

Research Article

Edi Karni, Quitzé Valenzuela-Stookey and Marie-Louise Vierø*

Reverse Bayesianism: A Generalization

<https://doi.org/10.1515/bejte-2018-0176>

Received December 10, 2018; accepted March 19, 2020

Abstract: This paper studies an environment in which a decision maker choosing between acts may initially be unaware of certain consequences. We follow the approach of Karni and Vierø (2013) to modeling increasing awareness, which allows for the decision maker's state space to expand as she becomes aware of new possible consequences. We generalize the main result in Karni and Vierø (2013) by allowing the discovery of new consequences to nullify some states that were non-null before the discovery. We also provide alternative assumptions which strengthen the predictions of the belief updating model.

Keywords: growing awareness, reverse Bayesianism, falsification

1 Introduction

Karni and Vierø (2013) provides an approach to modeling increasing awareness of a decision maker. The approach allows for the decision maker's state space to expand as she becomes aware of new possible actions and consequences. They named their approach to modeling increasing awareness "reverse Bayesianism," motivated by a consistency property of the decision maker's beliefs over the state spaces before and after the expansion of her awareness. Under reverse Bayesianism the likelihood ratios between originally non-null states remain unchanged upon the discovery of new consequences and the subsequent expansion of the state space.

*Corresponding author: Marie-Louise Vierø, Queen's University, Kingston, Canada,
E-mail: viero@econ.queensu.ca

Edi Karni: Johns Hopkins University, Baltimore, MD, USA, E-mail: karni@jhu.edu

Quitzé Valenzuela-Stookey: Northwestern University, Evanston, IL, USA,
E-mail: quitze.vs@u.northwestern.edu

Karni and Vierø (2013) implicitly assumes that any state that was non-null before the increase in awareness resulting from the discovery of a new consequence has to remain non-null after the discovery. Under monotonicity this is equivalent to assuming that the prior set of non-null states – what Karni and Vierø (2013) calls the feasible state space – is included in the posterior feasible state space. However, there are important situations in which such inclusion is untenable. For example, situations of scientific discoveries that falsify prior beliefs. To illustrate, consider the famous Michelson–Morley experiment. The experiment compared the velocity of light traveling in perpendicular directions in an attempt to detect difference in the return time that would indicate motion of matter through the substance aether, which was hypothesized to fill empty space. The failure to detect such difference provided strong evidence against the aether theory, contradicted the predictions of Newtonian mechanics, and prompted research that eventually led to Einstein’s special relativity theory.

In terms of reverse Bayesianism, the a-priori (i. e., before the Michelson–Morley experiment) feasible state-space includes states in which the velocity of light obeys the of rules of Newtonian mechanics. The Michelson-Morley experiment resulted in a consequence that required a revision of the Newtonian outlook. This revision nullified some of the a-priori feasible states simultaneous to an expansion of the conceivable state space. The present paper modifies the axiomatization of Karni and Vierø (2013) so as to allow some states that were non-null before the discovery of new consequences to become null after it.

To illustrate more concretely, imagine an injured athlete who is contemplating whether or not to take a pain killer. The athlete knows that taking the drug has influence on when he will be ready to compete again. In particular, it influences whether he will be ready to compete next month. Suppose that the athlete believes that regardless of whether or not he takes the drug, he will be ready to compete next year. He thus faces a set of feasible actions $F = \{\text{take drug (D), abstain from taking drug (A)}\}$, while the set of consequences he is aware of are $C_0 = \{\text{ready to compete next month (M), ready to compete next year but not next month (Y)}\}$. Then the set of possible resolutions of uncertainty the athletes faces are given by the conceivable state space C_0^F , illustrated by the matrix (1):

$$\begin{array}{cccc}
 & S_1 & S_2 & S_3 & S_4 \\
 D & M & M & Y & Y \\
 A & M & Y & M & Y
 \end{array} \tag{1}$$

Suppose that the athlete considers it to be impossible to compete next month if he doesn’t take the drug and for there to be a 50-50 chance of being ready in a month if

he takes the drug. Thus, the set of non-null states, referred to as the feasible state space, is $\{s_2, s_4\}$ with $\pi_0(s_2) = \pi_0(s_4) = .5$.

Consider the possibility that the athlete discovers that the drug is highly addictive and taking it means that he will not be able to compete next month due to the need for rehabilitation and possibly he will never be able to compete again. The new set of feasible consequences is $C_1 = \{\text{ready to compete next month (M), ready to compete next year but not next month (Y), never compete again (N)}\}$. The new conceivable state space C_1^F is

$$\begin{array}{cccccccccc}
 & s_1 & s_2 & s_3 & s_4 & s_5 & s_6 & s_7 & s_8 & s_9 \\
 D & M & M & Y & Y & M & Y & N & N & N \\
 A & M & Y & M & Y & N & N & M & Y & N
 \end{array} \quad (2)$$

In the wake of the discovery of the addictive nature of the drug, the new feasible state space is $\{s_4, s_8\}$. Therefore, simultaneously to the expansion of the conceivable state space, the prior feasible state s_2 becomes null.

The paper is organized as follows: Section 2 provides the framework, assumptions, and results, while Section 3 discusses the results and provides an example of extreme belief revision within the model. Proofs are collected in the appendix.

2 The Main Result

2.1 Preliminaries

We briefly restate the framework of Karni and Vierø (2013). Let F be a finite, nonempty set of *feasible actions*, and 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.

On this conceivable state space, we define what we refer to as *conceivable acts*. Formally,

$$\hat{F} := \{f : C^F \rightarrow \Delta(C)\}, \quad (3)$$

where $\Delta(C)$ is the set of all lotteries over C . As is usually done, we abuse notation and use p to also denote the constant act that returns the lottery p in each state. We use both c and δ_c to denote the lottery that returns consequence c with probability 1, depending on the context.

Discovery of new consequences expands the conceivable state space. Let C denote the initial set of consequences and suppose that a new consequence, \tilde{c} , is discovered. The set of consequences of which the decision maker is aware then expands to $C' = C \cup \{\tilde{c}\}$. As a result, the conceivable state space expands to $(C')^F$. The corresponding expanded set of conceivable acts is given by

$$\hat{F}^* := \left\{ f : (C')^F \rightarrow \Delta(C') \right\}. \quad (4)$$

We consider a decision maker whose choice behavior is characterized by a preference relation $\succsim_{\hat{F}}$ on the set of conceivable acts \hat{F} . We denote by $\succ_{\hat{F}}$ and $\sim_{\hat{F}}$ the asymmetric and symmetric parts of $\succsim_{\hat{F}}$ with the interpretations of strict preference and indifference, respectively. For any $f \in \hat{F}$, $p \in \Delta(C)$, and $E \subseteq C^F$, let $p_{E}f$ be the act in \hat{F} obtained from f by replacing its s -th coordinate with p for all $s \in E$. A state $s \in C^F$ is said to be *null* if $p_s f \sim_{\hat{F}} q_s f$ for all $p, q \in \Delta(C)$ and 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. Henceforth we refer to $S(F, C)$ as the *feasible state space*.

When the state space expands so does the set of conceivable acts, which means that the preference relation must be redefined on the extended domain. Specifically, if \hat{F}^* is the expanded set of conceivable acts in the wake of discoveries of new feasible consequences, then the corresponding preference relation is denoted by $\succsim_{\hat{F}^*}$. Let \mathcal{F} be a family of sets of conceivable acts corresponding to increasing awareness of consequences.¹

In Karni and Vierø (2013) it is implicitly assumed that upon the expansion of the state space following the discovery of a new consequence, non-null states remain non-null. Formally, for all $f \in \hat{F}$ and $f' \in \hat{F}^*$, if $p_s f \succ_{\hat{F}} q_s f$ then $p_s f' \succ_{\hat{F}^*} q_s f'$, for all $s \in S(F, C)$. Under monotonicity this is equivalent to assuming that $S(F, C) \subseteq S(F, C')$. However, as we discussed in the introduction, there are important situations in which such inclusion is untenable.

For each $\hat{F} \in \mathcal{F}$, $f, g \in \hat{F}$, 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 in a linear space.² We assume that, for each $\hat{F} \in \mathcal{F}$, $\succsim_{\hat{F}}$ abides by the axioms of Anscombe and Aumann (1963). Formally,

¹ For the preference relation $\succsim_{\hat{F}}$ as well as for preference relations associated with any other awareness levels, we make the corresponding definitions to those stated in the previous paragraph.

² Throughout this paper we use Fishburn's (1970) formulation of Anscombe and Aumann (1963). According to this formulation, mixed acts, (that is, $\alpha f + (1 - \alpha)g$) are, by definition, conceivable acts.

- (A.1) (*Weak order*) For all $\widehat{F} \in \mathcal{F}$, the preference relation $\succsim_{\widehat{F}}$ is transitive and complete.
- (A.2) (*Archimedean*) For all $\widehat{F} \in \mathcal{F}$ and $f, g, h \in \widehat{F}$, if $f \succ_{\widehat{F}} g$ and $g \succ_{\widehat{F}} h$ then $\alpha f + (1 - \alpha)h \succ_{\widehat{F}} g$ and $g \succ_{\widehat{F}} \beta f + (1 - \beta)h$, for some $\alpha, \beta \in (0, 1)$.
- (A.3) (*Independence*) For all $\widehat{F} \in \mathcal{F}$, $f, g, h \in \widehat{F}$, and $\alpha \in (0, 1]$, $f \succ_{\widehat{F}} g$ if and only if $\alpha f + (1 - \alpha)h \succ_{\widehat{F}} \alpha g + (1 - \alpha)h$.
- (A.4) (*Monotonicity*) For all $\widehat{F} \in \mathcal{F}$, $f \in \widehat{F}$, $p, q \in \Delta(C)$ and nonnull event $E \subseteq C^F$, $p_E f \succ_{\widehat{F}} q_E f$ if and only if $p \succ_{\widehat{F}} q$.
- (A.5) (*Nontriviality*) For all $\widehat{F} \in \mathcal{F}$, $\succ_{\widehat{F}} \neq \emptyset$.

Karni and Vierø (2013) postulated the following awareness consistency axiom to characterize the decision maker's reaction to expansion in his awareness of consequences:

- (A.7) (*Awareness consistency*) For every given F , for all C, C' with $C \subset C'$, $S(F, C) \subseteq S(F, C')$, $f, g \in \widehat{F}$, and $f', g' \in \widehat{F}^*$, such that $f' = f$ and $g' = g$ on $S(F, C)$ and $f' = g'$ on $S(F, C') - S(F, C)$ it holds that $f \succ_{\widehat{F}} g$ if and only if $f' \succ_{\widehat{F}^*} g'$.

In words, axiom (A.7) posits that preferences conditional on events that are considered feasible before the expansion of awareness remain unchanged when awareness expands.

To allow for the possibility that some non-null states become null upon the discovery of new consequences, we modify the awareness consistency axiom (A.7) by restricting the events conditional on which preferences are required to remain unchanged when awareness expands to events that are considered feasible both before and after the awareness expansion.

- (A.7r) (*Revised awareness consistency*) For every given F , for all C, C' with $C \subset C'$, and for $f, g \in \widehat{F}$, and $f', g' \in \widehat{F}^*$, such that $f' = f$ and $g' = g$ on $S(F, C) \cap S(F, C')$, $f = g$ on $S(F, C) - [S(F, C) \cap S(F, C')]$ and $f' = g'$ on $S(F, C') - [S(F, C) \cap S(F, C')]$ it holds that $f \succ_{\widehat{F}} g$ if and only if $f' \succ_{\widehat{F}^*} g'$.

To grasp the nature of the extension of the “Reverse Bayesianism” result, consider the awareness consistency axiom (A.7). It is more restrictive than axiom (A.7r) in the sense that it requires that equalities $f' = f$ and $g' = g$ hold on the set $S(F, C)$ while (A.7r) requires that these equalities only hold on the event $S(F, C') \cap S(F, C)$. This difference is significant because if, in (A.7r), we do not insist that $f = g$ on $S(F, C) - [S(F, C) \cap S(F, C')]$, there could arise a preference reversal $g \succ_{\widehat{F}} f$ and $f' \succ_{\widehat{F}^*} g'$. This reversal reflects the fact that the decision maker initially believes that all the states in $S(F, C)$ are nonnull, and following the revision of her beliefs in the wake of the discovery of new consequences, some

states in $S(F, C)$ become null. Hence, it is possible that initially $g \succ_{\hat{F}} f$ because the payoffs of g in the states that would later become null make it more attractive than f but once these states are nullified, f dominates g . The awareness consistency axiom (A.7) is compelling if none of the states in $S(F, C)$ becomes null after the new consequences are discovered.

2.2 Representation theorem

Dominiak and Tserenjigmid (2018) have shown that the invariant risk preferences axiom (A.6) in Karni and Vierø (2013) is redundant. Therefore, in addition to invoking the revised awareness consistency axiom, we state and prove the theorem below without the invariant risk preferences axiom. Since the proof in Dominiak and Tserenjigmid (2018) was based on the awareness consistency axiom, the revised awareness consistency axiom requires a proof of the below theorem that is slightly different from the proofs in both Karni and Vierø (2013) and Dominiak and Tserenjigmid (2018).

Theorem. For each $\hat{F} \in \mathcal{F}$, let $\succ_{\hat{F}}$ be a binary relation on \hat{F} then, for all $\hat{F}, \hat{F}^* \in \mathcal{F}$, the following two conditions are equivalent:

- i. Each $\succ_{\hat{F}}$ satisfies (A.1) – (A.5) and, jointly, $\succ_{\hat{F}}$ and $\succ_{\hat{F}^*}$ satisfy (A.7r).
- ii. There exist real-valued, non-constant, affine functions, U on $\Delta(C)$ and U^* on $\Delta(C')$, and for any two $\hat{F}, \hat{F}^* \in \mathcal{F}$, there are probability measures, $\pi_{\hat{F}}$ on $C^{\hat{F}}$ and $\pi_{\hat{F}^*}$ on $(C')^{\hat{F}^*}$, such that for all $f, g \in \hat{F}$,

$$f \succ_{\hat{F}} g \Leftrightarrow \sum_{s \in C^{\hat{F}}} U(f(s)) \pi_{\hat{F}}(s) \geq \sum_{s \in C^{\hat{F}}} U(g(s)) \pi_{\hat{F}}(s). \quad (5)$$

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

$$f' \succ_{\hat{F}^*} g' \Leftrightarrow \sum_{s \in (C')^{\hat{F}^*}} U^*(f'(s)) \pi_{\hat{F}^*}(s) \geq \sum_{s \in (C')^{\hat{F}^*}} U^*(g'(s)) \pi_{\hat{F}^*}(s). \quad (6)$$

Moreover, U and U^* are unique up to positive linear transformations, and there exists such transformations for which $U(p) = U^*(p)$ for all $p \in \Delta(C)$. The probability distributions $\pi_{\hat{F}}$ and $\pi_{\hat{F}^*}$ are unique, $\pi_{\hat{F}}(S(F, C)) = \pi_{\hat{F}^*}(S(F, C')) = 1$, and

$$\frac{\pi_{\hat{F}}(s)}{\pi_{\hat{F}}(s')} = \frac{\pi_{\hat{F}^*}(s)}{\pi_{\hat{F}^*}(s')}, \quad (7)$$

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

Proof: See the appendix.

Note that if $S(F, C) \subseteq S(F, C')$, then $S(F, C) \cap S(F, C') = S(F, C)$, and we have the result in Theorem 1 in Karni and Vierø (2013).

The generalization in the theorem allows for a wider range of applications of reverse Bayesianism than Karni and Vierø (2013). The generalization permits a particular type of belief revision on the prior feasible state space, namely extreme belief revision in which a state is nullified. For prior feasible states that are still considered feasible after the expansion in awareness, beliefs are updated according to reverse Bayesianism. This can be justified on the ground that it is possible to falsify a hypothesis, but one can only gain evidence that supports that something is true. Therefore, it is reasonable that one would nullify a state when presented with evidence that falsifies it but that with other types of evidence one maintains the relative beliefs.

The simultaneous expansion of the state space and nullification of prior feasible events may reflect an implicit belief of the DM in an underlying theory, or causal relation, under which the newly discovered consequences are inconsistent with particular actions resulting in some previously known consequences. Such theories are themselves subjective views, or interpretations, of the world that are manifested in the revision of beliefs. Because of their subjective nature these subjective interpretations are not formalized in this model. The examples in the introduction illustrates this point. The observations about the velocity of light in the Michelson-Morley experiment are inconsistent with Newtonian mechanics, therefore, once accepted as valid, it compelled nullifying states in which the experiment yield consequences implied by the aether theory. In the athlete example, there is an underlying cause, namely addiction, that simultaneously adds the new consequence of never competing again and nullifies the existing state in which the athlete can compete next month.

Another example is the theory of evolution and the subsequent discovery of genetics. These findings contradict Lamarck's hypothesis that an organism can pass on characteristics acquired through use or disuse during its lifetime to its offspring. In the 1930s, long after the discovery of genetics, Lysenko revived the ideas of Lamarck in the Soviet Union with Lysenkoism. Lysenkoism influenced Soviet agricultural policy and was later blamed for crop failures. This example illustrates that while the modern theory of genetics is interpreted by some as ruling out Lamarckism, others find it compatible with Lamarckism. Similarly, the evidence supporting evolution produced by research in molecular biology convinced some people that Darwin's theory is valid, while others continue to believe in intelligent design. These examples serve to illustrate the subjective nature of belief revision and nullification of states. Such beliefs can presumably be detected and

quantified by the odds a decision maker would be willing to place on the outcomes of experiments testing the prediction of the underlying hypothesis.

Karni and Vierø (2013) imposed the assumption that $S(F, C) \subseteq S(F, C')$. A strengthening of axiom (A.7r), in conjunction with monotonicity, will imply that nullification of prior feasible states will not occur. Axiom (A.7r') below gives a preference-based condition under which the assumption in Karni and Vierø (2013) is satisfied.

(A.7r') For every given F , for all C, C' with $C \subset C'$, $f, g \in \widehat{F}$ and $f', g' \in \widehat{F}^*$, such that $f' = f$ and $g' = g$ on $S(F, C') \cap S(F, C)$ and $f' = g'$ on $S(F, C') - [S(F, C) \cap S(F, C')]$ it holds that $f \succsim_F g$ if and only if $f' \succsim_F g'$.

Proposition 1

If \succsim_F and \succsim_{F^*} satisfy axioms (A.7r') and (A.4) then $S(F, C) \subseteq S(F, C')$.

Proof: See the appendix.

The difference between axioms (A.7r) and (A.7r') is that axiom (A.7r) allows for preference reversals for acts that differ on $S(F, C) - [S(F, C') \cap S(F, C)]$, while axiom (A.7r') does not. By forcing the prior and posterior preference relations to agree on the ranking of acts, even if the acts differ in that event, axiom (A.7r') together with monotonicity, has the implication that $S(F, C) - S(F, C')$ must be null under the posterior preference relation.

Karni and Vierø (2013) allows for the possibility of states that were previously conceivable but null to become feasible upon the discovery of new consequences. For example, before the discovery that mosquitoes are carriers of malaria and yellow fever, the consequence of developing these diseases as a result of exposure to mosquitoes was considered a null event. The discovery of the germs mosquitoes carry, which was a new consequence of the exposure, turned a null event into a non-null event.

The following revision of (A.7r') rules out such belief revisions. That is, it implies that prior conceivable but null states remain null. In this case the likelihood ratios of all original conceivable states remain unchanged when beliefs are updated according to reverse Bayesianism.

(A.7r'') For every given F , for all C, C' with $C \subset C'$, $f, g \in \widehat{F}$ and $f', g' \in \widehat{F}^*$, such that $f' = f$ and $g' = g$ on $S(F, C') \cap S(F, C)$ and $f' = g'$ on $S(F, C') - [S(F, C') \cap C^F]$ it holds that $f \succsim_F g$ if and only if $f' \succsim_{F^*} g'$.

Proposition 2

Let $C \subset C'$, then (A.7r'') and (A.4) imply that $S(F, C') \cap C^F = S(F, C)$.

Proof: See the appendix.

Since states in $C^F - S(F, C)$ are $\succsim_{\hat{F}}$ -null, $f_{S(F,C)}h \succsim_{\hat{F}} g_{S(F,C)}h'$ if and only if $f_{S(F,C)}k \succsim_{\hat{F}} g_{S(F,C)}k'$ for all $h, h', k, k' \in \hat{F}$. In other words, if acts differ on $C^F - S(F, C)$, it does not affect their ranking according to the prior preferences. However, if a state $s \in C^F - S(F, C)$ is non-null according to the posterior, what the acts return in that state affect their ranking. In contrast to (A7r'), (A7r'') rules out preference reversals for acts that differ on $C^F - S(F, C)$. Therefore, any such state must be $\succsim_{\hat{F}^*}$ -null to preserve the ranking.

Axioms (A.7r), (A.7r'), and (A.7r'') illustrate how assuming consistency of $\succsim_{\hat{F}}$ and $\succsim_{\hat{F}^*}$ on a larger set of acts shrinks the set of states on which extreme belief revisions may occur. Which of these assumptions is relevant will depend on the setting.

3 Discussion

This paper provides a model of belief updating under increasing awareness that is due to the discovery of new consequences. Within the framework of reverse Bayesianism we axiomatize a representation consistent with the observation that the discovery of new consequences may contradict the belief that certain states may occur. The result is nullification of prior feasible states simultaneous to the addition of newly discovered states.

An alternative to the simultaneous addition and nullification of states is a two-step procedure according to which the first step expands the feasible state space to include the newly discovered states, and the second step nullifies prior feasible states that are considered no longer feasible and, at the same time, updates the probabilities of the feasible states using Bayes rule. Given a prior, the simultaneous addition and nullification of states and the two-step procedure can generate identical posterior beliefs. However, whether updating is done in one or two steps could matter depending on the timing of subsequent decisions. Consider the example of the athlete from the introduction. Given the posterior, taking the drug is dominated by not taking it, since if the athlete takes the drug, he will either compete in a year or never, while if he doesn't take the drug, he will for sure compete in a year. With the two-step procedure, if a decision is to be made between the moment when the athlete discovers that taking the drug might prevent him from competing again (that is, the time when the new states are incorporated into the athlete's perception of the world) and the realization, sometime later, that taking the drug will prevent him from competing in a month (which leads to posterior updating) the athlete may still prefer to take the drug, since he assigns a positive probability to the event that will have him race ready in a month.

The issue of whether the nullification of the states occurs simultaneously with the discovery of new consequences or after a delay may depend on the particularity of the situation at hand. In principle, given a particular situation, the issue can be resolved by conducting an experiment designed to elicit the odds placed on the states immediately following the discovery of the new consequence. Such experiment would determine if some of the original states become null immediately. If not, repeating the experiment after a delay could show if updating of beliefs occurs separately from the discovery of the new states.

Karni and Vierø (2013) considered discovery of new consequences without severing existing links between acts and consequences. They also considered situations in which links between acts and consequences that the decision maker believed possible are discovered to be impossible, without a simultaneous discovery of a new consequence. The latter realization renders null some states that were considered feasible without altering the conceivable state space. This paper expands the results to situations in which the discovery of a new consequence simultaneously severs a link. Thus, in this paper the severance of act-consequence links is the result of a discovery of a new consequence that entails an expansion of the conceivable state space.

These extreme belief revisions differ conceptually from intermediate revisions. Whereas new information may provide logical evidence falsifying a null hypothesis, it will only provide supporting evidence that it is true. Propositions 1 and 2 illustrate the subtle nature of such extreme revisions. These results highlight the distinction between extreme belief revisions that support the possibility of a state and those that nullify it.

Acknowledgments: Vierø acknowledges financial support from SSHRC grant number 435-2014-0642.

A Proofs

A.1 Proof of Theorem

(Sufficiency) Fix F and C . By (A.1)–(A.5) and the von Neumann-Morgenstern expected utility theorem, there exists a real-valued, non-constant, function u_F^\wedge on C such that for all $p, q \in \Delta(C)$

$$p \succcurlyeq_F q \Leftrightarrow \sum_{c \in \text{Supp}(p)} u_F^\wedge(c) p(c) \geq \sum_{c \in \text{Supp}(q)} u_F^\wedge(c) q(c). \quad (8)$$

Let $C' \supset C$ and $\widehat{F}^* \in \mathcal{F}$. Then, by the same argument as above, there exists a real-valued function $u_{\widehat{F}^*}$ on C' such that for all $p', q' \in \Delta(C')$

$$p' \succsim_{\widehat{F}^*} q' \Leftrightarrow \sum_{c \in \text{Supp}(p')} u_{\widehat{F}^*}(c) p'(c) \geq \sum_{c \in \text{Supp}(q')} u_{\widehat{F}^*}(c) q'(c). \quad (9)$$

The functions $u_{\widehat{F}}$ and $u_{\widehat{F}^*}$ are unique up to positive linear transformations. Define $U(f(s)) := \sum_{c \in \text{Supp}(f(s))} u_{\widehat{F}}(c) f(s)(c)$, for all $f \in \widehat{F}$ and $s \in S(\widehat{F}, C)$. Thus, $U(f(s))$ is the von Neumann-Morgenstern utility of the lottery $f(s)$. Define also $U^*(f(s)) := \sum_{c \in \text{Supp}(f(s))} u_{\widehat{F}^*}(c) f(s)(c)$, for all $f \in \widehat{F}^*$ and $s \in S(\widehat{F}^*, C')$.

The remainder of the proof follows steps similar to steps in Dominiak and Tserenjigmid (2018), but needs more elaboration to prove different properties on different subsets of states. In particular, the steps in (11) through (15) show that the properties hold on the the intersection of the feasible state spaces. Let b and w be a best and worst consequence in C , respectively. Without loss of generality, normalize $u_{\widehat{F}}(b) = u_{\widehat{F}^*}(b) = 1$ and $u_{\widehat{F}}(w) = u_{\widehat{F}^*}(w) = 0$.

Take any lottery $q \in \Delta(C)$ and acts $f, g \in \widehat{F}$ and $f', g' \in \widehat{F}^*$ such that $f' = f$ and $g' = g = q$ on $S(F, C) \cap S(F, C')$, $f = g = \delta_w$ on $S(F, C) - [S(F, C) \cap S(F, C')]$ and $f' = g' = \delta_w$ on $S(F, C') - [S(F, C) \cap S(F, C')]$. Then, by the Anscombe and Aumann (1963) theorem and axiom (A.7r), we have that $f \succsim_{\widehat{F}} g$ if and only if $f' \succsim_{\widehat{F}^*} g'$, which implies that

$$\begin{aligned} \sum_{s \in S(F, C) \cap S(F, C')} U(f(s)) \pi_{\widehat{F}}(s) &= U(q) \pi_{\widehat{F}}(S(F, C) \cap S(F, C')) \\ \Leftrightarrow \sum_{s \in S(F, C) \cap S(F, C')} U^*(f(s)) \pi_{\widehat{F}^*}(s) &= U^*(q) \pi_{\widehat{F}^*}(S(F, C) \cap S(F, C')) \end{aligned} \quad (10)$$

The next step will show that beliefs are updated according to reverse Bayesianism for all states in $S(F, C) \cap S(F, C')$. For any $s \in S(F, C) \cap S(F, C')$, let $f(s) = \delta_c$ for some $c \in C$ and $f(\tilde{s}) = \delta_w$ for all $\tilde{s} \neq s$. Let $q = \frac{\pi_{\widehat{F}}(s)}{\pi_{\widehat{F}}(S(F, C) \cap S(F, C'))} \delta_c + 1 - \frac{\pi_{\widehat{F}}(s)}{\pi_{\widehat{F}}(S(F, C) \cap S(F, C'))} \delta_w$. Then the utility of f is given by

$$\sum_{s \in S(F, C)} U(f(s)) \pi_{\widehat{F}}(s) = \pi_{\widehat{F}}(s) U(\delta_c) \quad (11)$$

while the utility of g is given by

$$\sum_{s \in S(F, C)} U(g(s)) \pi_{\widehat{F}}(s) = \frac{\pi_{\widehat{F}}(s)}{\pi_{\widehat{F}}(S(F, C) \cap S(F, C'))} U(\delta_c) \pi_{\widehat{F}}(S(F, C) \cap S(F, C')). \quad (12)$$

Equations (11) and (12) imply that $f \sim_{\widehat{F}} g$.

The utility of f^* is given by

$$\sum_{s \in S(F, C')} U^*(f'(s)) \pi_F^{\sim}(s) = \pi_F^{\sim}(s) U^*(\delta_c) \quad (13)$$

while the utility of g' is given by

$$\sum_{s \in S(F, C')} U^*(g'(s)) \pi_F^{\sim}(s) = \frac{\pi_F^{\sim}(s)}{\pi_F^{\sim}(S(F, C) \cap S(F, C'))} U^*(\delta_c) \pi_F^{\sim}(S(F, C) \cap S(F, C')). \quad (14)$$

By axiom (A.7r), $f' \sim_F g'$, since $f \sim_F g$. Hence, (10) implies that

$$\pi_F^{\sim}(s) U^*(\delta_c) = \frac{\pi_F^{\sim}(s)}{\pi_F^{\sim}(S(F, C) \cap S(F, C'))} U^*(\delta_c) \pi_F^{\sim}(S(F, C) \cap S(F, C')). \quad (15)$$

Equivalently,

$$\frac{\pi_F^{\sim}(s)}{\pi_F^{\sim}(S(F, C) \cap S(F, C'))} = \frac{\pi_F^{\sim}(s)}{\pi_F^{\sim}(S(F, C) \cap S(F, C'))} \quad (16)$$

Next we show that under the normalization of the utility functions u_F^{\sim} and $u_{F^*}^{\sim}$ it holds that $U^*(\delta_c) = U(\delta_c)$ for all $c \in C$. For any $c \in C$ and $s \in S(F, C) \cap S(F, C')$, let $f(s) = \alpha \delta_c + (1 - \alpha) \delta_w$ and $f(\tilde{s}) = \delta_w$ for all $\tilde{s} \neq s$. Let $q = \beta \delta_b + (1 - \beta) \delta_w$. By the normalization, $U(\delta_b) = U^*(\delta_b) = 1$. Hence, (10) implies that

$$\alpha U(\delta_c) \pi_F^{\sim}(s) = \beta \pi_F^{\sim}(S(F, C) \cap S(F^*, C')) \quad (17)$$

$$\Leftrightarrow \alpha U^*(\delta_c) \pi_F^{\sim}(s) = \beta \pi_F^{\sim}(S(F, C) \cap S(F^*, C')) \quad (18)$$

Equations (16), (17), and (18) now imply that $U^*(\delta_c) = U(\delta_c)$ for all $c \in \Delta(C)$.

(Necessity) The necessity of (A.1)–(A.5) is an implication of the Anscombe and Aumann (1963) theorem. The necessity of (A.7r) is immediate.

The uniqueness part is an implication of the uniqueness of the utility and probability in Anscombe and Aumann (1963).

A.2 Proof of Proposition 1

Suppose there exists s such that $s \in S(F, C)$, $s \notin S(F, C')$. Let $p, q \in \Delta(C)$ be such that $p \succ_F q$ (Note that if no such p, q exist then by Monotonicity $S(F, C) = \emptyset$, so the

result is trivial). Define the conceivable acts f, g , and f', g' as follows: $f = g$ on $S(F, C) \setminus s$, $f(s) = p$, $g(s) = q$. Furthermore $f' = f$, $g' = g$ on $S(F, C)$, and $f' = g'$ on $S(F, C') - [S(F, C') \cap S(F, C)]$. Note that f' and g' coincide exactly on $S(F, C')$, so $f' \sim_F g'$. Moreover, by Monotonicity $f \succ_F g$. Since f, g, f', g' satisfy the conditions of (A.7r') we have a contradiction.

A.3 Proof of Proposition 2

That $S(F, C) \subseteq S(F, C') \cap C^F$ follows from Proposition 1, since (A7r'') is a strengthening of (A7r').

It remains to show that $S(F, C') \cap C^F \subseteq S(F, C)$. Suppose there exists s such that $s \in S(F, C') \cap C^F$, $s \notin S(F, C)$. Let $p, q \in \Delta(C)$ be such that $p \succ_F q$ (if no such p, q exist then, by (A.4), $S(F, C') = \emptyset$, so the result is trivial). Define the conceivable acts f, g , and f', g' as follows: $f = g$ on $C^F \setminus s$, $f(s) = p$, $g(s) = q$. Furthermore $f' = f$, $g' = g$ on C^F , and $f' = g'$ on $S(F, C') - [S(F, C') \cap C^F]$. But f and g coincide on $S(F, C)$, hence, $f \sim_F g$. By (A.4) $f' \succ_F g'$. Since f, g, f', g' satisfy the conditions of (A.7r'') we have a contradiction.

References

- Anscombe, Francis J. and Robert J. Aumann. 1963. "A Definition of Subjective Probability." *Annals of Mathematical Statistics* 43: 199–205.
- Dominiak, Adam and Gerelt Tserenjigmid. 2018. "Belief Consistency and Invariant Risk Preferences." *Journal of Mathematical Economics* 79: 157–162. <https://doi.org/10.1016/j.jmatheco.2018.03.005>.
- Fishburn, Peter. 1970. *Utility Theory for Decision Making*. New York: Wiley.
- Karni, Edi and Marie-Louise Vierø. 2013. "Reverse Bayesianism": A Choice-Based Theory of Growing Awareness. *American Economic Review* 103, 2790–810. <http://doi.org/10.1257/aer.103.7.2790>.