Queen's Economics Department Working Paper No. 1303

Probabilistic Sophistication and Reverse Bayesianism

Edi Karni
Johns Hopkins University

Marie-Louise VierÃÿ
Queen's University

Department of Economics
Queen's University
94 University Avenue
Kingston, Ontario, Canada
K7L 3N6

2-2013

Probabilistic Sophistication and Reverse Bayesianism *

Edi Karni †
Johns Hopkins University and
University of Warwick

Marie-Louise Vierø ‡
Queen's University

May 29, 2015

Abstract

This paper extends our earlier work on reverse Bayesianism by relaxing the assumption that decision makers abide by expected utility theory, assuming instead weaker axioms that merely imply that they are probabilistically sophisticated. We show that our main results, namely, (modified) representation theorems and corresponding rules for updating beliefs over expanding state spaces and null events that constitute "reverse Bayesianism," remain valid.

Keywords: Awareness, unawareness, reverse Bayesianism, probabilistic sophistication

JEL classification: D8, D81, D83

*We thank Graeme Doole, Asen Kochov, and Jacob Sagi for comments.

†Johns Hopkins University, Department of Economics and University of Warwick, Warwick Business School, E-mail: karni@jhu.edu.

‡Queen's University, Department of Economics, Dunning Hall Room 312, 94 University Avenue, Kingston, Ontario K7L 3N6, Canada, viero@econ.queensu.ca.

1 Introduction

The theory of "reverse Bayesianism" is intended to depict the response of Bayesian decision makers to expansions of their universe in the wake of discoveries of new consequences and/or acts, and improved understanding of the links between acts and consequences. In particular, we are interested in those aspects of the structure of the preferences that persist as the universe expands, since those aspects allow one to infer from existing preferences something about the preferences in the expanded environment. This feature of our model is worth emphasizing against the backdrop of consumer theory under certainty. This issue was discussed in Lancaster (1966) who made the following statement:[1]

> "Perhaps the most important aspects of consumer behavior relevant to an economy as complex as that of the United States are those of consumer reactions to new commodities and to quality variations. Traditional theory has nothing to say on these. In the case of new commodities, the theory is particularly helpless. We have to expand from a commodity space of dimension n to one of dimension *n+1,* replacing the old utility function by a completely new one, and *even a complete map of the consumer's preferences among the n goods provides absolutely no information about the new preference map. A theory which can make no use of such information is a remarkably empty one.*" (Lancaster (1966) p. 133, the italics are ours).

In our earlier work on reverse Bayesianism (see Karni and Vierø (2013)) the aspects of the preference structure that remain intact when the decision maker's state space expands are his risk attitudes and beliefs regarding the relative likelihoods of the events in the original state space. Thus, in the context of decision making under uncertainty, this theory does provide some information about the updated preferences.

This paper extends our earlier work on reverse Bayesianism, replacing the subjective expected utility model with probabilistic sophistication. Probabilistically sophisticated choice is characterized by a unique subjective probability measure on a state space by which acts (that is, mappings from the set of states to the set of lotteries over consequences) are transformed to lotteries over consequences, a utility function on the set of these lotteries, and choice behavior that maximizes the utility over the lotteries corresponding to feasible

[1] We thank Graeme Doole for calling our attention to this paragraph in Lancaster's article.

sets of acts. In two seminal papers, Machina and Schmeidler (1992, 1995) axiomatize probabilistically sophisticated choice in the analytical frameworks of Savage (1954) and Anscombe and Aumann (1963). The main contribution of these works is to break the link between the subjective expected utility model and the existence of choice-based subjective probabilities. This is an important extension of Bayesian theory.

Karni and Vierø (2013) introduced a model describing the evolution of the beliefs of subjective expected utility maximizing decision makers as they discover new acts, consequences, and information pertaining to links between acts and consequences. We introduced the notions of conceivable and feasible state spaces and showed how the former state space is constructed from what the decision maker perceives to be the feasible acts and consequences while the latter is derived from his revealed preferences. In addition, we showed how these state spaces expand in response to the discoveries of new consequences and/or acts. In this paper, we extend our earlier work showing that the reverse Bayesianism model is not predicated on subjective expected utility maximizing behavior. This inquiry is motivated, in part, by the large body of experimental work documenting systematic departures from the expected utility model. Accordingly, we relax the strictures of expected utility theory, assuming instead that decision makers are probabilistically sophisticated, and show that the results of Karni and Vierø (2013) hold in this, more general, framework. In particular, the model pursued here admits choice behavior that displays reversals of conditional preferences as the universe expands, a phenomenon that is inconsistent with reverse Bayesianism founded on the subjective expected utility model. As we show below (see section 3.4), the conditional preference reversals are the result of the nonseparability of preferences across states. In particular, our model admits nonseparability with respect to outcomes on null events that are rendered nonnull as a result of the discovery of links between acts and consequences that the decision maker thought impossible. Because our concrete example relies on the analytical framework exposed in the next section and the results in section 3.3, we relegate it to a later section.

Invoking the analytical framework of Anscombe and Aumann (1963), we demonstrate that our main results in Karni and Vierø (2013), namely, (modified) representations of preferences and corresponding rules for updating beliefs over expanding state spaces that constitute "reverse Bayesianism," hold when preferences are probabilistically sophisticated. We accomplish this by replacing, at each level of awareness, the axioms of expected utility theory with the axioms of Machina and Schmeidler (1995), and introduce new axioms to connect the preferences across distinct levels of awareness. Unlike the axioms of Machina and Schmeidler which are well known, the "connecting" axioms are new and merit brief description. The first axiom, invariant risk preferences, introduced in Karni and Vierø (2013), asserts that growing awareness does not alter the decision maker's risk attitudes. The formal expression of the other "connecting" axioms varies with the context. In essence, however, they all assert that the ranking of risky vs. uncertain prospects, conditional on a feasible state space, remains intact when the state space expands. This is in contrast to the corresponding axioms in Karni and Vierø (2013) which require that the entire conditional preference relation remains unchanged.

The results of this work depict the evolution of a decision maker's beliefs in the wake of discoveries of new consequences and acts. In a nutshell, the results assert that, as his state space expands, a probabilistically sophisticated decision maker updates his beliefs in a way that preserves the likelihood ratios of events in the original state space. In addition, we show how a probabilistically sophisticated decision maker updates his beliefs when he arrives at new understanding of the links between acts and consequences. Specifically, when he discovers that links that he believed possible are in fact impossible, he updates his beliefs according to Bayes' rule. When he discovers that links between feasible acts and consequences that he believed impossible are, in fact, possible, he updates the zero probability events using a formula that is analogous to Bayes' rule and is best described as a reverse Bayesian updating rule.

Since the analytical framework and the related literature were discussed in Karni and Vierø (2013), in what follows, we review briefly those aspects of the model necessary to make the exposition self-contained, underscoring instead the adjustment necessary for the transition from expected utility to probabilistically sophisticated choice. This is done in the next section. In section 3 we expose the representation theorems and analyze the evolution of beliefs in the wake of discovery of new consequences, acts and links between them. Concluding remarks appear in section 4. The proofs are collected in section 5.

2 The Analytical Framework

Building upon Schmeidler and Wakker (1987) and Karni and Schmeidler (1991), Karni and Vierø (2013) introduced a unifying framework within which growing awareness due to the discovery of new acts and consequences as well as revising beliefs in light of new information regarding their links may be described and analyzed. We briefly recall this framework below, but refer the reader to Karni and Vierø (2013) for details.

2.1 Conceivable states and acts

Let F be a finite, nonempty set of *feasible acts*, and let C be a finite, nonempty set of *feasible consequences.* Together these sets determine a *conceivable state space*, C^F, whose elements depict the resolutions of uncertainty. In other words, a *state* is a function from the set of feasible acts to the set of feasible consequences. As an illustration, let there be two feasible acts, $F = \{f_1, f_2\}$, and two consequences, $C = \{c_1, c_2\}$. The resulting conceivable state space is C^F, consisting of four states as depicted in the following matrix:

$$\begin{array}{ccccc} F\backslash C^F & s_1 & s_2 & s_3 & s_4 \\ f_1 & c_1 & c_2 & c_1 & c_2 \\ f_2 & c_1 & c_1 & c_2 & c_2 \end{array} \tag{1}$$

Once the set of conceivable states is fixed, we can define the set of *conceivable acts*, which consists of all the mappings from the conceivable state space to lotteries on the set of consequences.[2] Formally, the set of conceivable acts is given by:

$$\hat{F} := \{f : C^F \to \Delta(C)\}, \tag{2}$$

where $\Delta(C)$ is the set of all lotteries with consequences in C as prizes.[3] Conceivable acts are imaginable given the decision maker's awareness of feasible acts and consequences and the corresponding conceivable state space.

A decision maker's conceivable state space expands due to discovery of new feasible acts and/or consequences. Consider the two-act two-consequences example depicted above and imagine that a third feasible consequence was discovered, so that the new set of feasible consequences is $C' = \{c_1, c_2, c_3\}$. The feasible acts need to be redefined, since choosing the act $f_i, i = 1, 2$ conceivably may result in any of the three consequences. We denote the redefined set of feasible acts by F^*. The corresponding conceivable state space is given by

[2] Fishburn (1970) Ch. 12 discusses a construction of a state space along similar lines. He does not, however, discuss an extension of the set of acts to include conceivable acts.

[3] Formally, $p \in \Delta(C)$ is a function $p : C \to [0, 1]$ satisfying $\Sigma_{c \in C} p(c) = 1$. Notice that with this definition of $\Delta(C)$ we have that, for any $C \subset C'$, any $p \in \Delta(C)$ is also an element of $\Delta(C')$ with $p(c) = 0$ for all $c \in C' - C$. Likewise, $q \in \Delta(C')$ is an element of $\Delta(C)$ if $q(c) = 0$ for all $c \in C' - C$.

$(C')^{F^*}$ and consists of nine states. The event $(C')^{F^*} - C^F$ represents the expansion of the decision maker's conceivable state space due to the discovery of the new consequence.

Discovery of new feasible acts also alters the conceivable state space, albeit in a different way. Consider again the two-act two-consequences example above and suppose that a new feasible act, f_3, is discovered. The new set of feasible acts is $F' = \{f_1, f_2, f_3\}$ and the corresponding conceivable state space, $C^{F'}$, consists of eight states and is a finer partition of the original state space C^F.

As the decision maker's conceivable state space expands, so does the set of conceivable acts. In the wake of the discovery of a new consequence, the new set of conceivable acts is $\hat{F}^* := \{f : (C')^{F^*} \to \Delta(C')\}$, and the decision maker's posterior preference relation over these acts is denoted by $\succcurlyeq_{\hat{F}^*}$. In the aftermath of the discovery of a new feasible act, the new set of conceivable acts is $\hat{F}' := \{f : C^{F'} \to \Delta(C)\}$, and the decision maker's posterior preference relation over these acts is denoted by $\succcurlyeq_{\hat{F}'}$.

At this point the reader may find it disturbing that we expand the conceivable state space in the wake of the discovery of new feasible acts but we do not expand the state space when we introduce new conceivable acts. The reason is that conceivable acts are bets on the conceivable states (or events). As we have seen, the discovery of new feasible acts expands the conceivable state space by assigning to every existing state the set of all the consequences, thereby "splitting" it to generate a refined state space. By contrast a new conceivable act assigns to every existing state a unique outcome. Hence, for conceivable acts the subjective uncertainty regarding the payoffs of all acts, feasible or otherwise, is completely resolved once the original state is known. Consequently, the introduction of conceivable acts does not change the conceivable state space.

Decision makers are characterized by preference relations over the conceivable acts. Because the set of conceivable acts is a variable in our model, we denote the preference relation on $\hat{F}$ by $\succcurlyeq_{\hat{F}}$, and denote by $\succ_{\hat{F}}$ and $\sim_{\hat{F}}$ the asymmetric and symmetric parts of $\succcurlyeq_{\hat{F}}$, respectively. These derived relations are given the usual interpretation of strict preference and indifference, respectively. With the usual abuse of notation, we denote by p the constant act that assigns p to each $s \in C^F$ and by c the degenerate lottery δ^c that assigns the unit probability mass to the consequence c.

2.2 Feasible states

Decision makers entertain beliefs about the possible links between feasible acts and their potential consequences. These beliefs manifest themselves in, and may be inferred from, the decision makers' choice behavior.

Consider a decision maker whose choices are characterized by a preference relation $\succcurlyeq_{\hat{F}}$ on $\hat{F}$. For any $f \in \hat{F}$, $p \in \Delta(C)$, and $E \subset C^F$, let $p_E f$ be the act in $\hat{F}$ obtained from f by replacing its sth coordinate with p for all $s \in E$. Following Savage (1954), a state $s \in C^F$ is said to be *null* if $p_{\{s\}} f \sim_{\hat{F}} q_{\{s\}} f$, for all $p, q \in \Delta(C)$, for all $f \in \hat{F}$. A state is said to be *nonnull* if it is not null. Denote by E^N the set of null states and let $S(F,C) = C^F - E^N$ be the set of all nonnull states, which we refer to as the *feasible state space*. Note that a conceivable state is null if it includes an assignment of a feasible consequence to a feasible act that the decision maker believes to be impossible.[4]

New information may change the decision maker's beliefs concerning the links between feasible acts and consequences and, consequently, his perception of the feasible state space. Unlike the discovery of new feasible consequences and/or new feasible acts, changes of the decision maker's beliefs concerning the links between them expand or contract only the set of feasible states without affecting the conceivable state space.

Expansion of the feasible state space entails updating zero probability events, while contraction of it entails nullifying positive probability events that are no longer considered possible. When new links become possible, the decision maker includes the consequences $f(s)$, for all $f \in F$ and some $s \in C^F - S(F,C)$, in the ranges he considers possible of the feasible acts. Vice versa when old links are eliminated. We denote the newly defined feasible state space by $S'(F,C)$, the corresponding set of conceivable acts by $\hat{F}_{S'}$, and the decision maker's posterior preference relation by $\succcurlyeq_{\hat{F}_{S'}}$.

2.3 Basic preference structure

Let F and C be finite, nonempty sets of feasible acts and consequences, respectively. The set of conceivable states is given by C^F and the corresponding set of conceivable acts by $\hat{F} := \{f : C^F \to \Delta(C)\}$ as described above. For all $f, g \in \hat{F}$ and $\alpha \in [0,1]$ define the convex combination $\alpha f + (1-\alpha) g \in \hat{F}$ by: $(\alpha f + (1-\alpha) g)(s) = \alpha f(s) + (1-\alpha) g(s)$, for all $s \in C^F$. Then, $\hat{F}$ is a convex subset of a linear space.

[4] Fishburn's (1970) notion of excluded states is analogous to our non-feasible states.

We assume throughout that each set of consequences has a most preferred and a least preferred element. Formally, there exist $c^*(C), c_*(C) \in C$ such that the constant act that assigns $c^*(C)$ to every state is strictly preferred over any other constant act in $\hat{F}$ and the constant act that assigns $c_*(C)$ to every state is strictly less preferred than any other constant act in $\hat{F}$.

For all $p, q \in \Delta(C)$, p dominates q according to first-order stochastic dominance if $\Sigma_{\{i|c_i \preccurlyeq c\}} p(c_i) \leq \Sigma_{\{i|c_i \preccurlyeq c\}} q(c_i)$ for all $c \in C$, and p strictly dominates q according to first-order stochastic dominance if p dominates q according to first-order stochastic dominance and, in addition, $\Sigma_{\{i|c_i \preccurlyeq c\}} p(c_i) < \Sigma_{\{i|c_i \preccurlyeq c\}} q(c_i)$ for some $c \in C$. We denote these domination relations by $p \geq^1 q$ and $p >^1 q$, respectively.

As described above, when the state space expands in the wake of discoveries of new feasible consequences, the set of conceivable acts must be expanded and the preference relations must be redefined on the extended domain.

Following Machina and Schmeidler (1995), we assume that, for each $\hat{F}$, $\succcurlyeq_{\hat{F}}$ adheres to the following axioms, which ensure probabilistic sophistication.

(A.1) (**Weak order**) For every $\hat{F}$, the preference relation $\succcurlyeq_{\hat{F}}$ is transitive and complete.

(A.2) (**Mixture continuity**) For each $\hat{F}$ and all $f, g, h \in \hat{F}$, if $f \succ_{\hat{F}} g$ and $g \succ_{\hat{F}} h$ then there exist $\alpha \in (0,1)$ such that $\alpha f + (1-\alpha) h \sim_{\hat{F}} g$.

(A.3) (**Monotonicity**) For every $\hat{F}$ and $p, q \in \Delta(C)$, if $p \geq^1 q$ then $p_{E_i} h \succcurlyeq_{\hat{F}} q_{E_i} h$, for all partitions $\{E_1, ..., E_n\}$ of C^F and all $h \in \hat{F}$, with $p_{E_i} h \succ_{\hat{F}} q_{E_i} h$ if $p >^1 q$ and E_i is nonnull.

(A.4) (**Replacement**) For every $\hat{F}$ and any partition $\{E_1, ..., E_n\}$ of C^F, if

$$\delta^{c^*(C)}\ _{E_i} \left(\delta^{c_*(C)}\ _{E_j} \delta^{c_*(C)} \right) \sim_{\hat{F}} \left(\alpha \delta^{c^*(C)} + (1-\alpha) \delta^{c_*(C)} \right)_{E_i \cup E_j} \delta^{c_*(C)}$$

for some $\alpha \in [0,1]$ and pair of events E_i, E_j, then

$$p_{E_i} \left(q_{E_j} h \right) \sim_{\hat{F}} \left(\alpha p + (1-\alpha) q \right)_{E_i \cup E_j} h$$

for all $p, q \in \Delta(C)$ and $h \in \hat{F}$.

(A.5) (**Nontriviality**) For every $\hat{F}$, $\succ_{\hat{F}} \neq \varnothing$.

To link the preference relations across expanding sets of conceivable acts, we invoke the invariant risk preferences axiom introduced in Karni and Vierø (2013), asserting the commonality of risk attitudes across levels of awareness.

(A.6) (**Invariant risk preferences**) For every given $\hat{F}, \hat{F}'$, if C and C' are the sets of consequences associated with $\hat{F}$ and $\hat{F}'$, respectively, then $p \succcurlyeq_{\hat{F}} q$ if and only if $p \succcurlyeq_{\hat{F}'} q$ for all $p, q \in \Delta\left(C \cap C'\right)$.

When new consequences are discovered, i.e. $C \subset C'$, then $C \cap C' = C$. When new feasible acts are discovered, the invariant risk preferences axiom may be stated as follows: For all F, F' and $p, q \in \Delta(C)$, $p \succcurlyeq_{\hat{F}} q$ if and only if $p \succcurlyeq_{\hat{F}'} q$. When new links are discovered (or old links eliminated) between the original sets of acts, F, and consequences, C, the invariant risk preferences axiom asserts that, for all $p, q \in \Delta(C)$, $p \succcurlyeq_{\hat{F}} q$ if and only if $p \succcurlyeq_{\hat{F}_{S'}} q$.

3 The Main Results

As in Karni and Vierø (2013), we divide the analysis of the effects of growing awareness on choice behavior and the evolution of decision makers' beliefs into three parts. First, we explore the implications of the discovery of new consequences. Second, we explore the implications of the discovery of new feasible acts. Third, we explore the implications of new information regarding acts-consequences links. The discovery of new acts or consequences increases the number of conceivable states and, in general, also that of feasible states. However, unlike the discovery of new consequences, the discovery of new feasible acts increases the number of conceivable states by refining the original state space. By contrast, the discovery of new acts-consequences links changes the set of feasible states without affecting the conceivable state space.

To explore the implications of these sources of growing awareness, we introduce additional axioms, each of which modifies a corresponding axiom in Karni and Vierø (2013). These modifications are needed in order to accommodate the possibility that preferences are not necessarily separable. As it is the case for the consistency axioms in Karni and Vierø (2013), the consistency axioms below ensure robustness of the decision maker's preferences with respect to future discoveries.

3.1 Discovery of new consequences and its representation

The following axiom requires that when a decision maker discovers new consequences his ranking of subjective versus objective uncertainty, conditional on the original set of feasible states, remains intact. To formalize this idea, let $C' \supset C$, F^*, and $S(F^*, C')$ denote, respectively, the new set of consequences, the set of feasible acts redefined to accommodate the new consequences, and the resulting new feasible state space.[5]

(A.7) (**Replacement consistency I**) For every given F, for all C, C' with $C \subset C'$ and $S(F,C) \subseteq S(F^*, C')$, for all $s \in S(F,C)$, $\eta \in [0,1]$, $f, g \in \hat{F}$, and $f', g' \in \hat{F}^*$, if $f = \delta^{c^*(C)}{}_{\{s\}}\delta^{c_*(C)}$, $g = \eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}$ on C^F, $g' = \eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}$ on $(C')^{F^*}$, $f' = f$ on $S\,(F,C)$ and $f' = g'$ on $S\,(F^*, C') - S(F,C)$, then it holds that $f \succcurlyeq_{\hat{F}} g$ if and only if $f' \succcurlyeq_{\hat{F}^*} g'$.

Axiom (A.7) concerns bets that involve only the best and worst consequences in C. The act g is defined on the original state space and is an objective bet that pays off with probability η, while the act f, also defined on the original state space, is a subjective bet that pays off in state s. The act g' extends the objective bet to the expanded state space such that it continues to only involve objective uncertainty. Finally, f' is an extension of f that agrees with g' on the new event. The axiom therefore asserts that, conditional on the prior subjective state space, the decision maker's ranking of the subjective bet that pays off in state s and the objective bet that pays off with probability η is preserved when he discovers new feasible consequences which expand the conceivable state space. It thus ensures consistency between the ranking of subjective versus objective uncertainty given awareness of conceivable acts that correspond to nested sets of feasible consequences.

The corresponding axiom, awareness consistency, in Karni and Vierø (2013) requires that when the decision maker discovers new consequences, his entire preference relation conditional on the prior feasible state space remains unchanged. Thus, awareness consistency is reminiscent of Savage's (1954) sure thing principle in that it requires that preference between acts is independent of the aspects on which the acts agree. Put differently, the awareness consistency axiom requires consistency of the conditional ranking of *all* acts.

[5]Below, $f' = f$ on an event E means that $f'(s) = f(s)$ for all $s \in E$ (i.e., it is defined pointwise for the states in E).

To the contrary, the present Axiom (A.7) only requires consistency for a subset of initial acts and extensions. In particular, it only requires consistency for initial acts f that involve subjective uncertainty regarding a single state and g that involve purely objective uncertainty, and extensions f' and g' of these that agree on the new event such that g' continues to only involve objective uncertainty. Put differently, it only requires consistency of the conditional ranking when one of the acts is a perfect insurance against subjective uncertainty and continues to be that after the expansion of the state space.

Our first result describes the evolution of a decision maker's beliefs in the wake of discoveries of new consequences. Like Theorem 1 in Karni and Vierø (2013) this theorem asserts that, as he becomes aware of new consequences, the decision maker updates his beliefs in a way that preserves the likelihood ratios of events in the original state space. Unlike in Karni and Vierø (2013) the decision maker is not necessarily an expected utility maximizer, he is merely probabilistically sophisticated. Hence, reverse Bayesianism is independent of the expected utility hypothesis.

Theorem 1 For each set, $\hat{F}$, of conceivable acts let $\succcurlyeq_{\hat{F}}$ be a binary relation on $\hat{F}$ then, for all $\hat{F}, \hat{F}^*$, the following two conditions are equivalent:

(i) $\succcurlyeq_{\hat{F}}$ and $\succcurlyeq_{\hat{F}^*}$ each satisfy (A.1) - (A.5) and jointly, $\succcurlyeq_{\hat{F}}$ and $\succcurlyeq_{\hat{F}^*}$ satisfy (A.6) and (A.7).

(ii) There exist real-valued, mixture continuous, strictly monotonic[6] functions, V on $\Delta(C)$ and V^* on $\Delta(C')$, and probability measures, $\pi_{\hat{F}}$ on C^F and $\pi_{\hat{F}^*}$ on $(C')^{F^*}$, such that for all $f, g \in \hat{F}$,

$$f \succcurlyeq_{\hat{F}} g \Leftrightarrow V\left(\Sigma_{s\in S(F,C)}\pi_{\hat{F}}(s) f(s)\right) \geq V\left(\Sigma_{s\in S(F,C)}\pi_{\hat{F}}(s) g(s)\right). \quad (3)$$

and, for all $f', g' \in \hat{F}^*$,

$$f' \succcurlyeq_{\hat{F}^*} g' \Leftrightarrow V^*\left(\Sigma_{s\in S(F^*,C')}\pi_{\hat{F}^*}(s) f'(s)\right) \geq V^*\left(\Sigma_{s\in S(F^*,C')}\pi_{\hat{F}^*}(s) g'(s)\right). \quad (4)$$

The functions V and V^* are unique up to positive transformations and V^* is an extension of V, the probability measures $\pi_{\hat{F}}$ and $\pi_{\hat{F}^*}$ are unique, $\pi_{\hat{F}}(S(F,C)) = \pi_{\hat{F}^*}(S(F^*,C')) = 1$,

[6] A function V is strictly monotonic if $V(p) \geq V(q)$ whenever p dominates q according to first-order stochastic dominance, with strict inequality in the case of strict dominance, and V is mixture continuous if $V(\alpha p + (1-\alpha) q)$ is continuous in α for all p and q.

and, for all $s \in S(F,C)$,

$$\pi_{\hat{F}}(s) = \frac{\pi_{\hat{F}^*}(s)}{\pi_{\hat{F}^*}(S(F,C))}. \tag{5}$$

3.2 Discovery of new feasible acts and its representation

The discovery of new feasible acts expands the conceivable state space and increases the number of coordinates defining a state. To state the next axiom, which is analogous to Axiom (A.7), we introduce the following additional notations: If $F \subset F'$ then for each $s \in C^F$ there corresponds an event $E(s) \subset C^{F'}$ defined by $E(s) = \{s' \in C^{F'} \mid \boldsymbol{P}_{C^F}(s') = s\}$, where $\boldsymbol{P}_{C^F}(\cdot)$ is the projection of $C^{F'}$ on C^F.[7] For $s \in C^F$, we refer to the set $E(s)$ as the inverse image under $\boldsymbol{P}_{C^F}$ of s on $C^{F'}$.

(A.8) (**Replacement consistency II**) For every given C, all pairs of feasible acts F and F' such that $F \subset F'$, all $s \in S(F,C)$, $\eta \in [0,1]$, $f,g \in \hat{F}$, and $f',g' \in \hat{F}'$, if $f = \delta^{c^*(C)}{}_{\{s\}}\delta^{c_*(C)}$, $g = g' = \eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}$, and $f' = \delta^{c^*(C)}{}_{E(s)}\delta^{c_*(C)}$, then it holds that $f \succcurlyeq_{\hat{F}} g$ if and only if $f' \succcurlyeq_{\hat{F}'} g'$.

Like Axiom (A.7), Axiom (A.8) concerns bets that involve only the best and worst consequences in C. It asserts that the ranking of subjective bets that pay off in state s and the objective bets that pay off with probability η, conditional on a given set of conceivable acts, is the same as the ranking of subjective bets that pay off in the event $E(s)$ and the objective bets that pay off with probability η conditional on the set of conceivable acts spanned by the discovery of new feasible acts. In other words, the axiom asserts that the decision maker's ranking of subjective versus objective uncertainty is independent of the detail with which the subjective uncertainty is described.

The representation theorem below describes how a decision maker's beliefs evolve as he becomes aware of new feasible acts. Specifically, the decision maker updates his beliefs so that the probability of each state in the original state space is equal to that of its inverse image under $\boldsymbol{P}_{C^F}$ on $C^{F'}$. In other words, since the event $E(s)$ in $C^{F'}$ is a refinement of the state s in C^F, its probability is equal to that of s.

Theorem 2 For each set $\hat{F}$ of conceivable acts let $\succcurlyeq_{\hat{F}}$ be a binary relation on $\hat{F}$, then, for all $\hat{F}, \hat{F}'$, the following two conditions are equivalent:

[7] Suppose that $\mid F \mid= r$ and $\mid F' \mid= k > r$. Let $s = (c_1, ..., c_k) \in C^{F'}$, then $\boldsymbol{P}_{C^F}(s) = (c_1, ..., c_r) \in C^F$.

(i) $\succcurlyeq_{\hat{F}}$ and $\succcurlyeq_{\hat{F}'}$ each satisfy (A.1) - (A.5) and jointly, $\succcurlyeq_{\hat{F}}$ and $\succcurlyeq_{\hat{F}'}$ satisfy (A.6) and (A.8).

(ii) There exist a real-valued, mixture continuous, strictly monotonic function V on $\Delta(C)$ and probability measures, $\pi_{\hat{F}}$ on C^F and $\pi_{\hat{F}'}$ on $C^{F'}$, such that for all $f, g \in \hat{F}$,

$$f \succcurlyeq_{\hat{F}} g \Leftrightarrow V\left(\Sigma_{s \in S(F,C)} \pi_{\hat{F}}(s) f(s)\right) \geq V\left(\Sigma_{s \in S(F,C)} \pi_{\hat{F}}(s) g(s)\right). \tag{6}$$

and, for all $f', g' \in \hat{F}'$,

$$f' \succcurlyeq_{\hat{F}'} g' \Leftrightarrow V\left(\Sigma_{s \in S(F',C)} \pi_{\hat{F}'}(s) f'(s)\right) \geq V\left(\Sigma_{s \in S(F',C)} \pi_{\hat{F}'}(s) g'(s)\right). \tag{7}$$

The function V is unique up to positive transformations, the probability measures $\pi_{\hat{F}}$ and $\pi_{\hat{F}'}$ are unique, $\pi_{\hat{F}}(S(F,C)) = \pi_{\hat{F}'}(S(F',C)) = 1$, and, for all $s \in S(F,C)$,

$$\pi_{\hat{F}}(s) = \pi_{\hat{F}'}(E(s)) \tag{8}$$

where $E(s)$ is the inverse image under $\boldsymbol{P}_{C^F}$ of s on $C^{F'}$.

3.3 Discovery of new feasible states and the nullification of existing feasible states and their representations

When links between feasible acts and consequences that were believed to exist are discovered to be non-existent, the feasible state space contracts. Similarly, when such links that were believed not to exist are discovered to exist, the feasible state space expands. In the first instance, an event that was believed to be nonnull and was assigned positive probability becomes a null event and must be assigned zero probability. In the second instance, an event that was believed to be null and was assigned zero probability becomes a nonnull event and must be assigned positive probability.

The next axiom depicts the evolution of the preference relation in these circumstances. Clearly, the first instance described above corresponds to the usual Bayesian updating. The second instance, in which the posterior of a zero-probability event is positive, does not admit Bayesian updating. It is, however, consistent with our model of reverse Bayesianism. In fact, in our model the two instances are treated symmetrically, which is reassuring given that they depict symmetrically opposing discoveries.

(A.9) (**Replacement consistency III**) For all pairs of conceivable acts $\hat{F}$ and $\hat{F}_{S'}$, all $f, g \in \hat{F}$ and $f', g', \in \hat{F}_{S'}$, if $g = g' = \eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}$ for some $\eta \in [0,1]$ and

$f = f' = \delta^{c^*(C)}{}_{\{s\}} \left(\delta^{c_*(C)}{}_{S(F,C) \cap S'(F,C)} g\right)$, for some $s \in S(F,C) \cap S'(F,C)$, then it holds that $f \succcurlyeq_{\hat{F}} g$ if and only if $f' \succcurlyeq_{\hat{F}_{S'}} g'$.

The next representation theorem describes how a decision maker's beliefs are updated as he discovers that links between feasible acts and consequences that he believed impossible are in fact possible and when he discovers that links that he believed possible are in fact impossible.

Theorem 3 For each set of conceivable acts $\hat{F}$, let $\succcurlyeq_{\hat{F}}$ be a binary relation on $\hat{F}$ then, for all $\hat{F}$ and $\hat{F}_{S'}$, the following two conditions are equivalent:

(*i*) Each $\succcurlyeq_{\hat{F}}$ and $\succcurlyeq_{\hat{F}_{S'}}$ satisfy (A.1) - (A.5) and jointly $\succcurlyeq_{\hat{F}}$ and $\succcurlyeq_{\hat{F}_{S'}}$ satisfy (A.6) and (A.9).

(*ii*) There exist a real-valued, mixture continuous, strictly monotonic function V on $\Delta(C)$ and, for all $\hat{F}$ and $\hat{F}_{S'}$, there are probability measures $\pi_{\hat{F}}$ and $\pi_{\hat{F}_{S'}}$ on C^F such that, for all $f, g \in \hat{F}$,

$$f \succcurlyeq_{\hat{F}} g \Leftrightarrow V\left(\Sigma_{s \in S(F,C)} \pi_{\hat{F}}(s) f(s)\right) \geq V\left(\Sigma_{s \in S(F,C)} \pi_{\hat{F}}(s) g(s)\right). \tag{9}$$

and, for all $f', g' \in \hat{F}_{S'}$,

$$f' \succcurlyeq_{\hat{F}_{S'}} g' \Leftrightarrow V\left(\Sigma_{s \in S'(F,C)} \pi_{\hat{F}_{S'}}(s) f'(s)\right) \geq V\left(\Sigma_{s \in S'(F,C)} \pi_{\hat{F}_{S'}}(s) g'(s)\right). \tag{10}$$

The function V is unique up to positive transformations, the probability measures $\pi_{\hat{F}}$ and $\pi_{\hat{F}_{S'}}$ are unique, $\pi_{\hat{F}}(S(F,C)) = \pi_{\hat{F}_{S'}}(S'(F,C)) = 1$, and

$$\frac{\pi_{\hat{F}}(s)}{\pi_{\hat{F}}(s')} = \frac{\pi_{\hat{F}_{S'}}(s)}{\pi_{\hat{F}_{S'}}(s')} \tag{11}$$

for all $s, s' \in S(F,C) \cap S'(F,C)$.

If $S'(F,C) \subset S(F,C)$ then Theorem 3 describes Bayesian updating (that is, (11) may be written as $\pi_{\hat{F}_{S'}}(s) = \pi_{\hat{F}}(s) / \pi_{\hat{F}}(S'(F,C))$, for all $s \in S'(F,C)$). If $S'(F,C) \supset S(F,C)$ then $\pi_{\hat{F}}(s) = \pi_{\hat{F}_{S'}}(s) / \pi_{\hat{F}_{S'}}(S(F,C))$ for all $s \in S(F,C)$.

3.4 Preference reversals

Consider next an example of the reversal of conditional preferences we alluded to in the introduction. Let there be three outcomes, c_1, c_2, and c_3, with c_1 and c_3 being the most and least preferred outcomes, respectively. Let there be two acts, f which perfectly insures the decision maker by always returning the outcome c_2, and g which may return either c_1 or c_3. Consequently, there are two feasible states, $s_1 = (c_2, c_1)$ and $s_2 = (c_2, c_3)$. Suppose that the decision maker strictly prefers f over g. The decision maker now makes a discovery that leads him to believe that the act f may also result in the consequence c_3 if act g results in c_3. This implies an expansion of the feasible state space to include the state $s_3 = (c_3, c_3)$. The discovery also redefines the set of conceivable acts, which now includes the act f' that returns c_2 in states s_1 and s_2 and c_3 in state s_3, and the act g' that returns c_1 in state s_1 and c_3 in states s_2 and s_3. The act f' thus agrees with f conditional on the initial feasible state space, as does the act g' with g. Since f' and g' agree on s_3, the axioms in Karni and Vierø (2013) force f' to be preferred to g'. (This is because the representation in Karni and Vierø (2013) is separable across states, so the terms that concern the state s_3 in which the acts agree cancel out). However, the decision maker may have preferred f over g initially either because f provided perfect insurance or because of disappointment aversion. In other words, before the discovery the decision maker can avoid the ex post sense of disappointment (or elation) by choosing f, while by choosing g he exposes himself to disappointment (if s_2 obtains) or elation (if state s_1 obtains). After the state space expands, since f' might also result in disappointment (if s_3 obtains) the decision maker can no longer avoid the disappointment and as a result may prefer the potential higher gain from g'. The latter type of behavior is permitted by the axioms in the present paper.[8]

As the preceding example shows, the behavioral implications of the model in the present paper are different from those of the model in Karni and Vierø (2013), since the present paper's model is compatible with conditional preference reversals, while our previous model is not. However, the two models have in common that both growing awareness of consequences and the discovery of new links between acts and consequences can lead to reversal of *unconditional* preferences between two acts. To grasp this, consider the conceivable state space with four states depicted in the matrix (1), and suppose that $S(F, C) = \{s_2, s_3\}$,

[8] Gul (1991) axiomatized a model of disappointment aversion under risk. The argument here is analogous to that of Gul except that here it is cast in terms of uncertainty.

that $\pi(s_2) < \pi(s_3)$, that $c_1 \succ c_2$, and that we are in the special case of Theorem 1 in which the representation has an expected utility form. Then $f_1 \succ_{\hat{F}} f_2$. Suppose now that a new consequence, $c_3 \prec c_2$, is discovered, resulting in a conceivable state space that includes an additional five states. Our update rules imply that beliefs under the expanded level of awareness have $\frac{\pi_{\hat{F}^*}(s_2)}{\pi_{\hat{F}^*}(s_3)} = \frac{\pi_{\hat{F}}(s_2)}{\pi_{\hat{F}}(s_3)}$ and thus that $\pi_{\hat{F}^*}(s_2) < \pi_{\hat{F}^*}(s_3)$. However, the following is consistent with our update rules: $S(F^*, C') = \{s_2, s_3, s_5\}$, where $s_5 = (c_3, c_1)$, and $\pi_{\hat{F}^*}(s_5) > (\pi_{\hat{F}^*}(s_3) - \pi_{\hat{F}^*}(s_2))\frac{u^*(c_1)-u^*(c_2)}{u^*(c_1)-u^*(c_3)}$. Then, by the representation (5) in Theorem 1, $f_2 \succ_{\hat{F}^*} f_1$. That is, the change in awareness has led to a reversal of unconditional preferences over the two acts. Enough probability mass is shifted to the newly emerged feasible state that the worse outcome to f_1 in that state outweighs the ranking conditional on the initial state space. This causes the preference reversal.

Unlike growing awareness of consequences or the discovery of new links between acts and consequences, growing awareness of feasible acts cannot lead to preference reversals among existing feasible acts. The reason is twofold. First, our update rule in Theorem 2 implies that the probability of each state in the original state space equals the probability of the corresponding event in the new state space. Second, all existing feasible acts remain measurable with respect to the original state space under the expanded level of awareness. Therefore, the ranking of feasible acts must be unchanged.

4 Concluding Remarks

Grant and Polak (2006) propose an alternative axiomatization of probabilistically sophisticated choice behavior, which is equivalent to that in Machina and Schmeidler (1995) and "decomposes" the independence assumptions that are built into the replacement axiom of Machina and Schmeidler. The axioms of Grant and Polak, together with our axioms (A.6) through (A.9), would result in the analogues of Theorems 1, 2, and 3 in the axiomatic framework of Grant and Polak (2006).

Kochov (2010) shows that a decision maker's beliefs when he is aware that his perception of the environment is incomplete are represented by a non-singleton set of priors. As the decision maker's perception of the environment becomes more precise, Kochov's axioms imply that the updating of the multiple priors is rectangular as in Epstein and Schneider (2003). Kochov's model corresponds to our analysis of the discovery of new feasible acts.

The issue of updating the prior probability of null (that is, zero-probability) events was recently addressed by Ortoleva (2012). Ortoleva proposes an interesting model, according to which decision makers are characterized by a set of priors and a prior on the set of priors. Before the arrival of information, decision makers maximize their subjective expected utility with respect to the most likely prior in the support of the prior on priors. When new information indicates that an event that was assigned zero probability according to this prior obtains, the prior is discarded, the prior on priors is updated according to Bayes' rule, and the most likely prior in the support of the updated distribution is used to evaluate acts, by the subjective expected utility criterion. For this method to work every null event must be assigned positive probability by some prior in the support of the prior on priors. This begs the question of updating beliefs on events that are believed impossible. Ortoleva's (2012) approach to updating the probabilities of null events is different from ours in that it is non-Bayesian. If reduction of compound lotteries was applied in Ortoleva's model to obtain a unique prior, then all the null events in his model would be assigned positive probabilities.

5 Proofs

5.1 Proof of Theorem 1

$(i) \Rightarrow (ii)$. Since $\succcurlyeq_{\hat{F}}$ and $\succcurlyeq_{\hat{F}^*}$ satisfy (A.1) - (A.5), the Theorem of Machina and Schmeidler (1995) implies (3) and (4) as well as the uniqueness of V and V^* and of $\pi_{\hat{F}}$ and $\pi_{\hat{F}^*}$. By (3) and (4), the restriction of $\succcurlyeq_{\hat{F}}$ and $\succcurlyeq_{\hat{F}^*}$ to the constant acts $p \in \Delta(C)$ imply that $V(p) \geq V(q)$ if and only if $p \succcurlyeq_{\hat{F}} q$ and $p \succcurlyeq_{\hat{F}^*} q$ if and only if $V^*(p) \geq V^*(q)$. By (A.6), $p \succcurlyeq_{\hat{F}} q$ if and only if $p \succcurlyeq_{\hat{F}^*} q$. Thus, by the uniqueness of the representations, V and V^* can be chosen so that $V = V^*$ on $\Delta(C)$.

To prove (5) suppose that, for some $s \in S(F, C)$,

$$\pi_{\hat{F}}(s) \neq \frac{\pi_{\hat{F}^*}(s)}{\pi_{\hat{F}^*}(S(F,C))}.$$

Without loss of generality, let

$$\pi_{\hat{F}}(s) > \frac{\pi_{\hat{F}^*}(s)}{\pi_{\hat{F}^*}(S(F,C))} := \pi_{\hat{F}^*}(s \mid S(F,C)).$$

Then there is $\eta \in \left(\pi_{\hat{F}^*}(s \mid S(F,C)), \pi_{\hat{F}}(s)\right)$. By the representation in (3), $f = \delta^{c^*(C)}{}_{\{s\}}\delta^{c_*(C)} \sim_{\hat{F}} \pi_{\hat{F}}(s)\,\delta^{c^*(C)} + \left(1 - \pi_{\hat{F}}(s)\right)\delta^{c_*(C)}$.

Since $\pi_{\hat{F}}(s) > \eta > \pi_{\hat{F}^*}(s \mid S(F,C))$, by Axiom (A.3), we have the following ranking of lotteries:

$$\pi_{\hat{F}}(s)\,\delta^{c^*(C)} + \left(1 - \pi_{\hat{F}}(s)\right)\delta^{c_*(C)} \succ_{\hat{F}} \eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)} \tag{12}$$

and

$$\eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)} \succ_{\hat{F}^*} \pi_{\hat{F}^*}(s \mid S(F,C))\,\delta^{c^*(C)} + \left(1 - \pi_{\hat{F}^*}(s \mid S(F,C))\right)\delta^{c_*(C)} \tag{13}$$

for all $\eta \in \left(\pi_{\hat{F}^*}(s \mid S(F,C)), \pi_{\hat{F}}(s)\right)$.

Now, by (13) and Axiom (A.3),

$$\begin{aligned} & \left(\eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}\right)_{S(F,C)}\left(\eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}\right) \\ \succ_{\hat{F}^*} \quad & \left(\pi_{\hat{F}^*}(s \mid S(F,C))\delta^{c^*(C)} + \left(1 - \pi_{\hat{F}^*}(s \mid S(F,C))\right)\delta^{c_*(C)}\right)_{S(F,C)}\left(\eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}\right), \end{aligned}$$

which is equivalent to

$$\begin{aligned} & \eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)} \\ \succ_{\hat{F}^*} \quad & \left(\pi_{\hat{F}^*}(s \mid S(F,C))\,\delta^{c^*(C)} + \left(1 - \pi_{\hat{F}^*}(s \mid S(F,C))\right)\delta^{c_*(C)}\right)_{S(F,C)}\left(\eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}\right), \end{aligned}$$

which, by (4), is equivalent to

$$\begin{aligned} & V^*\left(\eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}\right) \\ > \quad & V^*\Bigg(\pi_{\hat{F}^*}(S(F,C))\left(\pi_{\hat{F}^*}(s \mid S(F,C))\delta^{c^*(C)} + \left(1 - \pi_{\hat{F}^*}(s \mid S(F,C))\right)\delta^{c_*(C)}\right) \\ & + \left(1 - \pi_{\hat{F}^*}(S(F,C))\right)\left(\eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}\right)\Bigg) \qquad (14) \\ = \quad & V^*\Bigg(\pi_{\hat{F}^*}(s)\delta^{c^*(C)} + \left(\pi_{\hat{F}^*}(S(F,C)) - \pi_{\hat{F}^*}(s)\right)\delta^{c_*(C)} + \left(1 - \pi_{\hat{F}^*}(S(F,C))\right)\left(\eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}\right)\Bigg). \end{aligned}$$

By (4), inequality (14) is equivalent to

$$\eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)} \succ_{\hat{F}^*} {\delta^{c^*(C)}}_{\{s\}}\left({\delta^{c_*(C)}}_{S(F,C)}\left(\eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}\right)\right). \tag{15}$$

Now, by (12) and (3),

$$V\left(\pi_{\hat{F}}(s)\,\delta^{c^*(C)} + \left(1 - \pi_{\hat{F}}(s)\right)\delta^{c_*(C)}\right) > V\left(\eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}\right),$$

which is equivalent to

$$\delta^{c^*(C)}{}_{\{s\}}\delta^{c_*(C)} \succ_{\hat{F}} \eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}. \tag{16}$$

But the act on the left hand side of (15) is the act g' in Axiom (A.7), while the act on the right hand side of (15) is the act f'. Also, the act on the left hand side of (16) is the act f in Axiom (A.7), while the act on the right hand side of (16) is the act g. Expressions (15) and (16) thus imply that $f \succ_{\hat{F}} g$ and $g' \succ_{\hat{F}^*} f'$, a contradiction of Axiom (A.7).

$(ii) \to (i)$. That $\succcurlyeq_{\hat{F}}$ and $\succcurlyeq_{\hat{F}^*}$ satisfy (A.1) - (A.5) is an implication of the Theorem of Machina and Schmeidler (1995). Invariant risk preferences, (A.6), follows from the equality of V and V^* on $\Delta(C)$.

To show that (A.7) holds, let $f, g \in \hat{F}$ and $f', g' \in \hat{F}^*$ be as in (A.7). By (3),

$$f \succcurlyeq_{\hat{F}} g \Leftrightarrow V\left(\pi_{\hat{F}}(s)\,\delta^{c^*(C)} + \left(1-\pi_{\hat{F}}(s)\right)\delta^{c_*(C)}\right) \geq V\left(\eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}\right).$$

By the equality of V and V^* on $\Delta(C)$ and (5), the last inequality holds if and only if

$$V^*\left(\frac{\pi_{\hat{F}^*}(s)}{\pi_{\hat{F}^*}(S(F,C))}\delta^{c^*(C)} + \left(1-\frac{\pi_{\hat{F}^*}(s)}{\pi_{\hat{F}^*}(S(F,C))}\right)\delta^{c_*(C)}\right) \geq V^*\left(\eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}\right). \tag{17}$$

which, by (4), is equivalent to

$$\frac{\pi_{\hat{F}^*}(s)}{\pi_{\hat{F}^*}(S(F,C))}\delta^{c^*(C)} + \left(1-\frac{\pi_{\hat{F}^*}(s)}{\pi_{\hat{F}^*}(S(F,C))}\right)\delta^{c_*(C)} \succcurlyeq_{\hat{F}^*} \eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}. \tag{18}$$

Now, since the left-hand-side lottery in (18) first-order stochastically dominates the right-hand side lottery, by Axiom (A.3)

$$\begin{aligned}&\left(\pi_{\hat{F}^*}(s \mid S(F,C))\,\delta^{c^*(C)} + \left(1-\pi_{\hat{F}^*}(s \mid S(F,C))\right)\delta^{c_*(C)}\right)_{S(F,C)}\left(\eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}\right)\\ \succcurlyeq\ &_{\hat{F}^*}\left(\eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}\right)_{S(F,C)}\left(\eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}\right),\end{aligned}$$

Hence, (17) holds if and only if $V^*\left(\xi\delta^{c^*(C)} + (1-\xi)\,\delta^{c_*(C)}\right) \geq V^*\left(\eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}\right)$, where

$$\xi := \left(\pi_{\hat{F}^*}(S(F,C))\frac{\pi_{\hat{F}^*}(s)}{\pi_{\hat{F}^*}(S(F,C))} + \left(1-\pi_{\hat{F}^*}(S(F,C))\right)\eta\right) = \pi_{\hat{F}^*}(s) + \left(1-\pi_{\hat{F}^*}(S(F,C))\right)\eta.$$

Since $\xi\delta^{c^*(C)} + (1-\xi)\,\delta^{c_*(C)} \in \Delta(C)$ is the constant act whose payoff is $\Sigma_{s\in S(F^*,C')}\pi_{\hat{F}^*}(s)\,f'(s)$, the representation result in (4) implies that (17) holds if and only if $f' \succcurlyeq_{\hat{F}^*} g'$.♠

5.2 Proof of Theorem 2

$(i) \Rightarrow (ii)$. Since $\succcurlyeq_{\hat{F}}$ and $\succcurlyeq_{\hat{F}'}$ satisfy (A.1) - (A.5), the Theorem of Machina and Schmeidler (1995) implies a representation as in (6) as well as the uniqueness of V and of $\pi_{\hat{F}}$ for each level of awareness. By (A.6), $p \succcurlyeq_{\hat{F}} q$ if and only if $p \succcurlyeq_{\hat{F}'} q$. Thus, by the uniqueness of the representations, V can be chosen to be invariant to the level of awareness.

To prove (8), suppose that, for some $s \in S(F, C)$, $\pi_{\hat{F}}(s) \neq \pi_{\hat{F}'}(E(s))$. Without loss of generality, let $\pi_{\hat{F}}(s) > \pi_{\hat{F}'}(E(s))$. Then there exists $\eta \in \left(\pi_{\hat{F}'}(E(s)), \pi_{\hat{F}}(s)\right)$, and by Axiom (A.3), we have the following ranking of lotteries:

$$\pi_{\hat{F}}(s)\, \delta^{c^*(C)} + \left(1 - \pi_{\hat{F}}(s)\right) \delta^{c_*(C)} \succ_{\hat{F}} \eta \delta^{c^*(C)} + (1-\eta)\, \delta^{c_*(C)} \tag{19}$$

and

$$\eta \delta^{c^*(C)} + (1-\eta)\, \delta^{c_*(C)} \succ_{\hat{F}'} \pi_{\hat{F}'}(E(s)) \delta^{c^*(C)} + \left(1 - \pi_{\hat{F}'}(E(s))\right) \delta^{c_*(C)} \tag{20}$$

for all $\eta \in \left(\pi_{\hat{F}'}(E(s)), \pi_{\hat{F}}(s)\right)$.

However, by (6),

$$f = \delta^{c^*(C)}{}_{\{s\}} \delta^{c_*(C)} \sim_{\hat{F}} \pi_{\hat{F}}(s)\, \delta^{c^*(C)} + \left(1 - \pi_{\hat{F}}(s)\right) \delta^{c_*(C)}.$$

Also, by (7),

$$f' = \delta^{c^*(C)}{}_{E(s)} \delta^{c_*(C)} \sim_{\hat{F}'} \pi_{\hat{F}'}(E(s))\, \delta^{c^*(C)} + \left(1 - \pi_{\hat{F}'}(E(s))\right) \delta^{c_*(C)}.$$

Thus, the preference in (19) is equivalent to $f \succ_{\hat{F}} g$, while the preference (20) is equivalent to $g' \succ_{\hat{F}'} f'$. This contradicts Axiom (A.8).

$(ii) \rightarrow (i)$ That $\succcurlyeq_{\hat{F}}$ and $\succcurlyeq_{\hat{F}'}$ satisfy (A.1) - (A.5) is an implication of the Theorem of Machina and Schmeidler (1995). Invariant risk preferences, (A.6), follows from the function V being independent of $\hat{F}$.

To show that (A.8) holds, let $f, g \in \hat{F}$ and $f', g' \in \hat{F}'$ be as in (A.8). By (6),

$$f \succcurlyeq_{\hat{F}} g \Leftrightarrow V\left(\pi_{\hat{F}}(s)\, \delta^{c^*(C)} + \left(1 - \pi_{\hat{F}}(s)\right) \delta^{c_*(C)}\right) \geq V\left(\eta \delta^{c^*(C)} + (1-\eta)\, \delta^{c_*(C)}\right).$$

By (8), the last inequality holds if and only if

$$V\left(\pi_{\hat{F}'}(E(s))\, \delta^{c^*(C)} + \left(1 - \pi_{\hat{F}'}(E(s))\right) \delta^{c_*(C)}\right) \geq V\left(\eta \delta^{c^*(C)} + (1-\eta)\, \delta_{c_*(C)}\right). \tag{21}$$

By (7), the expression in (21) is equivalent to $f' \succ_{\hat{F}'} g'$. Thus, Axiom (A.8) must hold.♠

5.3 Proof of Theorem 3

$(i) \Rightarrow (ii)$. Since $\succcurlyeq_{\hat{F}}$ and $\succcurlyeq_{\hat{F}_{S'}}$ satisfy (A.1) - (A.5), the Theorem of Machina and Schmeidler (1995) implies a representation as in (9) as well as the uniqueness of V and of $\pi_{\hat{F}}$ for each level of awareness. By (A.6), $p \succcurlyeq_{\hat{F}} q$ if and only if $p \succcurlyeq_{\hat{F}_{S'}} q$. Thus, by the uniqueness of the representations, V can be chosen to be invariant to the level of awareness.

For some $s \in S(F,C) \cap S'(F,C)$, let $g = g' = \eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}$, and f, f' be as in Axiom (A.9). Suppose that $f \sim_{\hat{F}} g$. But $f \sim_{\hat{F}} g$ if and only if

$$\delta^{c^*(C)}{}_{\{s\}}\left(\delta^{c_*(C)}{}_{S(F,C)\cap S'(F,C)}\left(\eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}\right)\right) \sim_{\hat{F}} \eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}. \quad (22)$$

By the representation in (9) the last indifference holds if and only if

$$V\Bigg(\pi_{\hat{F}}(s)\delta^{c^*(C)} + \left(\pi_{\hat{F}}(S(F,C)\cap S'(F,C)) - \pi_{\hat{F}}(s)\right)\delta^{c_*(C)} + \quad (23)$$

$$(1-\pi_{\hat{F}}(S(F,C)\cap S'(F,C)))\left(\eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}\right)\Bigg)$$

$$= V\left(\eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}\right)$$

But (23) holds if and only if $\pi_{\hat{F}}(s) + (1-\pi_{\hat{F}}(S(F,C)\cap S'(F,C)))\eta = \eta$. Hence,

$$\eta = \frac{\pi_{\hat{F}}(s)}{\pi_{\hat{F}}(S(F,C)\cap S'(F,C))}. \quad (24)$$

By Axiom (A.9), $f \sim_{\hat{F}} g$ if and only if $f' \sim_{\hat{F}_{S'}} g'$, which is equivalent to

$$\delta^{c^*(C)}{}_{\{s\}}\left(\delta^{c_*(C)}{}_{S(F,C)\cap S'(F,C)}\left(\eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}\right)\right) \sim_{\hat{F}_{S'}} \eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}. \quad (25)$$

By the representation in (10), (25) holds if and only if

$$V\Bigg(\pi_{\hat{F}_{S'}}(s)\delta^{c^*(C)} + \left(\pi_{\hat{F}_{S'}}(S(F,C)\cap S'(F,C)) - \pi_{\hat{F}_{S'}}(s)\right)\delta^{c_*(C)} \quad (26)$$

$$+(1-\pi_{\hat{F}_{S'}}(S(F,C)\cap S'(F,C)))\left(\eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}\right)\Bigg)$$

$$= V\left(\eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}\right)$$

But (26) holds if and only if $\pi_{\hat{F}_{S'}}(s) + (1-\pi_{\hat{F}_{S'}}(S(F,C)\cap S'(F,C)))\eta = \eta$. Thus, $f' \sim_{\hat{F}_{S'}} g'$ if and only if

$$\eta = \frac{\pi_{\hat{F}_{S'}}(s)}{\pi_{\hat{F}_{S'}}(S(F,C)\cap S'(F,C))}. \quad (27)$$

By (24) and (27) we have that

$$\frac{\pi_{\hat{F}}(s)}{\pi_{\hat{F}}(S(F,C)\cap S'(F,C))} = \frac{\pi_{\hat{F}_{S'}}(s)}{\pi_{\hat{F}_{S'}}(S(F,C)\cap S'(F,C))}. \tag{28}$$

An analogous argument applies for any $s' \in S(F,C) \cap S'(F,C)$. We therefore also have that, for any $s' \in S(F,C) \cap S'(F,C)$,

$$\frac{\pi_{\hat{F}}(s')}{\pi_{\hat{F}}(S(F,C)\cap S'(F,C))} = \frac{\pi_{\hat{F}_{S'}}(s')}{\pi_{\hat{F}_{S'}}(S(F,C)\cap S'(F,C))}. \tag{29}$$

Together (28) and (29) imply that

$$\frac{\pi_{\hat{F}_{S'}}(s)}{\pi_{\hat{F}_{S'}}(s')} = \frac{\pi_{\hat{F}}(s)}{\pi_{\hat{F}}(s')}. \tag{30}$$

$(ii) \to (i)$. That $\succcurlyeq_{\hat{F}}$ and $\succcurlyeq_{\hat{F}_{S'}}$ satisfy (A.1) - (A.5) is an implication of the Theorem of Machina and Schmeidler (1995). Invariant risk preferences, (A.6), follows from the function V being independent of $\hat{F}$.

To show that (A.9) holds, let $f, g \in \hat{F}$ and $f', g' \in \hat{F}_{S'}$ be as in (A.9). By (9), $f \succcurlyeq_{\hat{F}} g$ if and only if

$$V\Bigg(\pi_{\hat{F}}(s)\delta^{c^*(C)} + \left(\pi_{\hat{F}}(S(F,C)\cap S'(F,C)) - \pi_{\hat{F}}(s)\right)\delta^{c_*(C)}$$

$$+(1-\pi_{\hat{F}}(S(F,C)\cap S'(F,C)))\left(\eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}\right)\Bigg)$$

$$\geq V\left(\eta\delta^{c^*(C)} + (1-\eta)\,\delta^{c_*(C)}\right).$$

By first order stochastic dominance, the last inequality holds if and only if

$$\frac{\pi_{\hat{F}}(s)}{\pi_{\hat{F}}(S(F,C)\cap S'(F,C))} \geq \eta. \tag{31}$$

Suppose that $g' \succ_{\hat{F}_{S'}} f'$. By (10), $g' \succ_{\hat{F}_{S'}} f'$ if and only if

$$V\Bigg(\pi_{\hat{F}_{S'}}(s)\delta^{c^*(C)} + \left(\pi_{\hat{F}_{S'}}(S(F,C)\cap S'(F,C)) - \pi_{\hat{F}_{S'}}(s)\right)\delta^{c_*(C)}$$

$$+(1-\pi_{\hat{F}_{S'}}(S(F,C)\cap S'(F,C)))\left(\eta\delta^{c^*(C)}+(1-\eta)\,\delta^{c_*(C)}\right)\Big)$$

$$< V\left(\eta\delta^{c^*(C)}+(1-\eta)\,\delta^{c_*(C)}\right).$$

By first order stochastic dominance, this holds if and only if $\pi_{\hat{F}_{S'}}(s)+(1-\pi_{\hat{F}_{S'}}(S(F,C)\cap S'(F,C)))\eta<\eta$. Hence,

$$\eta > \frac{\pi_{\hat{F}_{S'}}(s)}{\pi_{\hat{F}_{S'}}(S(F,C)\cap S'(F,C))}. \tag{32}$$

Now, expressions (31) and (32) imply that

$$\frac{\pi_{\hat{F}}(s)}{\pi_{\hat{F}}(S(F,C)\cap S'(F,C))} > \frac{\pi_{\hat{F}_{S'}}(s)}{\pi_{\hat{F}_{S'}}(S(F,C)\cap S'(F,C))}. \tag{33}$$

However, by (11),

$$\frac{\pi_{\hat{F}}(s')}{\pi_{\hat{F}}(s)} = \frac{\pi_{\hat{F}_{S'}}(s')}{\pi_{\hat{F}_{S'}}(s)} \tag{34}$$

for all $s, s' \in S(F,C)\cap S'(F,C)$. Summing over $s' \in S(F,C)\cap S'(F,C)$ and rearranging, (34) implies that

$$\frac{\pi_{\hat{F}}(s)}{\pi_{\hat{F}}(S(F,C)\cap S'(F,C))} = \frac{\pi_{\hat{F}_{S'}}(s)}{\pi_{\hat{F}_{S'}}(S(F,C)\cap S'(F,C))}$$

which contradicts (33).♠

References

[1] Anscombe, F. J. and R. J. Aumann (1963). A Definition of Subjective Probability. *Annals of Mathematical Statistics*, 43, 199–205.

[2] Epstein, L. and M. Schneider (2003). Recursive Multiple-Priors. *Journal of Economic Theory*, 113, 1–31

[3] Fishburn, P. C. (1970). *Utility Theory for Decision Making.* John Wiley & Sons, New York.

[4] Grant, S. and B. Polak (2006). Bayesian Beliefs with Stochastic Monotonicity: An Extension of Machina and Schmeidler. *Journal of Economic Theory*, 130, 264–282.

[5] Gul, F. (1991). A Theory of Disappointment Aversion. *Econometrica*, 59, 667-686.

[6] Karni, E. and D. Schmeidler (1991). Utility Theory with Uncertainty. In W. Hildenbrand and H. Sonnenschein (Eds.), *Handbook of Mathematical Economics,* vol. IV. Elsevier Science Publishers B.V., Amsterdam.

[7] Karni, E. and M.-L. Vierø (2013). "Reverse Bayesianism": A Choice-Based Theory of Growing Awareness. *American Economic Review*, 103, 2790–2810.

[8] Kochov, A. (2010). A Model of Limited Foresight. Working paper, University of Rochester.

[9] Lancaster, K. J. (1966). A New Approach to Consumer Theory. *Journal of Political Economy*, 74, 132-157.

[10] Machina, M. and D. Schmeidler (1992). A More Robust Definition of Subjective Probability. *Econometrica*, 60, 745–780.

[11] Machina, M. and D. Schmeidler (1995). Bayes without Bernoulli: Simple Conditions for Probabilistically Sophisticated Choice. *Journal of Economic Theory,* 67, 106–128.

[12] Ortoleva, P. (2012). Modeling the Change of Paradigm: Non-Bayesian Reaction to Unexpected News. *American Economic Review*, 102, 2410–2436.

[13] Savage, L. J. (1954). *The Foundations of Statistics.* Dover Publications, New York.

[14] Schmeidler, D. and P. Wakker (1987). Expected Utility and Mathematical Expectation. In J. Eatwell, M. Milgate, and P. Newman (Eds.) *The New Palgrave: A Dictionary of Economics*. Macmillan Press, New York.