,1,0)\}$ in the first choice problem and $\{(1,0,1),(0,1,1)\}$ in the second choice problem. If payoff vectors are what the DM cares about, then this provides a strong argument that choices in the dynamic version of the Ellsberg problem should be the same as in the original problem since the feasible payoff vectors are the same. This is closely related to ensuring that the value of information is always non-negative, another desirable principle for decision making (see e.g., Wakker 1988). If choices in the dynamic version of the problem differed from those in the original problem (where the choices are made before the information is revealed) then the DM would strictly prefer to face the original problem, thus declining the (free) information.

What kind of rules for updating ambiguous beliefs imply that the choices in the dynamic version are the same as in the original problem (not only for this example, but in general)? In Hanany and Klibanoff 2007 we provide such rules, and also show that any such rules must depend on the feasible set of payoff vectors available in the problem and/or on an ex-ante optimal act for the problem. To make this clear, consider the following specification of MEU preferences that allow the Ellsberg choices in the example above. 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, $u : \mathbb{R} \to \mathbb{R}$, such that $\forall f, g \in \mathbb{R}^3$, $f \succsim h \iff \min_{q \in C} \int (u \circ f) dq \geq \min_{q \in C} \int (u \circ h) dq$. Let $u(x) = x$ for all $x \in \mathbb{R}$

and let $C = \left\{ \left(\frac{1}{3}, \alpha, \frac{2}{3} - \alpha \right) \mid \alpha \in \left[\frac{1}{4}, \frac{5}{12} \right] \right\}$, a set of measures symmetric with respect to the probabilities of red and yellow and consistent with the information that 40 of the 120 balls are black. Observe that, indeed, $(1, 0, 0)$ is preferred to $(0, 1, 0)$ and $(0, 1, 1)$ is preferred to $(1, 0, 1)$ according to these preferences. If we apply full Bayesian updating (Bayesian conditioning of each measure in C) conditional on the event $E = \{B, R\}$, the updated set of measures is $C_E = \left\{ (\alpha, 1 - \alpha, 0) \mid \alpha \in \left[\frac{4}{9}, \frac{4}{7} \right] \right\}$. According to these updated preferences, "Bet on B" is strictly preferred to "Bet on R" conditional on learning $E = \{B, R\}$, leading $(1, 0, 0)$ to be selected over $(0, 1, 0)$ in choice problem 1, in agreement with the unconditional preferences, but also leading $(1, 0, 1)$ to be selected over $(0, 1, 1)$ in choice problem 2, in conflict with the unconditional preferences. It follows that for an update rule to maintain the choices according to the unconditional preferences for both choice problems, the rule must depend on the feasible set and/or on an ex-ante optimal act. A dynamically consistent update in this example corresponds to Bayesian updating of measures in a particular strict subset of the ex-ante set of measures, specifically the conditional set of measures is $C_E^1 = \left\{ (\alpha, 1 - \alpha, 0) \mid \alpha \in \left[\frac{1}{2}, \frac{4}{7} \right] \right\}$ for choice problem 1 and $C_E^2 = \left\{ (\alpha, 1 - \alpha, 0) \mid \alpha \in \left[\frac{4}{9}, \frac{1}{2} \right] \right\}$ in choice problem 2. This emphasizes that although natural analogues to updating beliefs under EU exist for updating beliefs under ambiguity, dynamic consistency requires novel procedures that operate somewhat differently.

The prior literature on updating ambiguous beliefs in MEU proposes and explores rules that, in fact, fail dynamic consistency for at least some MEU preferences. This literature includes many well-known rules, such as full (or generalized) Bayesian updating, maximum likelihood updating, or Dempster-Shafer updating. Full Bayesian updating 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). Maximum likelihood updating says that, of the set of priors, only those measures assigning the highest probability to the observed event should be updated using Bayes' rule, and the other priors should be discarded (see Gilboa and Schmeidler 1993). Kriegler 2009 advocates a hybrid approach applying full Bayesian updating to a set of measures formed through ε-contamination where additionally ε is updated through a maximum likelihood procedure. Both full Bayesian and maximum likelihood updating are given an interpretation in terms of epistemic belief hierarchies by Walliser and Zwirn 2011. For Dempster-Shafer updating see Dempster 1968, Shafer 1976. Jaffray 1999 suggests that the inconsistency between unconditional and conditional preferences might be resolved in a way that is 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. Nielsen and Jaffray 2006 construct algorithms for implementing the approach suggested by Jaffray 1999 in the context of risk.

In contrast to this literature, Hanany and Klibanoff 2007 identify update rules that are dynamically consistent for any MEU preferences upon observing any non-null event. Given the necessary dependence of the consistent rules on the feasible set and/or an ex-ante optimal act, the practicality of their implementation is an issue. This motivates the present paper, where we provide algorithms for computing the updated beliefs determined by these consistent rules. In developing the algorithms and proving that they implement the various rules, we draw on techniques from convex analysis.

In the next section we describe the framework for our analysis. Section 3 describes algorithms to compute updated beliefs. Section 4 provides a short summary. Proofs are collected in Appendix A. Appendix B collects some useful results from polyhedral theory. Code for an implementation of the algorithms in the paper using Wolfram *Mathematica* is available as an online supplement on the journal webpage or from the authors.

2 Framework

Consider the set X of all simple (i.e., finite-support) lotteries over a set of consequences Z, a finite set of states of nature S endowed with the algebra Σ of all events, and the set $\mathcal{A}$ of all acts, i.e. functions $f : S \to X$. Consider a non-degenerate max-min expected utility (MEU, Gilboa and Schmeidler 1989) preference relation $\succsim$ over $\mathcal{A}$, for which there exists a compact and convex set of probability measures with a finite set of extreme points, C, and a vonNeumann-Morgenstern (vN-M) EU function, $u : X \to \mathbb{R}$, such that $\forall f, h \in \mathcal{A}$, $f \succsim h \iff \min_{q \in C} \int (u \circ f) dq \geq \min_{q \in C} \int (u \circ h) dq$. Let $\mathcal{P}^{MEU}$ denote the set of all such preference relations. If $\succsim$ is non-degenerate, C is unique and u is unique (among vN-M EU functions) up to positive affine transformations. As usual, $\sim$ and $\succ$ denote the symmetric and asymmetric parts of $\succsim$.

Let $\mathcal{N}(\succsim)$ denote the set of events $E \in \Sigma$ for which $\forall q \in C$, $q(E) > 0$. We limit attention to updating on events that are non-null in this sense. For $E \in \Sigma$, let $\Delta(E)$ denote the set of all 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 1963 mixtures) and compact with a finite set of extreme points. Elements of $\mathcal{B}$ are considered feasible sets and their convexity could be justified, for example, by randomization over acts. Compactness is needed to ensure the existence of optimal acts.

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 $U : \mathcal{T} \to \mathcal{P}^{MEU}$, producing for each $(\succsim, E, g, B) \in \mathcal{T}$ a MEU conditional preference, denoted $\succsim_{E,g,B}$, representable using the same (up to normalization) vN-M utility function u as $\succsim$ and a non-empty, closed and convex set of conditional measures $C_{E,g,B} \subseteq C_E \equiv \{q_E \mid q \in C\}$, with a finite set of extreme points. Such a conditional preference is viewed as governing choice upon the realization of the conditioning event E. Let $\mathcal{U}^{Bayes}$ denote the set of all such update rules.

Abusing notation in the standard way, $x \in X$ is also used to denote the constant act for which $\forall s \in S$, $f(s) = x$. For any $f, h \in \mathcal{A}$, we use $f_E h$ to denote the act equal to f on E and h on E^c, the complement of E. General vectors in $\mathbb{R}^{|S|}$ will be called utility acts. If a and b are utility acts, we use $a_E b$ to denote the utility act equal to a on E and b on E^c. Since S is finite, we sometimes identify probability measures with vectors in $\mathbb{R}^{|S|}$ normalized to sum to 1. For an arbitrary convex, compact set of real vectors, A, denote by $ext(A)$ the set of extreme points of A. Let $\bar{C} = ext\,(C)$ and $\bar{B} = ext\,(u \circ B)$, thus $C = co\left(\bar{C}\right)$ and $u \circ B = co\left(\bar{B}\right)$, where co denotes the convex hull operator. This implies that C_E is the convex hull of a finite number of points. For each $a \in \mathbb{R}^{|S|}$ and $\xi \in \mathbb{R}$, denote the half-space $\left\{c \in \mathbb{R}^{|S|} \mid a \cdot c \geq \xi\right\}$ by W_a^ξ and the hyperplane $\left\{c \in \mathbb{R}^{|S|} \mid a \cdot c = \xi\right\}$ by H_a^ξ.

As discussed in the introduction, the update rules considered in this paper depend on an initially optimal act g and the feasible set B. Before introducing update rules, we describe a simple algorithm for computing an initially optimal act.

Algorithm 2.1 *Solve the linear program,*

$$\begin{array}{l} (o^*, \lambda^*) \in \arg\max_{(o,\lambda) \in \mathbb{R} \times [0,1]^{|\bar{B}|}} o \\ s.t. \\ \sum_{s \in S} \sum_{a \in \bar{B}} \lambda_a a_s q_s \geq o \quad , \forall q \in \bar{C} \\ \sum_{a \in \bar{B}} \lambda_a = 1. \end{array} \tag{P}$$

Let $b = \sum_{a \in \bar{B}} \lambda_a^* a$ *and* $g \in B$ *such that* $b = u \circ g$.

Proposition 2.1 *The act* g *is optimal in* B *according to* $\succsim$.

The following example will be used throughout the paper to demonstrate our results.

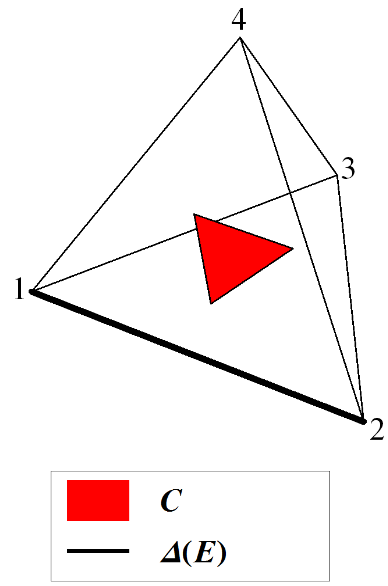

Figure 2.1: The set of probability measures C for our leading example (see Example 2.1). Updating of C must result in a subset of $\Delta(E)$.

Example 2.1 *Consider a state space $S = \{1, 2, 3, 4\}$ and two investment options l, n leading to state contingent monetary payments $(-5, -5, 20, 20)$ and $(10, -15, 10, -15)$, respectively. For simplicity, assume risk neutrality and take utility u over monetary outcomes in $Z = \mathbb{R}$ to be the identity function (our analysis applies to any assumption on risk attitudes). The ambiguous beliefs on the state space are represented by the set $C = co\{(0.4, 0.4, 0.1, 0.1), (0.1, 0.4, 0.25, 0.25), (0.4, 0.1, 0.25, 0.25)\}$. Before investing, it is possible to pay a cost of 1 to reveal whether the event $E = \{1, 2\}$ is true and make an investment decision based on this information. Figure 2.1 illustrates the sets C and $\Delta(E)$ within the simplex of all probability measures over S. Thus, before randomization, there are seven feasible alternatives: the two investment options l, n without buying the information, the non-investment option and four information contingent investment possibilities. This is summarized by the feasible utility set $u \circ B = co\{(-5, -5, 20, 20), (10, -15, 10, -15), (0, 0, 0, 0), (-6, -6, 19, 19), (9, -16, 9, -16), (-6, -6, 9, -16), (9, -16, 19, 19)\}$. Applying Algorithm 2.1 results in an optimum $o^* = 1$ and $\lambda^* = (0, 0, 0, 0, 0, 0, 1)$, i.e. buy the information and invest in n if E and in l if not E. Note that $b = u \circ g = (9, -16, 19, 19)$, evaluated by $\min_{q \in C} \sum_{s \in S} b_s q_s = 1$. In the next section we analyze the updated ambiguous beliefs in this problem. Throughout the paper we analyze the case where E is observed – similar analysis, omitted for brevity, can be done for the case where E^c is observed.*

3 Computing updated beliefs

3.1 Dynamically consistent update rules

Dynamic Consistency, an important and intuitive property of update rules, means that initial contingent plans should be respected after receiving information. In our framework, the act g should remain optimal also conditionally.

Axiom 3.1 ***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? 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 remained 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. An act $f = g$ on E^c will be called *comparable* to g.

Among rules satisfying desiderata such as dynamic consistency, there are those that are most *conservative* in the sense of maintaining the most ambiguity in the process of updating. In general, this means that the updated set of measures should be the largest possible subset of C_E, in the sense of set inclusion, while still satisfying the desiderata. 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. Note that if dynamic consistency is ignored, and all rules in $\mathcal{U}^{Bayes}$ are considered, ambiguity maximization would uniquely select full Bayesian updating.

Let us now observe the implications of **DC** for updating. The following algorithm can be used to compute the set of updated beliefs under the unique ambiguity maximizing update rule (denoted $U^{DC\max}$) within the dynamically consistent rules in $\mathcal{U}^{Bayes}$. Having found an initially optimal act g in Step **3.1.A**, the algorithm computes in Step **3.1.B** the set of feasible utility acts comparable to g. Step **3.1.C** computes Bayesian conditional measures normal to hyperplanes supporting this set at g. Only these measures can be used to evaluate g conditionally if dynamic consistency is required. If we were to stop the algorithm at this point and return $co(\bar{Q}_E)$, this would be a dynamically consistent update rule but need not maximize ambiguity. To maximize ambiguity, Step **3.1.D** finds the minimum evaluation of

g according to these measures, and computes the updated beliefs as the set of conditional measures that evaluate g at least as high as this minimum.

Algorithm 3.1 *Step* ***3.1.A****: Compute* $b = u \circ g$ *and* $o^* = \min_{q \in \bar{C}} \sum_{s \in S} b_s q_s$.[1]
Step ***3.1.B****: Compute* $\bar{L} \equiv ext[co(\bar{B}) \cap_{s \in E^c} H^{b_s}_{I(s)}]$, *where* $I(s) \in \mathbb{R}^{|S|}$ *denotes an s-indicator, i.e.* $I_s(s) = 1$ *and* $I_{\hat{s}}(s) = 0$ *for* $\hat{s} \neq s$.[2]
Step ***3.1.C****: Compute* $\bar{Q}_E \equiv ext[co(\bar{C}_E) \cap_{a \in \bar{L}} W^0_{b-a}]$, *where* $\bar{C}_E \equiv ext(C_E)$.
Step ***3.1.D****: Compute* $\bar{U}_E \equiv ext[co(\bar{C}_E) \cap W^0_{b-\beta_1 \mathbf{1}}]$, *where* $\beta_1 \equiv \min_{q \in \bar{Q}_E} \sum_{s \in S} b_s q_s$ *and* $\mathbf{1}$ *denotes the constant* $(1, ..., 1) \in \mathbb{R}^{|S|}$. *Return* $co(\bar{U}_E)$.

Proposition 3.1 *Algorithm 3.1 results in the updated set of measures,* $C_{E,g,B}$, *produced by the update rule* $U^{DC\max}$.

Example 3.1 *We apply the algorithm to our leading example. Since we already have an optimal utility act* $b = (9, -16, 19, 19)$ *and its value* $o^* = 1$, *we start by finding the set* $\bar{L}$, *the extreme points of the set of feasible acts comparable to* g. *The set of extreme points of* $co(\bar{B}) \cap H^{19}_{I(3)}$ *is* $\{(9, -16, 19, 19), (-6, -6, 19, 19), (-4.75, -4.75, 19, 19)\}$. *Then intersecting with* $H^{19}_{I(4)}$ *leaves this set unchanged, so this is* $\bar{L}$. *Note that this set includes the optimal utility act* b, *the utility act corresponding to the choice of buying the information and investing always in* l, *and the utility act* $(-4.75, -4.75, 19, 19)$ *resulting from a convex combination of the no investment option and investing in* l *without buying the information. Next we find the set* $\bar{C}_E$, *the extreme points of the set obtained by Bayesian conditioning of all measures in* C *(see the top part of Figure 3.1). The conditionals of the measures in* $\bar{C}$ *are* $\{(0.5, 0.5, 0, 0), (0.2, 0.8, 0, 0), (0.8, 0.2, 0, 0)\}$. *Since the first of these conditionals is in the convex hull of the other two,* $\bar{C}_E = \{(0.2, 0.8, 0, 0), (0.8, 0.2, 0, 0)\}$. *Notice that if one adopted this set as the updated ambiguous beliefs, i.e. followed the full Bayesian rule, the optimal choice given* E *would be not to follow the optimal act and instead invest in* l. *This would result in the utility act* $(-6, -6, 19, 19)$, *which is dominated (in every state) by the feasible act invest in* l *without buying the information. This demonstrates that full Bayesian updating may be an undesirable update rule. In contrast,* $U^{DC\max}$ *is dynamically consistent, and thus does not suffer from such phenomena. To see this, we continue to follow the algorithm and next find* $\bar{Q}_E$, *the extreme points of the set of measures supporting the conditional optimality of* g. *Since* $|\bar{L}| = 3$, $co(\bar{C}_E)$ *is intersected thrice, to ensure measures that support the feasible set comparable to* g. *Since* b *is always an element of* $\bar{L}$ *and* $b - b = 0$, *the intersection based on* b *never*

[1] If g is not known, one can apply Algorithm 2.1 to carry out this step.

[2] When computing the extreme points of the intersection of a list of hyperplanes/half-spaces with the convex hull of a finite set of points, this algorithm and all subsequent ones use standard procedures from polyhedral theory (e.g., as in Appendix B) and thus we do not detail these procedures here.

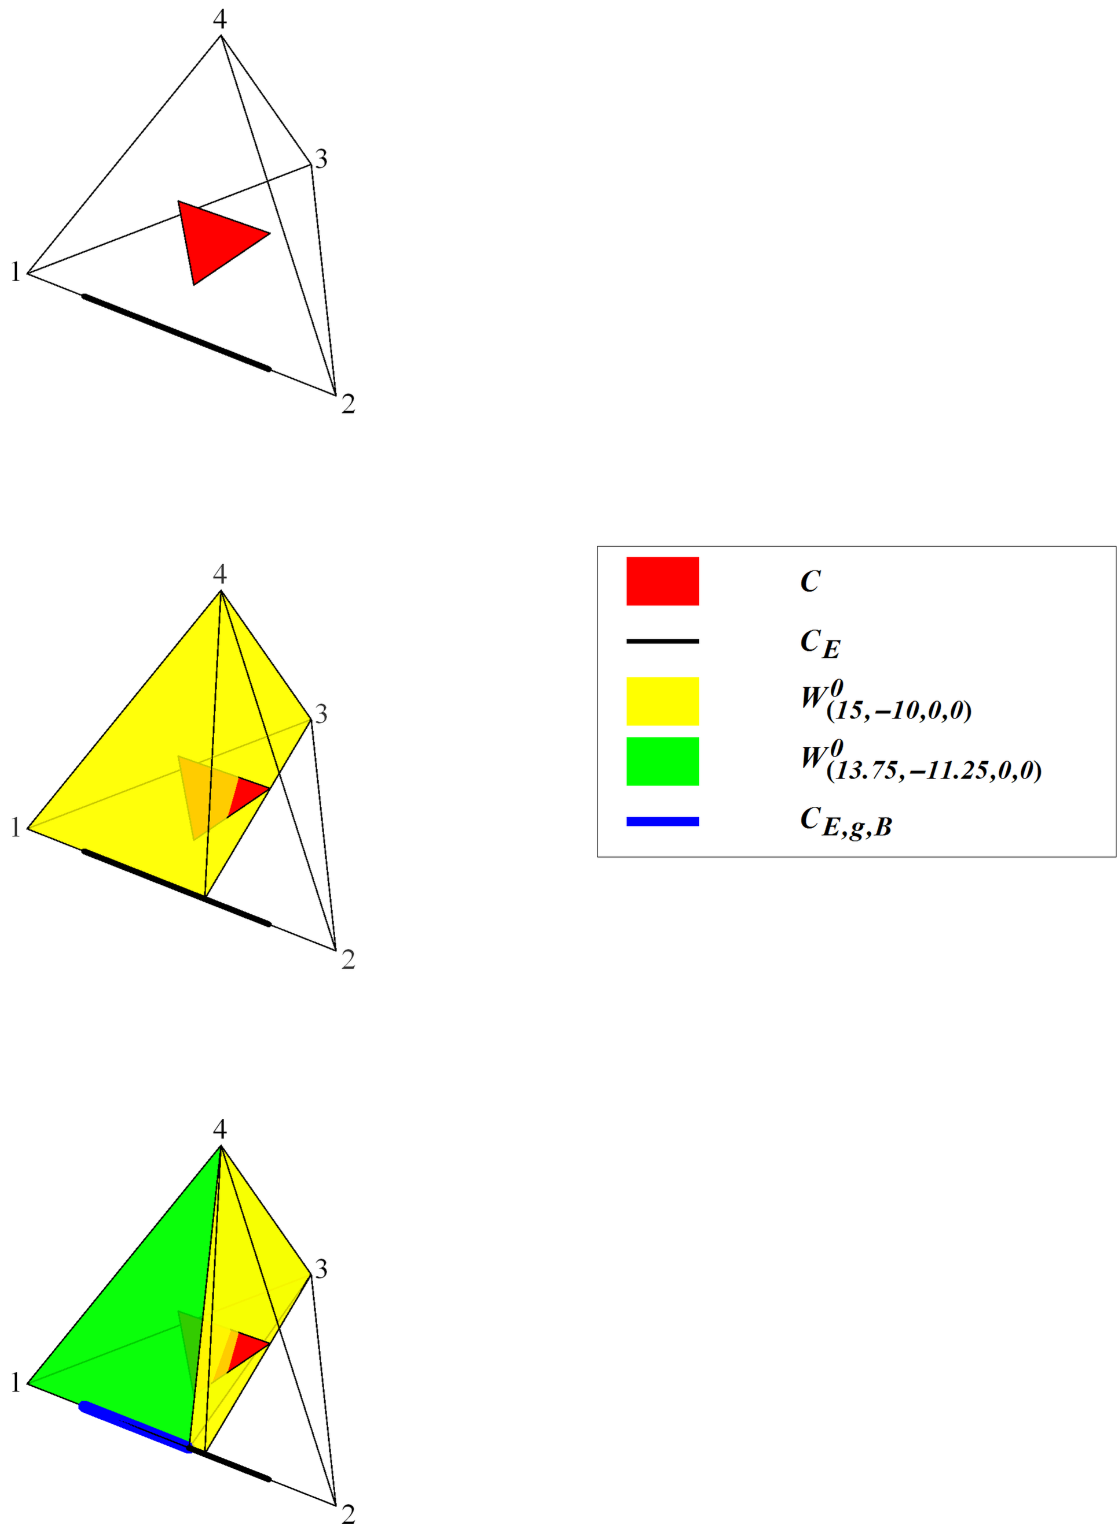

Figure 3.1: Illustration of Algorithm 3.1 for computing the update rule $U^{DC\max}$ (see Example 3.1).

imposes any restriction. Considering $b-(-6,-6,19,19)=(15,-10,0,0)$*,* $W^0_{(15,-10,0,0)}$ *corresponds to the set of measures such that* $15q_1-10q_2\geq 0$ *(see the middle part of Figure 3.1), so the set of extreme points of* $co(\bar{C}_E)\cap W^0_{(15,-10,0,0)}$ *is* $\{(0.4,0.6,0,0),(0.8,0.2,0,0)\}$*. The final intersection, with* $W^0_{(13.75,-11.25,0,0)}$*, requires* $13.75q_1-11.25q_2\geq 0$ *and thus shrinks the set of extreme points to* $\{(0.45,0.55,0,0),(0.8,0.2,0,0)\}$*, which is* $\bar{Q}_E$ *(see the bottom part of Figure 3.1). Note that any conditional measure putting weight lower than* 0.45 *to state* 1 *does not support the conditional optimality of* g *because it becomes worse than* $(-4.75,-4.75,19,19)$*. Finally we calculate* $\bar{U}_E$*, the set of extreme points of the set of updated ambiguous beliefs. Observe that* $\beta_1=\min\{(9,-16,19,19)\cdot(0.45,0.55,0,0),(9,-16,19,19)\cdot(0.8,0.2,0,0)\}=-4.75$*. Since* $q\in co(\bar{C}_E)$*,* $(b-\beta_1 1)\cdot q\geq 0$ *if and only if* $13.75q_1-11.25q_2\geq 0$*, which is not weaker than the restrictions already imposed, so* $\bar{U}_E\equiv co(\bar{C}_E)\cap W^0_{b-\beta_1\mathbf{1}}=\bar{Q}_E$*. Thus* $C_{E,g,B}=co\{(0.45,0.55,0,0),(0.8,0.2,0,0)\}$*. Note that given these updated ambiguous beliefs, the optimal act is conditionally equivalent to* -4.75*, which is better than* -6*, the payment if one invests instead in* l*. Thus dynamic consistency is satisfied. The example shows that it is necessary to eliminate some of the measures in the original set, in particular all measures that conditionally give weight lower than* 0.45 *to state* 1*. When a plan is made initially to follow the optimal act, the updated ambiguous beliefs represented by* $C_{E,g,B}$ *justify this plan: contingent on learning* E*, it will be optimal to carry it out.*

Note that computing the updated ambiguous beliefs is a hard problem because it involves computing the extreme points of intersections of convex hulls of finite sets of points, which is known to be hard. If these computations could be done efficiently, then all the algorithms in this paper would be polynomial in the size of the problem.

It is also worth noting that in cases where the feasible set has a special structure such that $f,f'\in B\implies f_Ef'\in B$, the restriction $f=g$ on E^c in the statement of **DC** is superfluous, thus the computation of Step **3.1.B** can be simplified to $\bar{L}\equiv\{a\in\bar{B}\mid a=b$ on $E^c\}$. Such feasible sets arise, for example, whenever one starts from a decision tree with branches corresponding to events E and E^c and derives B by first specifying what is feasible conditional on E, denoted B^E, and what is feasible conditional on E^c, denoted B^{E^c}, and then combining the two so that $B=\left\{h\in\mathcal{A}\mid h=f_Ef' \text{ for some } f\in B^E,\ f'\in B^{E^c}\right\}$.

3.2 More robust dynamic consistency

It may be desirable to strengthen dynamic consistency, so that all initially optimal acts comparable to g should remain optimal conditional on E.

Axiom 3.2 ***PFI*** *(Robust Dynamic Consistency).*[3] *For any* $(\succsim,E,g,B)\in\mathcal{T}$*, if* $f\in B$ *with*

[3] This axiom was named **PFI** (*Preservation of Feasible Optimal Indifference*) by Hanany and Klibanoff

$f = g$ on E^c and $f \sim g$, then $f \sim_{E,g,B} g$.

The following algorithm can be used to compute the set of updated beliefs under the unique ambiguity maximizing update rule (denoted $U^{DC \cap PFI \max}$) within the rules in $\mathcal{U}^{Bayes}$ satisfying **DC** and **PFI**. As in Algorithm 3.1, steps **3.2.A** and **3.2.B** find an initially optimal act g, the set of feasible acts comparable to g and the set of measures supporting the conditional optimality of g. Step **3.2.C** computes the measures normal to hyperplanes separating the feasible acts comparable to g from the acts strictly better than g. Step **3.2.D** uses one of these measures to compute the set of initially optimal acts comparable to g. Step **3.2.E** computes the Bayesian conditional measures normal to hyperplanes supporting this set at g. Only these measures can be used to evaluate g conditionally if robust dynamic consistency is required. To maximize ambiguity, Step **3.2.F** finds the minimum evaluation of g according to these measures, and computes the updated beliefs as the set of conditional measures that evaluate all initially optimal acts at least as high as this minimum.

Algorithm 3.2 *Step **3.2.A:** Compute $b = u \circ g$ and o^* as in Step **3.1.A** of Algorithm 3.1.*
*Step **3.2.B:** Compute $\bar{L}, \bar{Q}_E$ as in steps **3.1.B** and **3.1.C** of Algorithm 3.1, respectively.*
*Step **3.2.C:** Compute $\bar{R} \equiv ext[co(\bar{C}) \cap H_b^{o^*} \cap_{a \in \bar{L}} W_{b-a}^0]$.*
*Step **3.2.D:** Compute $\bar{J} \equiv ext[co(\bar{L}) \cap H_{q^g}^{o^*} \cap_{q \in \bar{C}} W_{q-q^g}^0]$, where $q^g \in \bar{R}$.*
*Step **3.2.E**: Compute $\bar{K}_E \equiv ext[co(\bar{Q}_E) \cap_{a \in \bar{J}} H_{b-a}^0]$.*
*Step **3.2.F**: Compute $\bar{U}_E \equiv ext[co(\bar{C}_E) \cap_{a \in \bar{J}} W_{a-\beta_2 \mathbf{1}}^0]$, where $\beta_2 \equiv \min_{q \in \bar{K}_E} \sum_{s \in S} b_s q_s$. Return $co(\bar{U}_E)$.*

Proposition 3.2 *Algorithm 3.2 results in the updated set of measures, $C_{E,g,B}$, produced by the update rule $U^{DC \cap PFI \max}$.*

The following example serves to demonstrate the comparative robustness of updating under **PFI** and **DC**, as compared to updating under only **DC**. It also illustrates the algorithm above.

Example 3.2 *We apply the algorithm to our leading example. It can be verified that the optimum found using Algorithm 2.1 is unique, thus PFI has no bite. Therefore we modify the example slightly and assume that the cost of buying the information is 2. Applying Algorithm 2.1 again, the previous optimal act – buy the information and invest in n if E and in l if not E – remains optimal. With the new cost, this is an optimal utility act*

2007. We call it *Robust Dynamic Consistency* here, as upon reflection, this seems a more informative name. To maintain continuity with our earlier nomenclature we will nonetheless continue to refer to it by the initials **PFI**.

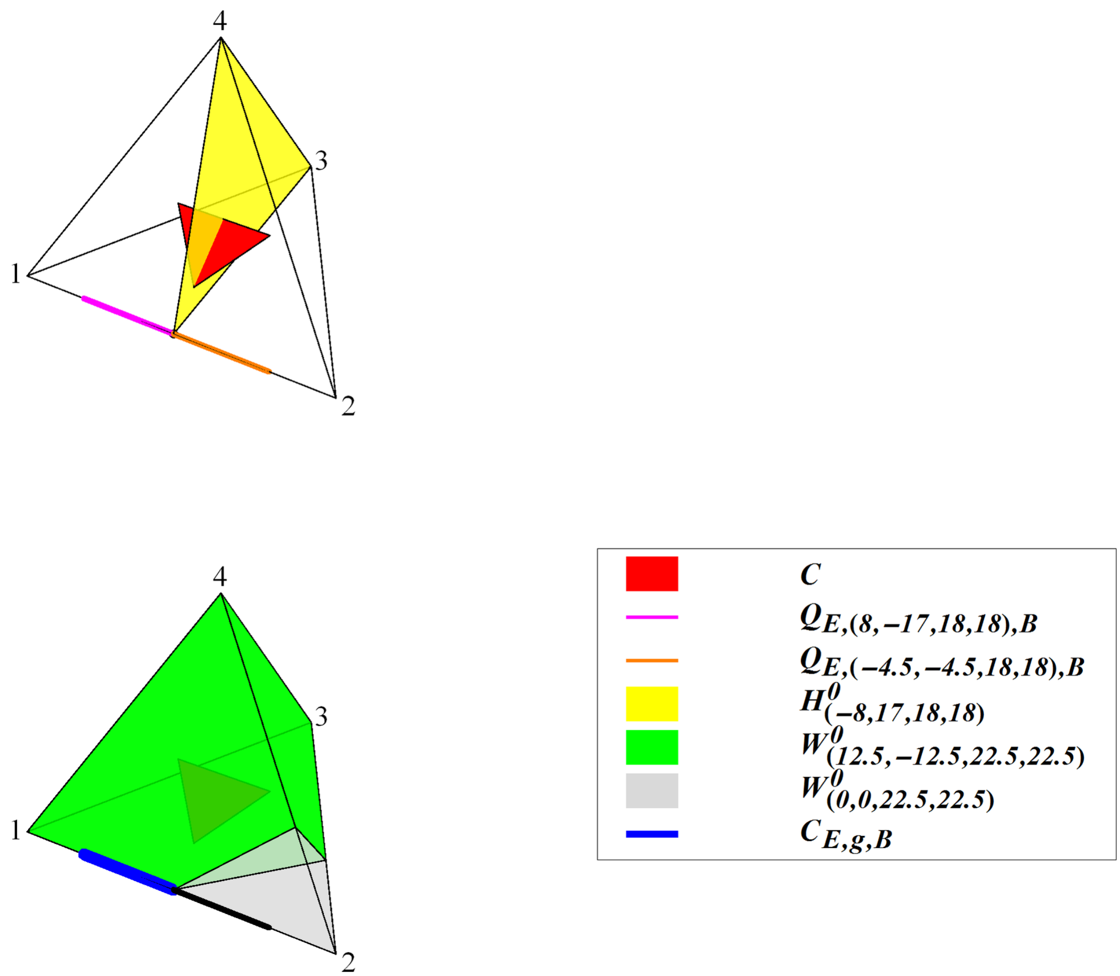

Figure 3.2: Illustration of Algorithm 3.2 for computing the update rule $U^{DC \cap PFI \max}$ (see Example 3.2).

$b = (8, -17, 18, 18)$ *with optimal value* $o^* = 0$. *The set of feasible utility acts comparable to* g *now has the extreme points* $\bar{L} = \{(8, -17, 18, 18), (-7, -7, 18, 18), (-4.5, -4.5, 18, 18)\}$. *The set of measures supporting the conditional optimality of* g *now has the extreme points* $\bar{Q}_E = \{(0.5, 0.5, 0, 0), (0.8, 0.2, 0, 0)\}$ *(see the top part of Figure 3.2). Note that the change in* $\bar{L}$ *caused a change in* $\bar{Q}_E$. *In particular, any conditional measure putting weight less than* 0.5 *to state* 1 *now does not support the conditional optimality of* g *because it becomes worse than* $(-4.5, -4.5, 18, 18)$. *Observe that* $co(\bar{C}) \cap H_b^{o^*} = \{(0.4, 0.4, 0.1, 0.1)\}$, *so* $\bar{R} = \{q^g\} = \{(0.4, 0.4, 0.1, 0.1)\}$. *Next we find* $\bar{J}$, *the extreme points of the initially optimal acts comparable to* g. *The intersection* $co(\bar{L}) \cap H_{q^g}^{o^*}$ *results in the extreme points* $\{(8, -17, 18, 18), (-4.5, -4.5, 18, 18)\}$. *Since* $\left|\bar{C}\right| = 3$, *we need 3 more intersections, which leave the set unchanged, so this is* $\bar{J}$. *Note that the second extreme point of the initially optimal acts results from a convex combination of the no investment option and investing in* l *without buying the information. We move on to compute* $\bar{K}_E$, *the extreme points of the set of measures supporting the conditional optimality of all initially optimal acts comparable to* g. *Since* $\left|\bar{J}\right| = 2$ *but the intersection based on* $b \in \bar{J}$ *never imposes any restriction, we intersect* $\bar{Q}_E$ *once, which leads to* $\bar{K}_E = \{(0.5, 0.5, 0, 0)\}$ *(see the top part of Figure 3.2). Finally, we compute* $\bar{U}_E$, *the set of extreme points of the set of updated beliefs. Observe that* $\beta_2 = \min\{(8, -17, 18, 18) \cdot (0.5, 0.5, 0, 0)\} = -4.5$. *Intersecting* $co(\bar{C}_E)$ *with* $W^0_{a-\beta_2\mathbf{1}}$ *for each* $a \in \bar{J}$ *leads to* $\bar{U}_E = \bar{Q}_E$ *(see the bottom part of Figure 3.2). Thus the updated ambiguous beliefs under* $U^{DC\cap PFI\max}$ *are* $C_{E,g,B} = co\{(0.5, 0.5, 0, 0), (0.8, 0.2, 0, 0)\}$. *Applying Algorithm 3.1 to compute the updated ambiguous beliefs under* $U^{DC\max}$ *for this example gives the same set* $C_{E,g,B}$. *To see the comparative robustness of updating under* ***PFI*** *and* ***DC,*** *consider the updated ambiguous beliefs given a different initially optimal act, corresponding to the utility act* $(-4.5, -4.5, 18, 18)$. *Under* $U^{DC\cap PFI\max}$, *the updated beliefs will be the same as above because this rule guarantees that all initially optimal acts in* $co\{(8, -17, 18, 18), (-4.5, -4.5, 18, 18)\}$ *remain optimal conditionally. However,* $U^{DC\max}$ *for this case coincides with full Bayesian updating, i.e. results in more conditional measures in the updated set compared to the case with the former initially optimal act. This additional ambiguity would have affected the conditional optimality of the utility act* $(8, -17, 18, 18)$, *while none of it affects the conditional optimality of* $(-4.5, -4.5, 18, 18)$ *due to the constancy on* E. *Incidentally, for* $U^{DC\max}$ *with the initially optimal act* $(-4.5, -4.5, 18, 18)$, *in contrast to Example 3.1, Step* ***3.1.D*** *has an effect: it expands the updated set from* $co(\bar{Q}_E) = co\{(0.2, 0.8, 0, 0), (0.5, 0.5, 0, 0)\}$ *(see the top part of Figure 3.2) to* $co(\bar{U}_E) = co\{(0.2, 0.8, 0, 0), (0.8, 0.2, 0, 0)\}$.

3.3 Reference dependent updating

Consider an additional condition which, in the presence of **DC**, implies **PFI** and is stronger than **PFI** only for infeasible acts.

Axiom 3.3 ***RA*** *(g as a Reference Act). 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* $f \succsim_{E,g,B} g$*.*

One way of viewing **RA** is as saying that updating must preserve or increase ambiguity affecting g more than ambiguity affecting any act f indifferent and comparable to g.[4] **RA** has the advantage of simplifying the updated beliefs in the following important special case. When the initially optimal act g is constant (in utilities) on E^c, we can use a simple threshold rule based on the probability each measure assigns to E to obtain the updated set. It follows that a sufficient condition for the threshold rule to apply occurs when E^c is a singleton, i.e., information is learned one state at a time, a not uncommon occurrence.

The strengthening of **PFI** to **RA** does not follow from dynamic consistency considerations, however the condition does not seem unreasonable and has the above-mentioned simplification as its main virtue. In **RA**, the part of the indifference curve through g agreeing with g on E^c is picked out for a special role in updating. This part of the indifference curve may be thought of as the portion where g is being used as a reference act. The axiom requires that g occupy an extremal position in the conditional preference relative to the other elements of this part of the (unconditional) indifference curve (through g) consisting of acts agreeing with g on E^c.

The following algorithm can be used to compute the set of updated beliefs in this special case under the unique ambiguity maximizing update rule (denoted $U^{DC \cap RA \max}$) within the rules in $\mathcal{U}^{Bayes}$ satisfying **DC** and **RA**. The key new step that provides the simplification to a threshold rule is Step **3.3.D**. This step selects the measures to update by comparing the weight they give to E with $q^g(E)$ and updates all measures giving weakly more weight than $q^g(E)$ (if $u \circ g > o^*$ on E^c), weakly less weight (if $u \circ g < o^*$ on E^c) or all measures in C (if $u \circ g = o^*$ on E^c).

Algorithm 3.3 *Step* ***3.3.A****: Compute* $b = u \circ g$ *and* o^* *as in Step* ***3.1.A*** *of Algorithm 3.1.*
Step ***3.3.B****: Compute* $\bar{L}, \bar{R}$ *as in Step* ***3.1.B*** *of Algorithm 3.1 and Step* ***3.2.C*** *of Algorithm 3.2, respectively.*
Step ***3.3.C****: Compute* $q_E^g \in \arg\min_{\{q_E | q \in \bar{R}\}} \sum_{s \in E} b_s (q_E)_s$*. Let* $q^g \in \bar{R}$ *such that* q_E^g *is its*

[4] An alternative strengthening of **PFI** is obtained by replacing $f \succsim_{E,g,B} g$ with $g \succsim_{E,g,B} f$ in **RA**. However, a small modification of the example in the proof of Proposition 13 in Hanany and Klibanoff 2007 can be used to show there is no update rule in $\mathcal{U}^{Bayes}$ satisfying the alternative condition (and in fact, no such rule exists in the larger family of update rules that allows updating measures outside C).

Bayesian conditional on E.
Step ***3.3.D***: *Compute $\bar{U}_E \equiv ext(\{q_E \mid q \in co(\bar{C}) \cap W^{\alpha q^g(E)}_{\alpha \sum_{s\in E} I(s)}\})$, where $\alpha \equiv sign[b(E^c) - o^*]$; return $co(\bar{U}_E)$.*

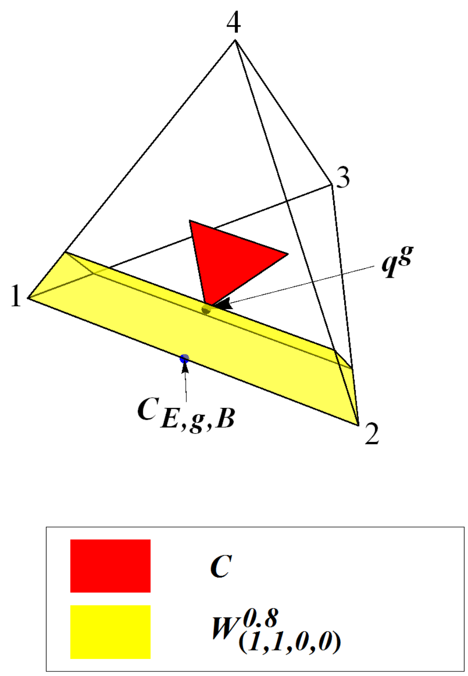

Figure 3.3: Illustration of Algorithm 3.3 for computing the update rule $U^{DC\cap RA\max}$ (see Example 3.3).

Example 3.3 *We apply this algorithm also to our leading example. Assume again that the cost of buying the information is 2. Recall the optimal act $g = (8, -17, 18, 18)$ with optimal value $o^* = 0$ and the measure in $\bar{R} = \{q^g\} = \{(0.4, 0.4, 0.1, 0.1)\}$, which uniquely separates the feasible acts comparable to g from the acts preferred to g (see Figure 3.3). Note that $u \circ g$ is constant on E^c. Since $sign(18 - o^*) = 1$, $W^{\alpha q^g(E)}_{\alpha \sum_{s\in E} I(s)}$ is equivalent to the condition $q(E) \geq q^g(E)$. The unique measure in $co(\bar{C})$ satisfying this condition is q^g, thus $\bar{U}_E$ includes only its Bayesian conditional. Therefore the updated beliefs are represented by $C_{E,g,B} = \{(0.5, 0.5, 0, 0)\}$. In general, the updated beliefs are not necessarily reduced to a single measure. For example, consider the optimal act $g = (0, 0, 0, 0)$. Since g is constant on E^c and $sign(0 - o^*) = 0$, $\bar{U}_E = \bar{C}_E$. Thus the updated beliefs coincide with those obtained by the full Bayesian update rule, i.e. $C_{E,g,B} = \{(0.2, 0.8, 0, 0), (0.8, 0.2, 0, 0)\}$.*

In fact, the algorithm producing the threshold rule is valid somewhat more generally than suggested above. It works whenever there is no ambiguity about the conditional expectation of g on E^c, in the sense that the initially optimal act g has the same conditional EU on E^c according to each measure in C giving positive probability to E^c.

Proposition 3.3 *Assume that $u(X) = \mathbb{R}$ and $\int_{E^c} b dq_{E^c}$ is the same for all $q \in C$ with $q(E^c) > 0$. Algorithm 3.3, with $b(E^c)$ replaced by $\int_{E^c} b dq_{E^c}$ if there exists $q \in C$ with $q(E^c) > 0$ and replaced by o^* otherwise, results in the updated set of measures, $C_{E,g,B}$, produced by the update rule $U^{DC \cap RA\max}$.*

The rule $U^{DC \cap RA\max}$ may also lead to intermediate updated ambiguous beliefs, i.e. to a strict subset of C_E which is not a singleton, as demonstrated in the following example.

Example 3.4 *Let $Z = \mathbb{R}$, $\bar{C} = \{(\frac{8}{24}, \frac{10}{24}, \frac{4}{24}, \frac{2}{24}), (\frac{5}{24}, \frac{4}{24}, \frac{10}{24}, \frac{5}{24})\}$, $\bar{B} = \{(2,1,3,0), (0,3,3,0)\}$ and $E = \{1,2\}$. For the unique initially optimal act, $b = u \circ g = (0,3,3,0)$ and $o^* = \frac{21}{12}$. Note that $u(X) = \mathbb{R}$ and for all $q \in C$, $q(E^c) > 0$ and $\int_{E^c} b dq_{E^c} = 2$. Applying Algorithm 3.3, $\bar{L} = \bar{B}$ and $\bar{R} = \{(\frac{8}{24}, \frac{10}{24}, \frac{4}{24}, \frac{2}{24}), (\frac{3}{12}, \frac{3}{12}, \frac{4}{12}, \frac{2}{12})\}$ (note that $\arg\min_{q \in C} \sum_{s \in S} b_s q_s = C$). Then we compute $\min_{\{q_E | q \in \bar{R}\}} \sum_{s \in E} b_s (q_E)_s = 1.5$ and $q^g = (\frac{3}{12}, \frac{3}{12}, \frac{4}{12}, \frac{2}{12})$. Since $sign(2 - o^*) = 1$, we intersect $co(\bar{C})$ with $q(E) \geq q^g(E)$, resulting in $\{(\frac{8}{24}, \frac{10}{24}, \frac{4}{24}, \frac{2}{24}), (\frac{6}{24}, \frac{6}{24}, \frac{8}{24}, \frac{10}{24})\}$, and compute the Bayesian conditionals to find $\bar{U}_E$. Thus the updated ambiguous beliefs are represented by $C_{E,g,B} = co\{(\frac{4}{9}, \frac{5}{9}, 0, 0), (\frac{1}{2}, \frac{1}{2}, 0, 0)\}$.*

For completeness, in Appendix A we provide a modification of Step **3.3.D** that we show allows computation of the updated beliefs under $U^{DC \cap RA\max}$ for the general case. To compute the extreme points of the updated beliefs, this modification considers pairs of extreme points of C and selects zero, one or two points from their convex hull to update.

4 Summary

This paper develops algorithms for updating ambiguous beliefs in the MEU model of decision making under ambiguity. The update rules all satisfy the desirable property of dynamic consistency as was shown in Hanany and Klibanoff 2007. Some of the rules also satisfy stronger and more robust consistency requirements as well. The algorithms are formulated concisely and are easy to implement, thus making dynamically consistent updating operational in the presence of ambiguity.

We close by mentioning two possible directions for future research. First, the algorithms in this paper deal only with finitely generated sets of beliefs and feasible acts. The question of whether these algorithms can be used to approximate dynamically consistent updating for arbitrary convex sets of beliefs or feasible acts is left open.

Second, although MEU is a popular theory of decision making with ambiguity aversion, it is far from the only one (see, among many, Hazen 1989, Klibanoff, Marinacci, Mukerji 2005, Maccheroni, Marinacci, Rustichini 2006, Nau 2006, Schmeidler 1989, Tversky and Kahneman 1992). In Hanany and Klibanoff 2009 we expand the approach in Hanany and Klibanoff 2007 to characterize dynamically consistent update rules for a very broad class of ambiguity averse preferences. Algorithmic implementation in these more general settings is left open as well.

Acknowledgements We thank Yigal Gerchak, Tal Raviv, two anonymous referees, the area editor and Thierry Denoeux, the Editor-in-Chief, for helpful discussions and comments. This research was partially supported by Grant No. 2006264 from the United States-Israel Binational Science Foundation (BSF), Jerusalem, Israel. All responsibility for the contents of the paper lies solely with the authors.

A Appendix: proofs and auxiliary lemmata

To conserve space in this Appendix, whenever we reference Hanany and Klibanoff 2007, we use the abbreviation HK.

Proof of Proposition 2.1. Let $q^* \in \bar{C}$ satisfy $\sum_{s\in S} b_s q^*_s = o^*$. For every $a \in \bar{B}$, there exists $f \in B$ such that $a = u \circ f$, so g exists by convexity of $u \circ B$. Moreover, g is optimal, otherwise there exists $\bar{g} \in B$ and $\bar{\lambda} \in [0,1]^{|\bar{B}|}$ such that $u \circ \bar{g} = \sum_{a\in\bar{B}} \bar{\lambda}_a a_s$ and $\bar{o} \equiv \min_{q\in\bar{C}} \sum_{s\in S} u[\bar{g}(s)]\, q_s > \min_{q\in\bar{C}} \sum_{s\in S} u[g(s)]\, q_s$, thus $\bar{o} > \sum_{s\in S} \sum_{a\in\bar{B}} \lambda^*_a a_s q^*_s = o^*$ in contradiction to the maximality of o^*. ■

Proof of Proposition 3.1. HK showed (p.270, Corollary 1) that $U^{DC\max}$ is defined by $C_{E,g,B} = \{q_E \mid q \in C \text{ and } \int (u\circ g)\, dq_E \geq \min_{p\in Q_E^{E,g,B}} \int (u \circ g)\, dp\}$, where $Q_E^{E,g,B}$ is the set of Bayesian conditionals on E of $Q^{E,g,B}$, the measures in C supporting the conditional optimality of g, defined by $Q^{E,g,B} \equiv \{q \in C \mid \int (u \circ g)\, dq_E \geq \int (u \circ f)\, dq_E$ for all $f \in B$ with $f = g$ on $E^c\}$. Since the hyperplanes used in step **3.1.B** are $\left\{c \in \mathbb{R}^{|S|} \mid c_s = b_s\right\}$, $co(\bar{L}) = u \circ \{f \in B \mid f = g \text{ on } E^c\}$. By definition, $Q_E^{E,g,B}$ is the intersection of C_E and all sets of the form $\left\{q \in C_E \mid \sum_{s\in S} (b_s - a_s)\, q_s \geq 0\right\}$ where $a \in u \circ \{f \in B \mid f = g \text{ on } E^c\}$, thus $Q_E^{E,g,B} = co(\bar{Q}_E)$. The fact that it suffices to apply a finite number of intersections ($\forall a \in \bar{L}$ and not $\forall a \in u \circ \{f \in B \mid f = g \text{ on } E^c\}$) follows from Proposition A.1 below. Finally, by definition of $U^{DC\max}$, $C_{E,g,B}$ is the set $\left\{q \in C_E \mid \sum_{s\in S} (b_s - \beta_1)\, q_s \geq 0\right\}$. ■

The following proposition shows that algorithms for computing the intersection of a compact, convex set with a finite number of half-spaces or hyperplanes (see e.g. Algorithm

B.2 in Appendix B) can also be used to compute the intersection of such a set with an *infinite* number of half-spaces/hyperplanes, when the latter are defined by normals taken from a compact, convex set having a finite number of extreme points.

Proposition A.1 *Fix $\xi \in \mathbb{R}$. Let $A, D \subset \mathbb{R}^{|S|}$ be convex, compact sets, each with a finite number of extreme points. Denote ext(D) by $\left\{a^k\right\}_{k=1}^r$. Then*

$$ext\left(A \bigcap_{k=1}^{r} W_{a^k}^{\xi}\right) = ext\left(A \bigcap_{a \in D} W_a^{\xi}\right).$$

The same is true when substituting W_a^ξ with H_a^ξ.

Proof. Let $D^\xi(a) = \{c \in A \mid c \cdot a \geq \xi\}$. By definition, $A \bigcap_{a \in D} W_a^\xi = \bigcap_{a \in D} D^\xi(a)$ and $A \bigcap_{k=1}^r W_a^\xi = \bigcap_{k=1}^r D^\xi(a^k)$, so it remains to show that $\bigcap_{a \in D} D^\xi(a) = \bigcap_{k=1}^r D^\xi(a^k)$. Since $\left\{a^k\right\}_{k=1}^r \subseteq D$, $\bigcap_{a \in D} D^\xi(a) \subseteq \bigcap_{k=1}^r D^\xi(a^k)$. On the other hand, let $c \in \bigcap_{k=1}^r D^\xi(a^k)$. Then $\forall k = 1, ..., r$, $c \cdot a^k \geq \xi$. Let $a \in D$, so $a = \sum_{k=1}^r \lambda_k a^k$ where $\lambda_k \in [0,1]$, $\sum_{k=1}^r \lambda_r = 1$. Then $c \cdot a = c \cdot \sum_{k=1}^r \lambda_k a^k = \sum_{k=1}^r \lambda_k \left(c \cdot a^k\right) \geq \sum_{k=1}^r \lambda_k \xi = \xi$, hence $c \in \bigcap_{a \in D} D^\xi(a)$, i.e. $\bigcap_{k=1}^r D^\xi(a^k) \subseteq \bigcap_{a \in D} D^\xi(a)$. The same arguments hold when substituting W_a^ξ with H_a^ξ and $\geq$ with $=$. ■

To show that Algorithm 3.2 provides the desired updated beliefs, the following notation is useful.

Notation A.1 *Given E and $h \in B$, define the set of feasible acts comparable and indifferent to h to be*

$$J_B^h = \{f \in B \mid f \sim h \text{ and } f = h \text{ on } E^c\}.$$

Proof of Proposition 3.2. HK showed (p.276, Proposition 5) that $U^{DC \cap PFI \max}$ is defined by $C_{E,g,B} = \{q_E \mid q \in C$ and $\int (u \circ f)\, dq_E \geq \min_{p \in K_E^{E,g,B}} \int (u \circ g)\, dp$ for all $f \in J_B^g\}$, where $K_E^{E,g,B}$ is the set of Bayesian conditionals on E of $K^{E,g,B}$, the measures supporting the conditional optimality of all acts initially optimal and comparable to g, defined by $K^{E,g,B} \equiv \{q \in C \mid \int (u \circ g)\, dq_E \geq \int (u \circ f)\, dq_E$ for all $f \in B$ with $f = g$ on E^c and $\int (u \circ g)\, dq_E = \int (u \circ f)\, dq_E$ if, in addition, $f \sim g\}$. We first show that $u \circ J_B^g = \{a \in u \circ B \mid a = b$ on E^c, $\min_{q \in C} \sum_{s \in S} a_s q_s = o^* = \sum_{s \in S} a_s q_s^g$, where $q^g \in Q^{E,g,B} \cap \arg\min_{q \in C} \sum_{s \in S} b_s q_s\}$ (see the proof of Proposition 3.1 for the definition of $Q^{E,g,B}$). To see this, note that if a is an element of the r.h.s then $a \in u \circ J_B^g$, because $a \in u \circ \{f \in B \mid f = g$ on $E^c\}$ and $\min_{q \in C} \sum_{s \in S} a_s q_s = \min_{q \in C} \sum_{s \in S} b_s q_s$. For the other direction, suppose $a \in u \circ J_B^g$. Existence of $q^g \in Q^{E,g,B} \cap \arg\min_{q \in C} \sum_{s \in S} b_s q_s$ is guaranteed by Lemma A.1 of HK stated below. Since $a \in u \circ J_B^g$, $\sum_{s \in S} b_s q_s^g \geq \sum_{s \in S} a_s q_s^g$ by definition of $Q^{E,g,B}$. Thus $\min_{q \in C} \sum_{s \in S} a_s q_s =$

$o^* = \sum_{s\in S} b_s q_s^g = \sum_{s\in S} a_s q_s^g$ and so a is an element of the r.h.s. In step **3.2.C**, the first intersection produces the measures in $\arg\min_{q\in C} \sum_{s\in S} b_s q_s$, and the remaining intersections ensure membership in $Q^{E,g,B}$. Given $q^g \in Q^{E,g,B} \cap \arg\min_{q\in C} \sum_{s\in S} b_s q_s$, the first intersection in step **3.2.D** results in the set $\hat{J} \equiv ext(\{a \in u \circ B \mid f = g \text{ on } E^c \text{ and } \sum_{s\in S} a_s q_s^g = o^*\})$. The remaining intersections in step **3.2.D** result in the set $\bar{J} = ext[co(\hat{J} \cap \{a \in \mathbb{R}^{|S|} \mid \forall q \in C,\ \sum_{s\in S} a_s (q_s - q_s^g) \geq 0\})] = ext\,(u \circ J_B^g)$. The fact that in step **3.2.D**, it suffices to apply a finite number of intersections ($\forall q \in \bar{C}$ and not $\forall q \in C$) follows from Proposition A.1. Next observe that since $K_E^{E,g,B} = Q_E^{E,g,B} \cap \{q \in C_E \mid \sum_{s\in S} (a_s - b_s)\, q_s = 0,\ \forall a \in u \circ J_B^g\}$, $\bar{K}_E = ext\left(K_E^{E,g,B}\right)$. Again, by Proposition A.1, it is sufficient in step **3.2.E** to compute a finite number of intersections ($\forall a \in \bar{J}$ and not $\forall a \in u \circ J_B^g$). Finally, by definition of $U^{DC\cap PFI\max}$ and Proposition A.1, $C_{E,g,B}$ is the intersection of C_E and the sets $\left\{q \in C_E \mid \sum_{s\in S} (a_s - \beta_2)\, q\,(s) \geq 0\right\}$ for all $a \in \bar{J}$. ■

Lemma A.1 *(HK, p.288, Lemma 3) For $(\succsim, E, g, B) \in \mathcal{T}$, $Q^{E,g,B} \cap \arg\min_{q\in C} \int (u \circ g)\, dq \neq \emptyset$.*

The following proposition is needed for the proof of Proposition 3.3.

Definition A.1 *Let*

$$R^{E,g,B} = Q^{E,g,B} \cap \arg\min_{q\in C} \int (u \circ g) dq,$$

and $R_E^{E,g,B}$ be the set of Bayesian conditionals on E of measures in $R^{E,g,B}$ (see the proof of Proposition 3.1 for the definition of $Q^{E,g,B}$).

Notation A.2 *Given E and $h \in \mathcal{A}$, define the set of acts comparable and indifferent to h to be*

$$J^h = \{f \in \mathcal{A} \mid f \sim h \text{ and } f = h \text{ on } E^c\}.$$

Proposition A.2 *$U^{DC\cap RAmax}$ is the update rule in $\mathcal{U}^{Bayes}$ such that*[5]

$$C_{E,g,B} = \{q_E \mid q \in C \text{ and } \int (u \circ f) dq_E \geq \min_{p \in R_E^{E,g,B}} \int (u \circ g) dp \text{ for all } f \in J^g\}.$$

Moreover, if $u(X) = \mathbb{R}$,

$$C_{E,g,B} = \overline{co} \left\{ \begin{array}{c} q_E \mid q \in \arg\min_{p\in C} \int (u \circ h)\, dp \text{ for some } h \in J^g \\ \text{and, for that } h,\ \int (u \circ h) dq_E \geq \min_{p\in R_E^{E,g,B}} \int (u \circ g) dp \end{array} \right\}.$$

[5]Without the assumption $u(X) = \mathbb{R}$, the set $\arg\min_{p\in C} \int (u \circ g) dp$ should be replaced everywhere with its superset $\left\{q \in C \mid \int (u \circ f) dq \geq \int (u \circ g) dq \text{ for all } f \in J^g\right\}$. As can be shown (proof available from the authors upon request), under the assumption $u(X) = \mathbb{R}$, the Bayesian conditionals formed from the two sets are the same. Notice that the latter set depends on E while the former does not.

Definition A.2 *An act f' is* convexly related *to an act f on E if $\exists x \in X$, $\lambda \in [0,1]$ such that $f' = \lambda f + (1-\lambda)x$ or $f = \lambda f' + (1-\lambda)x$ on E.*

Proof of Proposition A.2. The proof for the first characterization follows an argument similar to the one in the proof of Proposition 5 in HK (p. 291) with $R^{E,g,B}$ playing the role of $K^{E,g,B}$. To prove the second characterization, let $D^0_{E,g,B} = \{q_E \mid q \in \arg\min_{p\in C} \int (u \circ f)\, dp$ for some $f \in J^g$ and, for that f, $\int (u\circ f) dq_E \geq \min_{p\in R_E^{E,g,B}} \int (u \circ g) dp\}$. Let $C^0_{E,g,B} = \overline{co}(D^0_{E,g,B})$. For any $r \in R^{E,g,B}$, $r_E \in D^0_{E,g,B}$. Consider $q_E \in D^0_{E,g,B}$ and associated act $f \in J^g$ and $q \in \arg\min_{p\in C} \int (u\circ f)\, dp$ such that $\int (u \circ f) dq_E \geq \min_{p\in R_E^{E,g,B}} \int (u\circ g) dp$. For such q and f and any $h \in J^g$,

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

thus $\int (u\circ h) dq_E \geq \int (u \circ f)\, dq_E \geq \min_{p\in R_E^{E,g,B}} \int (u\circ g) dp$. Since these inequalities are preserved under convex combinations and closure, $C^0_{E,g,B} \subseteq C_{E,g,B}$ and $\min_{p\in C^0_{E,g,B}} \int (u \circ g) dp = \min_{p\in C_{E,g,B}} \int (u\circ g) dp$. Moreover, we show that $C_{E,g,B} \subseteq C^0_{E,g,B}$. Suppose that $C^0_{E,g,B} \neq C_{E,g,B}$. Since both sets are convex, there exists $\hat{q} \in C_{E,g,B} \setminus C^0_{E,g,B}$ on the boundary of $C_{E,g,B}$. Thus there exists $\hat{f} \in J^g$ such that $\hat{q} \in \arg\min_{p\in C_{E,g,B}} \int (u \circ \hat{f}) dp$. Since $\hat{q} \notin C^0_{E,g,B}$, there exists such $\hat{f}$ for which $\int (u \circ \hat{f}) d\hat{q} < \min_{p \in C^0_{E,g,B}} \int (u\circ \hat{f}) dp$. If there exists $h \in J^g$ convexly related to $\hat{f}$ on E such that $\min_{p\in C^0_{E,g,B}} \int (u\circ h) dp = \min_{p \in C_{E,g,B}} \int (u \circ g) dp$ then $\int (u\circ h) d\hat{q} < \min_{p\in C^0_{E,g,B}} \int (u\circ h) dp = \min_{p\in C_{E,g,B}} \int (u \circ g) dp$, contradicting $\hat{q} \in C_{E,g,B}$. Thus for any $h \in J^g$ convexly related to $\hat{f}$ on E, $\min_{p\in C^0_{E,g,B}} \int (u\circ h) dp > \min_{p\in C_{E,g,B}} \int (u\circ g) dp$. Therefore for any such h and $q^* \in \arg\min_{p\in C} \int (u \circ h) dp$, $\int (u\circ h) dq^*_E > \min_{p \in C_{E,g,B}} \int (u \circ g) dp$. Thus $q^*_E \in D^0_{E,g,B}$. It follows that $\min_{p\in C^0_{E,g,B}} \int (u\circ h) dp = \min_{p\in C} \int (u \circ h) dp_E$ by Lemma A.2. But then, $\min_{p\in C} \int (u\circ h) dp_E > \min_{p\in C_{E,g,B}} \int (u\circ g) dp$, a contradiction. Therefore $C^0_{E,g,B} = C_{E,g,B}$, as required. ■

Lemma A.2 *Fix any $f \in J^g$ and let $C^f = \{q \in \arg\min_{p\in C} \int (u\circ h) dp \mid h \in J^g$ convexly related to f on $E\}$. If $u(X) = \mathbb{R}$, then $\inf_{p\in C^f} \int (u \circ f) dp_E = \min_{p \in C} \int (u\circ f) dp_E$.*

Proof. If $u \circ f$ is constant on E, then $C = \arg\min_{p\in C} \int (u \circ f) dp_E$ and the lemma is trivially true. For the rest of the proof, we assume $u \circ f$ non-constant on E. For such an act $f \in J^g$, define the function $\beta : (0, \infty) \to \mathbb{R}$ at each $\alpha > 0$ by the solution to $\min_{p\in C} \int (\alpha(u\circ f) + \beta(\alpha))_E (u \circ g) dp = \min_{p\in C} \int (u \circ g) dp$. Such a function is well-defined because $p(E) > 0$ for all $p \in C$, and thus the left-hand side is strictly monotonic in $\beta(\alpha)$. Now define the function $q : (0,\infty) \to C$ such that, for $\alpha > 0$, $q(\alpha) \in \arg\min_{p\in C} \int (\alpha(u\circ f) + \beta(\alpha))_E (u\circ g) dp$. We denote $q(\alpha)$ by q^α. First we show that $\int(u \circ f)dq_E^\alpha$ is non-increasing with α. Let $\alpha_1 > \alpha_2 > 0$. Then

$$\begin{aligned}\int \left(\alpha_1(u \circ f) + \beta(\alpha_1)\right)_E (u \circ g)dq^{\alpha_1} &= \int \left(\alpha_2(u \circ f) + \beta(\alpha_2)\right)_E (u \circ g)dq^{\alpha_2} \\ &\leq \int \left(\alpha_2(u \circ f) + \beta(\alpha_2)\right)_E (u \circ g)dq^{\alpha_1},\end{aligned}$$

so $(\alpha_1 - \alpha_2)\int_E(u \circ f)dq^{\alpha_1} \leq [\beta(\alpha_2) - \beta(\alpha_1)]q^{\alpha_1}(E)$. Similarly, $(\alpha_2 - \alpha_1)\int_E(u \circ f)dq^{\alpha_2} \leq [\beta(\alpha_1) - \beta(\alpha_2)]q^{\alpha_2}(E)$. Thus $\int(u \circ f)dq_E^{\alpha_2} \geq \frac{\beta(\alpha_2)-\beta(\alpha_1)}{\alpha_1-\alpha_2} \geq \int(u \circ f)dq_E^{\alpha_1}$, establishing that $\int(u \circ f)dq_E^\alpha$ is non-increasing with α. By observing that it is bounded, it follows that $\inf_{\alpha>0}\{\int(u \circ f)dq_E^\alpha\}$ exists and equals $\lim_{\alpha\to\infty}\{\int(u \circ f)dq_E^\alpha\}$.
By definition of $\beta(\alpha)$ and q^α, for $\alpha > 0$,

$$\begin{aligned}0 &= \frac{1}{\alpha q^\alpha(E)}\left[\min_{p\in C}\int \left(\alpha(u \circ f) + \beta(\alpha)\right)_E (u \circ g)dp - \min_{p\in C}\int (u \circ g)dp\right] \\ &= \int(u \circ f)dq_E^\alpha + \frac{\beta(\alpha)}{\alpha} + \frac{1}{\alpha q^\alpha(E)}[\int_{E^c}(u \circ g)dq^\alpha - \min_{p\in C}\int(u \circ g)dp].\end{aligned}$$

From this, $\lim_{\alpha\to\infty}\{\int(u\circ f)dq_E^\alpha\} = \lim_{\alpha\to\infty}\{-\frac{\beta(\alpha)}{\alpha} - \frac{1}{\alpha q^\alpha(E)}[\int_{E^c}(u\circ g)dq^\alpha - \min_{p\in C}\int(u\circ g)dp]\}$. Since $\int_{E^c}(u \circ g)dq^\alpha$ is bounded, $\lim_{\alpha\to\infty}\{\frac{1}{\alpha q^\alpha(E)}[\int_{E^c}(u \circ g)dq^\alpha - \min_{p\in C}\int(u \circ g)dp]\} = 0$, and therefore, $\lim_{\alpha\to\infty}\{\int(u \circ f)dq_E^\alpha\} = \lim_{\alpha\to\infty}\{-\frac{\beta(\alpha)}{\alpha}\}$. From $\min_{p\in C}\int \left(\alpha(u \circ f) + \beta(\alpha)\right)_E (u \circ g)dp = \min_{p\in C}\int(u \circ g)dp$, subtracting $\beta(\alpha)$ from each state and dividing by α gives

$$\min_{p\in C}\int (u \circ f)_E \left(\frac{(u \circ g) - \beta(\alpha)}{\alpha}\right) dp = \frac{1}{\alpha}[\min_{p\in C}\int (u \circ g)dp - \beta(\alpha)]. \tag{A.1}$$

Taking the limit as $\alpha \to \infty$ on both sides of (A.1),

$$\min_{p\in C}\int (u \circ f)_E(\lim_{\alpha\to\infty}\{-\frac{\beta(\alpha)}{\alpha}\})dp = \lim_{\alpha\to\infty}\{-\frac{\beta(\alpha)}{\alpha}\},$$

so

$$\lim_{\alpha\to\infty}\{-\frac{\beta(\alpha)}{\alpha}\} = \min_{p\in C}\int (u \circ f)dp_E.$$

Hence $\inf_{\alpha>0}\{\int(u \circ f)dq_E^\alpha\} = \min_{p\in C}\int(u \circ f)dp_E$. If $u(X) = \mathbb{R}$, then acts h^α such that $u \circ h^\alpha = \left(\alpha(u \circ f) + \beta(\alpha)\right)_E (u \circ g)$ exist for all $\alpha > 0$ and each such h^α is convexly related to f on E and $h^\alpha \in J^g$. Therefore, in this case, $\inf_{p\in C^f}\int(u \circ f)dp_E = \inf_{\alpha>0}\{\int(u \circ f)dq_E^\alpha\}$ and the lemma is proved. ■

Proof of Proposition 3.3. If $q(E) = 1$ for all $q \in C$, then Step **3.3.D** produces $co(\bar{C}_E) = co(\bar{C})$, which by inspection of **DC** and **RA**, are the updated beliefs produced by $U^{DC\cap RAmax}$ in this case. From here on, suppose $q(E) < 1$ for some $\bar{q} \in C$. Let $x^{E^c} \in X$

be a constant act for which $u(x^{E^c}) = \int_{E^c} (u \circ g)\, d\bar{q}_{E^c}$. We need to prove that $C_{E,g,B}$ has as members the Bayesian updates q_E of all $q \in C$ satisfying (1) $q(E) \leq q^g(E)$ if $g \succ x^{E^c}$, (2) $q(E) \geq q^g(E)$ if $x^{E^c} \succ g$, or (3) no further restrictions if $g \sim x^{E^c}$. We use the second description of $C_{E,g,B}$ presented in Proposition A.2. We first show that measures q that violate (1)-(3) are not updated. Let $f \in J^g$ and let $q \in \arg\min_{p \in C} \int (u \circ f)\, dp$. Also let $\beta_3 \equiv \min_{p \in R_E^{E,g,B}} \int (u \circ g)\, dp$, i.e. $\beta_3 = \int (u \circ g)\, dq_E^g$. Observe that since $q^g \in R^{E,g,B}$, $q^g \in \arg\min_{p \in C} \int (u \circ g)\, dp$. Also note that $x^{E^c} \succsim g$ implies that $\beta_3 \leq u(x^{E^c})$, which is immediate if $q^g(E) = 1$, and follows also when $q^g(E) < 1$ because $0 \geq \int (u \circ g)\, dq^g - \int_{E^c} (u \circ g)\, dq^g_{E^c} = \int_E (u \circ g)\, dq^g - \frac{q^g(E)}{q^g(E^c)} \int_{E^c} (u \circ g)\, dq^g = q^g(E) \left(\int_E (u \circ g)\, dq^g_E - u(x^{E^c}) \right)$. Adding $f \in J^g$, we have

$$
\begin{aligned}
0 &= \min_{p \in C} \int (u \circ f)\, dp - \min_{p \in C} \int (u \circ g)\, dp \\
&= \left(\int_E (u \circ f)\, dq + \int_{E^c} (u \circ g)\, dq \right) - \left(\int_E (u \circ g)\, dq^g + \int_{E^c} (u \circ g)\, dq^g \right) \\
&= \left[q(E) \left(\int_E (u \circ f)\, dq_E - u(x^{E^c}) \right) + u(x^{E^c}) \right] - \left[q^g(E) \left(\beta_3 - u(x^{E^c}) \right) + u(x^{E^c}) \right] \\
&= q(E) \left(\int_E (u \circ f)\, dq_E - u(x^{E^c}) \right) - q^g(E) \left(\beta_3 - u(x^{E^c}) \right)
\end{aligned}
$$

by hypothesis. So $q(E) \left(\int_E (u \circ f)\, dq_E - u\left(x^{E^c}\right) \right) = q^g(E) \left(\beta_3 - u\left(x^{E^c}\right) \right) \leq 0$. Thus, if $g \sim x^{E^c}$, then $\int_E (u \circ f)\, dq_E = u\left(x^{E^c}\right) = \beta_3 = \min_{p \in R_E^{E,g,B}} \int (u \circ g)\, dp$. If $x^{E^c} \succ g$, all inequalities are strict and thus $\int_E (u \circ f)\, dq_E \geq \min_{p \in R_E^{E,g,B}} \int (u \circ g)\, dp \Leftrightarrow q(E) \geq q^g(E)$. A similar argument implies that if $g \succ x^{E^c}$, $\int_E (u \circ f)\, dq_E \geq \min_{p \in R_E^{E,g,B}} \int (u \circ g)\, dp \Leftrightarrow q(E) \leq q^g(E)$. We now show that measures q that satisfy (1)-(3) are updated. First note that for any $f \in \mathcal{A}$ such that $f = g$ on E^c and $q \in C$,

$$
\int (u \circ f)\, dq = \int_E (u \circ f)\, dq + q(E^c) u(x^{E^c}) = \int u \circ \left(f_E x^{E^c} \right) dq. \tag{A.2}
$$

If $g \sim x^{E^c}$, then by taking $f = x^{E^c}_E g$ and using equality (A.2) above we get $f \sim g$ and so $f \in J^g$, and moreover $\forall q \in C$, $q \in \arg\min_{p \in C} \int (u \circ f)\, dp$. So $\forall q \in C$, $q_E \in C_{E,g,B}$, thus establishing (3). To prove (2), Suppose that $x^{E^c} \succ g$. Note that for any $q \in C$, if there exist $\tilde{q} \in C, \lambda > 1 : \tilde{q}(F) = \lambda q(F), \forall F \subseteq E$, then q and $\tilde{q}$ have the same update, so it is sufficient to consider q such that $\nexists \tilde{q} \in C, \lambda > 1 : \tilde{q}(F) = \lambda q(F), \forall F \subseteq E$. By Lemma A.3 which is proven below, for such q, if $q(E) \geq q^g(E)$, then there exists $f \in \mathcal{A}$ constant on E^c, such that $q \in \arg\min_{q' \in co(C \cup \Delta(E^c))} \int (u \circ f)\, dq'$. So $\int (u \circ f)\, dq \leq u\left(f(E^c)\right)$ since $\Delta(E^c) \subseteq co\left(C \cup \Delta(E^c)\right)$. By closure of $C_{E,g,B}$ it is sufficient to consider the case of strict inequality, so $0 > \int \left((u \circ f) - u(f(E^c))\right) dq = b \left(I[u \circ g] - u\left(x^{E^c}\right) \right)$ for some $b > 0$. Let $f_E x^{E^c} \in \mathcal{A}$ be

such that $u \circ f_E x^{E^c} = \frac{1}{b}\left((u \circ f - f(E^c)) + u(x^{E^c})\right)$. Thus $q \in \arg\min_{q' \in C} \int u \circ (f_E x^{E^c}) dq'$ and $\int u \circ (f_E x^{E^c}) dq = I[u \circ g]$, so $f_E x^{E^c} \sim g$, which by equality (A.2) implies that $f_E g \sim g$. Denoting $f' = f_E g$, we get $f' \in J^g$. Then $\int_E (u \circ f')\, dq_E \geq \min_{p \in R_E^{E,g,B}} \int (u \circ g)\, dp$ by $q(E) \geq q^g(E)$ and the property initially proved, so $q_E \in C_{E,g,B}$.
The proof of case (1), where $g \succ x^{E^c}$, is similar with the change (in Lemma A.3 and in the rest of the proof) that q satisfies $q(E) \leq q^g(E)$ and $\nexists \tilde{q} \in C, \lambda \in [0,1) : \tilde{q}(F) = \lambda q(F), \forall F \subseteq E$, and $co\,(C \cup \Delta(E^c))$ is replaced by $C' \equiv \{\lambda q'' + (1-\lambda)\eta :\ q'' \in C,\ \eta \in \Delta(E^c),\ \lambda \geq 1\}$. Thus $\exists f' \in \mathcal{A}$ constant on E^c such that $q \in \arg\min_{q' \in C'} \int (u \circ f')\, dq'$, so $\int (u \circ f)\, dq > u\,(f(E^c))$ since f' separates C' from $\Delta(E^c)$ and then $0 < \int ((u \circ f) - u(f(E^c)))\, dq = b\left(I[u \circ g] - u\left(x^{E^c}\right)\right)$ for some $b > 0$. ■

Lemma A.3 *Let $q \in C$ such that $\nexists \tilde{q} \in C, \lambda > 1 : \tilde{q}(F) = \lambda q(F), \forall F \subseteq E$. Then there exists an act $f \in \mathcal{A}$ which is constant on E^c, such that $q \in \arg\min_{q' \in co(C \cup \Delta(E^c))} \int (u \circ f)\, dq'$.*

Proof of Lemma A.3. If $q(E) = 1$, then any $f \in \mathcal{A}$ such that $u \circ f = 0_E 1$ suffices, because $\int (u \circ f)\, dq' \geq 0$ for all $q' \in co(C \cup \Delta(E^c))$ and $\int (u \circ f)\, dq = 0$. From here on, suppose $q(E) < 1$. Consider the set $\Psi(q) \equiv \{q' : q'(F) = q(F), \forall F \subseteq E\}$ (note that $C \cap \Psi(q) \neq \emptyset$). First assume $Int(co\,(C \cup \Delta(E^c))) \cap \Psi(q) \neq \emptyset$, i.e. there exists an interior point of $co\,(C \cup \Delta(E^c))$, q'', which satisfies $q''(F) = q(F), \forall F \subseteq E$. Then $\exists \epsilon > 0 : \tilde{q} \equiv (1+\epsilon)q'' - \epsilon \left(\underbrace{0, ..., 0}_{E}, \underbrace{\frac{1}{|E^c|}, ..., \frac{1}{|E^c|}}_{E^c}\right) \in co\,(C \cup \Delta(E^c))$. Setting $\lambda = 1+\epsilon > 1$, we get $\tilde{q}(F) = \lambda q(F), \forall F \subseteq E$, which contradicts the initial condition. So $Int(co\,(C \cup \Delta(E^c))) \cap \Psi(q) = \emptyset$. Thus there exists a hyperplane H which contains $\Psi(q)$ and supports $co(C \cup \Delta(E^c))$ at any point in $co\,(C \cup \Delta(E^c)) \cap \Psi(q)$. Let us now take a normal $u \circ f$ to H. Since H contains $\Psi(q)$, it must be true that $\int (u \circ f)\, d\,(q - q') = 0$ for all $q' \in \Psi(q)$. So

$$0 = \int (u \circ f)\, dq - \int (u \circ f)\, dq' = \int_{E^c} (u \circ f)\, dq - \int_{E^c} (u \circ f)\, dq',$$

since q and q' are identical on E. Thus $\int_{E^c} (u \circ f)\, dq'$ is constant regardless of the choice of q'. Let $s \in E^c$, and let q'_s be such that $q'_s(F) = \begin{cases} q(F), & F \subseteq E \\ q(E^c), & s \in F \subseteq E^c \end{cases}$, so $q'_s \in \Psi(q)$, $\forall s \in E^c$. Thus $\forall s, \hat{s} \in E^c$, we have $\int_{E^c} (u \circ f)\, dq'_s = \int_{E^c} (u \circ f)\, dq'_{\hat{s}} \implies u \circ f(s) = u \circ f(\hat{s})$, i.e. $u \circ f$ is constant on E^c and therefore f can be chosen to be constant on E^c. ■

Algorithm A.1 (Step 3.3.D') *Fix some $\xi \in (0,1)$. Let $\bar{A} = \emptyset$. For each pair $q^1, q^2 \in \bar{C}$,* *solve the linear program,*

$$\begin{array}{ll} \max_{a\in\mathbb{R}^{|E|}} K^0 \cdot \sum_{s\in E}(a_s-\beta_3)q_s^2 & \\ s.t. & \\ \sum_{s\in S}(a_E b)_s q_s \geq o^* \qquad \forall q\in \bar{C} & \\ \sum_{s\in S}(a_E b)_s \left[\xi q_s^1+(1-\xi)q_s^2\right]=o^* & \\ \sum_{s\in E}(a_s-\beta_3)q_s^1 \geq 0, & \end{array} \tag{P*}$$

where $\beta_3 \equiv \sum_{s\in E} b_s \left(q_E^g\right)_s$ *and* $K^0 = 1$ *if* $K \geq 0$ *and* $K^0 = -1$ *if* $K < 0$, *for* $K \equiv o^* - \sum_{s\in S}[(\beta_3\mathbf{1})_E b]_s\left(\xi q_s^1+(1-\xi)q_s^2\right)$.
Consider only pairs (q^1,q^2) *for which* (P^*) *is feasible. If* (P^*) *is unbounded, then add* $\xi q^1 + (1-\xi)\, q^2$ *to* $\bar{A}$. *Otherwise, let* $a^* \in \arg\max(\mathrm{P}^*)$. *If* $\sum_{s\in E}(a_s^*-\beta_3)q_s^2 \geq 0$, *then add* q^1, q^2 *to* $\bar{A}$. *Otherwise, add* $\alpha^*_{(q^1,q^2)}q^1 + \left(1-\alpha^*_{(q^1,q^2)}\right)q^2$ *to* $\bar{A}$, *where* $\alpha^*_{(q^1,q^2)} = \frac{\sum_{s\in E}(a_s^*-\beta_3)q_s^2}{\sum_{s\in E}(a_s^*-\beta_3)(q_s^2-q_s^1)}$.
Compute $\bar{U}_E \equiv ext(\{q_E \mid q\in\bar{A})$. *Return* $co(\bar{U}_E)$.

Proposition A.3 *If* $u(X)=\mathbb{R}$, *Algorithm 3.3 with Step **3.3.D'** replacing Step **3.3.D** results in the updated set of measures,* $C_{E,g,B}$, *produced by the update rule* $U^{DC\cap RA\max}$.

The following proposition is needed for the proof of Proposition A.3. In Proposition A.2, the condition $[q\in\arg\min_{p\in C}\sum_{s\in S}a_s p_s$ for some $a\in u\circ J^g]$ will be referred to as condition 1. For q which satisfy condition 1 with the vector a, the condition $[\sum_{s\in E}a_s\left(q_E\right)_s \geq \beta_3 \equiv \min_{p\in R_E^{E,g,B}}\sum_{s\in E}b_s p_s]$ will be referred to as condition 2. The following proposition elucidates the geometric properties of the updated beliefs.

Proposition A.4 *If* $u(X)=\mathbb{R}$, $C_{E,g,B}$ *for* $U^{DC\cap RA\max}$ *equals the convex hull of Bayesian conditionals of*
(a) Extreme points of C *satisfying conditions 1 and 2.*
(b) Non-extreme Boundary points of C *which are convex combinations of two extreme points satisfying condition 1 with the same vector* a, *and for which condition 2 is satisfied with an equality.*

Proof of Proposition A.4. Denote by D the set of measures satisfying (a) or (b). Let $C^E=\{q\mid q\in\arg\min_{p\in C}\sum_{s\in S}a_s p_s$ for some $a\in u\circ J^g$ and $\sum_{s\in E}\left(a_s-\beta_3\right)q_s\geq 0\}$. Since condition 2 is equivalent to $\sum_{s\in E}\left(a_s-\beta_3\right)q_s\geq 0$, $D\subseteq C^E$, thus $co(D)\subseteq co(C^E)$. We will show that $co(D)=co(C^E)$. We first show that an updated boundary point must be a proper (positive weight) convex combination of extreme points of C satisfying condition 1 with the same a. Any such boundary point q can be expressed as a proper convex combination of a set A of extreme points of C, i.e. $q=\sum_{i\in A}\lambda_i q^i$. Suppose that q satisfies condition 1 using the vector a and assume by contradiction that $\sum_{s\in S} a_s q_s^j > \min_{p\in C} \sum_{s\in S} a_s p_s = \sum_{s\in S} a_s q_s$ for some $j \in A$. Then $\sum_{s\in S} a_s q_s' < \sum_{s\in S} a_s q_s$ for $q' \equiv \frac{1}{1-\lambda_j} \sum_{i\in A\setminus\{j\}} \lambda_i q^i$, contradicting $q \in \arg\min_{p\in C} \sum_{s\in S} a_s p_s$. Thus $q^i \in \arg\min_{p\in C} \sum_{s\in S} a_s p_s$ for all $i \in A$. Next we show that for each edge (convex hull of two extreme points entirely on the boundary) of C, the subset of updated points in this edge is convex. By the arguments above, it is sufficient to consider the case where both the edge's vertices satisfy condition 1. Let R be an edge of C, such that one of its vertices, q^1, satisfies conditions 1 and 2, and the other, q^2, satisfies only condition 1. If no point in R other than q^1, q^2 satisfies condition 1, then only q^1 gets updated, so this forms a convex set. Now suppose that there exists a point in R other than q^1, q^2 satisfying condition 1. Let $P(R)$ denote the set $\{a \in u \circ \{f \in \mathcal{A} \text{ with } f = g \text{ on } E^c\} \mid q' \in \arg\min_{p\in C} \sum_{s\in S} a_s p_s$ for all q' in $R\}$. By the arguments above, $P(R)$ is non-empty. For each $a \in P(R)$, there exists a half-space $\sum_{s\in E} (a_s - \beta_3)\, q_s \geq 0$ corresponding to condition 2. This half-space does not contain the entire edge R, for if it did, every point in R would get updated, contrary to our assumption. So the intersection between the half-space and R must be a set of the form $\{\alpha q^1 + (1-\alpha) q^2 \mid \alpha_R^a \leq \alpha \leq 1\}$ (the weak inequality follows from closure of $C_{E,g,B}$), where $\alpha_R^a > 0$ is the weight satisfying $\sum_{s\in E} (a_s - \beta_3)\left(\alpha_R^a q_s^1 + (1-\alpha_R^a) q_s^2\right) = 0$. Each vector $a \in P(R)$ has a unique corresponding weight α_R^a. Thus the set of points in R which are updated is $\{\alpha q^1 + (1-\alpha) q^2 \mid \alpha_R^* \leq \alpha \leq 1\}$, where $\alpha_R^* \equiv \inf_{a\in P(R)} \alpha_R^a$. This set is convex. Now consider a facet Q (convex hull of at least three extreme points entirely on the boundary) of C such that all of its vertices satisfy condition 1, and some vertices, but not all, satisfy condition 2. Let $D_1(Q)$ denote the set of updated extreme points of Q, and let $D_2(Q)$ denote the set of points of the form $\alpha_R^* q^1 + (1-\alpha_R^*) q^2$ in Q. We show that the set of points of Q which are updated is exactly $co\,(D_2(Q) \cup D_1(Q))$, implying that $co\,(D) = co(C^E)$. As we have seen above, it is impossible for a point to be updated, which is of the form $\alpha q^1 + (1-\alpha) q^2$, where q^1 and q^2 are the vertices of an edge R of Q, q^2 does not satisfy condition 2, and $\alpha < \alpha_R^*$, so assume $int\,[Q \setminus co\,(D_2(Q) \cup D_1(Q))]$ is non-empty. Consider a point $q \in int\,[Q \setminus co\,(D_2(Q) \cup D_1(Q))]$, and assume by contradiction that q is updated. Then there exists a vector $a \in u \circ \{f \in B \mid f = g \text{ on } E^c\}$ such that $q \in \arg\min_{p\in C} \sum_{s\in S} a_s p_s$ and $\sum_{s\in E} (a_s - \beta_3)\, q_s \geq 0$. It follows from the same argument used earlier, that any $q' \in Q$ also belongs to $\arg\min_{p\in C} \sum_{s\in S} a_s p_s$. Moreover, the half-space corresponding to condition 2 intersects any edge R of Q in $co\,(D_2(Q) \cup D_1(Q))$, so the half-space intersects the whole facet Q in $co\,(D_2(Q) \cup D_1(Q))$. But this contradicts the assumption on q, completing the proof. ■

Proof of Proposition A.3. By construction of q_E^g, $\beta_3 = \min_{p\in R_E^{E,g,B}} \int (u \circ g)\, dp$. Denote $\sum_{s\in E} (a_s - \beta_3) q_s^1$ by $\theta_1(a)$ and $\sum_{s\in E} (a_s - \beta_3) q_s^2$ by $\theta_2(a)$. The following is proved based on Proposition A.4.

(1) If $q^1 = q^2$, we show that (P*) is bounded and moreover $q^1 = q^2$ is updated if and only if (P*) is feasible with non-negative value. Since the second constraint in (P*) implies that $K = \theta_2(a)$, $\theta_2(a)$ is fixed, so the program is bounded. Assume it is feasible. The first and second constraints mean that there exists a vector $a_E b$ such that $q^1 = q^2 \in \arg\min_{p\in C} \sum_{s\in S} (a_E b)_s p_s$. The second constraint means that $\min_{q\in C} \sum_{s\in S} (a_E b)_s q_s = o^*$, hence $a_E b \in u \circ J^g$ and condition 1 holds. The third constraint means that condition 2 holds for $q^1 = q^2$, thus $q^1 = q^2$ is updated. On the other hand, if $q^1 = q^2$ is updated, it follows immediately from conditions 1 and 2 that (P*) is feasible.

(2) If $q^1 \neq q^2$, there are several mutually exclusive and exhaustive cases:

a. If q^1, q^2 fail to satisfy condition 1 with the same a (whether or not they are updated), we show that (P*) is infeasible. Assume by contradiction that (P*) is feasible. Then verity of the first and second constraints means that there exists a vector $a_E b \in u \circ \{f \in \mathcal{A}$ with $f = g$ on $E^c\}$ such that $\xi q^1 + (1-\xi) q^2 \in \arg\min_{p\in C} \sum_{s\in S} (a_E b)_s p_s$, and so due to arguments shown in the proof of Proposition A.4, both q^1, q^2 belong to $\arg\min_{p\in C} \sum_{s\in S} (a_E b)_s p_s$. Verity of the second constraint means that $\min_{p\in C} \sum_{s\in S} (a_E b)_s p_s = o^*$, so q^1, q^2 satisfy condition 1 with the same a - a contradiction.

b. If only q^2 is updated and q^1 satisfies condition 1 with equal a but violates condition 2, then (P*) is infeasible. This follows again from contradiction of the third constraint.

c. If only q^1 is updated and q^2 satisfies condition 1 with equal a but violates condition 2, observe the following. First, by the same arguments as above, (P*) is feasible. We minimize α_R^a for $R = co\{q^1, q^2\}$ (see the proof of Proposition A.4). Since $K = \xi\theta_1(a) + (1-\xi)\theta_2(a)$, we have $\theta_1(a) = \frac{1}{\xi}(K - (1-\xi)\theta_2(a))$. Thus $\sum_{s\in E} (a_s - \beta_3)(\alpha_R^a q_s^1 + (1-\alpha_R^a) q_s^2) = 0$ implies

$$\alpha_R^a = \frac{\sum_{s\in E}(a_s - \beta_3) q_s^2}{\sum_{s\in E}(a_s - \beta_3)(q_s^2 - q_s^1)} = \frac{\theta_2(a)}{\theta_2(a) - \theta_1(a)} = \frac{\theta_2(a)}{\theta_2(a) - \frac{K}{\xi} + \frac{1-\xi}{\xi}\theta_2(a)} = \frac{\xi\theta_2(a)}{\theta_2(a) - K}.$$

By the assumption on q^1 and q^2, for all feasible a, $\theta_1(a) \geq 0$ and $\theta_2(a) < 0$. Consequently K could be of any sign, depending on the choice of ξ, so the behavior of the above expression as a function of $\theta_2(a)$ depends on the sign of K. If $K \geq 0$ (and so $K^0 = 1$), then the function is non-increasing, so minimization of α_R^a is equivalent to maximization of $K^0 \cdot \theta_2(a)$. If $K < 0$ (thus $K^0 = -1$), the function is increasing, so minimization of α_R^a is equivalent to minimization of $\theta_2(a)$, or, again, maximization of $K^0 \cdot \theta_2(a)$. Observe that $\alpha_R^a \in (0, \xi]$ for all feasible a iff $K \geq 0$, in which case the problem is bounded from above by 0, otherwise there would exist a' for which $\theta_2(a') \geq 0$, contradicting the assumption on q^2. Similarly, $\alpha_R^a \in (\xi, 1]$ for all feasible a iff $K < 0$. If the problem is not bounded, $\theta_1(a)$ tends to ∞ and $\theta_2(a)$ tends to $-\infty$. Then $\frac{\xi\theta_2(a)}{\theta_2(a)-K}$ tends to ξ (from above), which means that $\alpha^*_{(q^1,q^2)} \equiv \inf_{a\in P(q^1,q^2)} \alpha_R^a = \xi$.

d. If q^1, q^2 are updated with equal a, we show that $K \geq 0$, and (P*) is feasible and bounded with non-negative value. To see that, first observe that $K = \xi\theta_1(a) + (1-\xi)\theta_2(a)$, and $\theta_1(a), \theta_2(a) \geq 0$, so $K \geq 0$. Since q^1 and q^2 satisfy condition 1 with the same a, and since $a_E b \in u \circ J^g$, the second constraint holds. The fact that both q^1, q^2 belong to $\arg\min_{p \in C} \sum_{s\in S} (a_E b)_s p_s$ implies $\sum_{s\in S} (a_E b)_s q^1_s \leq \sum_{s\in S} (a_E b)_s q^i_s, \quad \forall q^i \in \bar{C}$ and $\sum_{s\in S} (a_E b)_s q^2_s \leq \sum_{s\in S} (a_E b)_s q^i_s, \ \forall q^i$, so $\sum_{s\in S} (a_E b)_s (\xi q^1 + (1-\xi) q^2)_s \leq \sum_{s\in S} (a_E b)_s q^i_s, \ \forall q^i$, i.e. the first constraint holds. The third constraint holds because q^1 is updated. Thus a is feasible and so (P*) is feasible. Since the second constraint holds, $K = \xi\theta_1(a) + (1-\xi)\theta_2(a)$, and since $\theta_1(a) \geq 0$, $\theta_2(a)$ is bounded from above, and so the program is bounded. Note that since $K \geq 0$, $K^0 = 1$ and the program results in a^* such that $\sum_{s\in E}(a^*_s - \beta_3)q^2_s \geq 0$. ■

B Appendix: algorithms from polyhedral theory

This appendix includes several algorithms, known from polyhedral theory (Goodman and O'Rourke 2004), that can be used within the algorithms for computing updated beliefs.

B.1 The irredundancy problem; or, extracting a set's extreme points

Given a finite set $H \subseteq \mathbb{R}^{|S|}$, the following algorithm can be used to compute the set $ext[co(H)]$.

Algorithm B.1 *If $|H| \leq 2$, return H. Otherwise, for each $h \in H$, apply the linear feasibility program, with variable $\alpha \in [0,1]^{|H|}$,*

$$\begin{array}{ll} \sum_{\bar{h} \in H\setminus h} \alpha_{\bar{h}} \bar{h}_s = h_s & , \ \forall s \in S \\ \sum_{\bar{h} \in H\setminus h} \alpha_{\bar{h}} = 1, & \end{array} \qquad (P(h))$$

and return $\{h \in H \mid P(h) \text{ is infeasible}\}$.

B.2 A general algorithm for computing an intersection

Let $A \subseteq \mathbb{R}^{|S|}$ be a convex, compact set such that $ext(A)$ is finite. The following algorithm finds the set of extreme points of the intersection of A with a finite collection of half spaces $\{W^{\xi^k}_{a^k}\}_{k=1}^r$.

Algorithm B.2 *Set $k = 1$, $\bar{D}^k = ext(A)$. Apply the following subroutine:*
(i) Let $D^k_1 = \left\{c \in \bar{D}^k \mid a^k \cdot c \geq \xi^k\right\}$, $D^k_2 = \{c = \frac{\xi^k - a^k \cdot c^2}{a^k \cdot (c^1 - c^2)}(c^1 - c^2) + c^2 \mid (c^1, c^2) \in (\bar{D}^k)^2, \left(a^k \cdot c^1 - \xi^k\right) \cdot \left(a^k \cdot c^2 - \xi^k\right) < 0\}$.

(ii) Find $\bar{E}^k \equiv ext[co\left(D_1^k \cup D_2^k\right)]$ *using Algorithm B.1.*
(iii) If $k = r$ *stop and return* $\bar{E} = \bar{E}^r$*. Otherwise let* $\bar{D}^{k+1} = \bar{E}^k$*,* $k = k + 1$*. Return to (i).*

Lemma B.1 *Let* A *and* D *be convex sets such that* $A \cap D$ *is compact. Then any extreme point of* $A \cap D$ *is either an extreme point of* A *or* D*, or it belongs to the boundaries of both* A *and* D*.*

Proof. Let c be an extreme point of $A \cap D$ for which this is false. Since c does not belong to the boundary of one of the sets (say D), there exists a neighborhood $B_\varepsilon(c)$ of c such that $B_\varepsilon(c) \subset D$. Since c is not an extreme point of A, there exist $c^1, c^2 \in A \cap B_\varepsilon(c)$ such that c is a convex combination of c^1 and c^2. But $A \cap B_\varepsilon(c) \subset A \cap D$, so c cannot be an extreme point of $A \cap D$ - a contradiction. ■

Proposition B.1 *Algorithm B.2 finds the set of extreme points of* $A \cap_{k=1}^r W_{a^k}^{\xi^k}$*.*

Proof. It is sufficient to show that iteration k results in the set of extreme points of $A \cap_{j=1}^k W_{a^j}^{\xi^j}$. Consider $\bar{D}^k$, the set of extreme points of $A \cap_{j=1}^{k-1} W_{a^j}^{\xi^j}$, and consider the half-space $W_{a^k}^{\xi^k}$. For all $c \in D_2^k$, $a^k \cdot c = \frac{\xi^k - a^k \cdot c^2}{a^k \cdot (c^1 - c^2)} a^k \cdot (c^1 - c^2) + a^k \cdot c^2 = \xi^k$. Thus members of D_2^k lie on the boundary of $A \cap_{j=1}^k W_{a^j}^{\xi^j}$. In addition, any $q \in D_1^k$ is clearly in $A \cap_{j=1}^k W_{a^j}^{\xi^j}$, so $\bar{E}^k \subseteq A \cap_{j=1}^k W_{a^j}^{\xi^j}$. Since $A \cap_{j=1}^k W_{a^j}^{\xi^j}$ is a convex set, it follows that $co(\bar{E}^k) \subseteq A \cap_{j=1}^k W_{a^j}^{\xi^j}$. Assume by contradiction that $A \cap_{j=1}^k W_{a^j}^{\xi^j} \nsubseteq co(\bar{E}^k)$. Since both $A \cap_{j=1}^k W_{a^j}^{\xi^j}$ and $co(\bar{E}^k)$ are compact, there exists an extreme point c' of $A \cap_{j=1}^k W_{a^j}^{\xi^j}$ such that $c' \notin co(\bar{E}^k)$. c' cannot be an extreme point of $A \cap_{j=1}^{k-1} W_{a^j}^{\xi^j}$, since then it would have belonged to $D_1^k \subseteq \bar{E}^k$. Thus, by Lemma B.1, c' belongs to the intersection of the boundaries of $A \cap_{j=1}^{k-1} W_{a^j}^{\xi^j}$ and $W_{a^k}^{\xi^k}$, so it must belong to the hyperplane $H_{a^k}^{\xi^k} \equiv \left\{c \in \mathbb{R}^{|S|} \mid a^k \cdot c = \xi^k\right\}$. c' cannot be a convex combination of two points in $\bar{D}^k$, because then it would have belonged to D_2^k. Moreover, $c' \notin co(\bar{D}^k)$, because then $H_{a^k}^{\xi^k}$ would include other such combinations which are therefore in $A \cap_{j=1}^k W_{a^j}^{\xi^j}$, c' being their convex combination - a contradiction. Thus $c' \notin A \cap_{j=1}^{k-1} W_{a^j}^{\xi^j}$, a contradiction. It follows that $A \cap_{j=1}^k W_{a^j}^{\xi^j} \subseteq co(\bar{E}^k)$, thus $co(\bar{E}^k) = A \cap_{j=1}^k W_{a^j}^{\xi^j}$. Since $\bar{E}^k = ext[co(\bar{E}^k)]$, $\bar{E}^k = ext\left(A \cap_{j=1}^k W_{a^j}^{\xi^j}\right)$. Note that it is possible that the intersection sought is empty, in which case D_1^k and D_2^k are empty for all iteration starting from some iteration k'. ■

Corollary B.1 *To find the set of extreme points of* $A \cap_{k=1}^r H_{a^k}^{\xi^k}$*, Algorithm B.2 can be used with the change that* $D_1^k = \left\{c \in \bar{D}^k \mid a^k \cdot c = \xi^k\right\}$*. Alternatively, one can use Algorithm B.2 to find the set of extreme points of* $A \cap_{k=1}^r W_{a^k}^{\xi^k} \cap_{k=1}^r W_{-a^k}^{\xi^k}$*.*

References

Anscombe F, Aumann R. A Definition of Subjective Probability. Annals of Mathematical Statistics 1963;34; 199-025.

Dempster AP. A Generalization of Bayesian Inference. Journal of the Royal Statistical Society, Series B (Methodological) 1968;30; 205-247.

Ellsberg D. Risk, Ambiguity and the Savage Axioms. Quarterly Journal of Economics 1961;75 643-669.

Epstein LG, Schneider M. Recursive Multiple Priors. Journal of Economic Theory 2003;113; 1-31.

Fagin, R, Halpern J. A New Approach to Updating Beliefs. In: Uncertainty in Artifical Intelligence 6 (Piero P. Bonissone, Max Henrion, Laveen N. Kanal, and John F. Lemmer, eds.). North-Holland: Amsterdam; 1991. p. 347–374.

Gilboa I, Schmeidler D. Maxmin Expected Utility with Non-unique Prior. Journal of Mathematical Economics 1989;97; 141-153.

Gilboa I, Schmeidler D. Updating Ambiguous Beliefs. Journal of Economic Theory 1993;59; 33-49.

Goodman JE, O'Rourke J. Handbook of Discrete and Computational Geometry. CRC Press: New York; 2004.

Hacking I. Slightly More Realistic Personal Probability. Philosophy of Science 1967;34(4); 311-325.

Hanany E, Klibanoff P. Updating Preferences with Multiple Priors. Theoretical Economics 2007;2; 261-298.

Hanany E, Klibanoff P. Updating Ambiguity Averse Preferences. The B.E. Journal of Theoretical Economics (Advances) 2009;9(1); Article 37.

Hazen GB. Ambiguity Aversion and Ambiguity Content in Decision Making Under Uncertainty. Annals of Operations Research 1989;19; 415-434.

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

Jaffray J.-Y. Dynamic Decision Making with Belief Functions. In: Yager RR. et al. (Eds), Advances in the Dempster-Shafer Theory of Evidence. Wiley: New York; 1994. p. 331-352.

Jaffray J.-Y. Rational decision making with imprecise probabilities. In De Cooman G et al. (Eds), Proceedings of the First International Symposium on Imprecise Probabilities and Their Applications. Morgan Kaufmann: San Francisco; 1999. p. 183–188.
Available at http://decsai.ugr.es/~smc/isipta99/proc/027.html.

Klibanoff P, Marinacci M, Mukerji S. A Smooth Model of Decision Making under Ambiguity. Econometrica 2005;73; 1849-1892.

Kriegler E. Updating under unknown unknowns: An extension of Bayes' rule. International Journal of Approximate Reasoning 2009;50(4); 583-596.

Maccheroni F, Marinacci M, Rustichini A. Ambiguity Aversion, Robustness, and the Variational Representation of Preferences. Econometrica 2006;74; 1447-1498.

Modica S. Unawareness, priors and posteriors. Decisions in Economics and Finance 2008;31; 81–94.

Mukerji S, Tallon JM. An overview of economic applications of David Schmeidler's models of decision making under uncertainty. In: Gilboa I (Ed), Uncertainty in Economic Theory. Routledge: London and New York; 2004.

Nau R. Uncertainty Aversion with Second-Order Utilities and Probabilities. Management Science 2006;52(1); 136-145.

Nielsen TD, Jaffray J.-Y. Dynamic decision making without expected utility: An operational approach. European Journal of Operational Research 2006;169; 226–246.

Pires CP. A Rule for Updating Ambiguous Beliefs. Theory and Decision 2002;53; 137-152.

Sarin R, Wakker P. Dynamic Choice and Nonexpected Utility. Journal of Risk and Uncertainty 1998;17; 87-119.

Schipper B.C. Awareness-Dependent Subjective Expected Utility. University of California Davis 2010; mimeo.

Schmeidler D. Subjective Probability and Expected Utility without Additivity. Econometrica 1989;57; 571-587.

Shafer G. A Mathematical Theory of Evidence. Univ Princeton Press: New Jersey; 1976.

Siniscalchi M. Bayesian updating for general maxmin expected utility preferences. Northwestern University CMS-EMS Discussion Paper #1366; 2001.

Tversky A, Kahneman D. Advances in prospect theory: Cumulative representation of uncertainty. Journal of Risk and Uncertainty 1992;5; 297-323.

Wakker P., Nonexpected Utility as Aversion of Information. Journal of Behavioral Decision Making 1988;1; 169-175.

Walliser B., Zwirn D. Change rules for hierarchical beliefs. International Journal of Approximate Reasoning 2011;52(2); 166-183.

Walley P. Statistical Reasoning with Imprecise Probabilities. Chapman and hall: London; 1991.

Wang T. Conditional Preferences and Updating. Journal of Economic Theory 2003;108; 286-321.

Wasserman LA, Kadane J. Bayes' Theorem for Choquet Capacities. Annals of Statistics 1990;18; 1328-1339.

Winkler RL. An Introduction to Bayesian Inference and Decision. Holt, Rinehart and Winston; 1972.