Available online at www.sciencedirect.com

ScienceDirect

Journal of Economic Theory 168 (2017) 301–328

JOURNAL OF Economic Theory

www.elsevier.com/locate/jet

Awareness of unawareness: A theory of decision making in the face of ignorance ☆

Edi Karni [a,b,1], Marie-Louise Vierø [c,*,2]

[a] *Johns Hopkins University, Department of Economics, United States*

[b] *University of Warwick, Warwick Business School, UK*

[c] *Queen's University, Department of Economics, Dunning Hall Room 312, 94 University Avenue, Kingston, Ontario K7L 3N6, Canada*

Received 18 April 2016; final version received 12 December 2016; accepted 26 December 2016

Available online 3 January 2017

Abstract

In the wake of growing awareness, decision makers anticipate that they might acquire knowledge that, in their current state of ignorance, is unimaginable. Supposedly, this anticipation manifests itself in the decision makers' choice behavior. In this paper we model the anticipation of growing awareness, lay choice-based axiomatic foundations to a subjective expected utility representation of beliefs about the likelihood of discovering unknown consequences, and assign utility to consequences that are not only unimaginable but may also be nonexistent. In so doing, we maintain the flavor of reverse Bayesianism of Karni and Vierø (2013, 2015).

JEL classification: D8; D81; D83

Keywords: Awareness; Unawareness; Ignorance; Reverse Bayesianism; Utility of indescribable consequences

☆ We benefited from comments and suggestions of two anonymous referees and Marciano Siniscalchi, for which we are grateful.

* Corresponding author.

E-mail addresses: karni@jhu.edu (E. Karni), viero@econ.queensu.ca (M.-L. Vierø).

[1] Work on this project was done, in part, while I visited EIEF. I am grateful to the EIEF for its support and hospitality.

[2] I am grateful for financial support from SSHRC grant number 435-2014-0642.

http://dx.doi.org/10.1016/j.jet.2016.12.011

1. Introduction

Whether by chance or by design, actions and scientific experiments sometimes result in consequences that, prior to their discovery, were unimaginable or, for lack of appropriate language, indescribable. Examples include the discovery of HIV and the discovery of the structure of DNA. Habituation to such consequences is an important aspect of human experience, and the anticipation of additional such discoveries shapes our future outlook, manifesting itself in our choice behavior.

In this paper, which builds on Karni and Vierø (2013, 2015), we propose a choice-based theory that captures a decision maker's anticipation of becoming aware of consequences that she is currently unaware of and analyze its behavioral implications. Our presumptions are that although a decision maker cannot know what it is that she does not know, she can entertain the belief that there are unimaginable aspects of the universe yet to be discovered, and that this belief manifests itself in her choice behavior. Because we adhere to the revealed preference methodology, we require that the decision maker's choice set consists only of objects that are well-defined given her level of awareness. In other words, when uncertainty resolves, it must be possible to meaningfully settle any bet or trade that the decision maker may have engaged in.

The main thrust of Karni and Vierø (2013, 2015) is the evolution of decision makers' beliefs as they become aware of new acts, consequences, and the links among them. In these models, however, decision makers can be interpreted as being myopic, believing themselves, at every stage, to be fully aware of the scope of their universe. Formally, in these models, decision makers act as if they consider the state space that resolves the uncertainty associated with the feasible courses of action and consequences of which they are aware, to be a sure event. Consequently, even though it happened before, decision makers fail to anticipate the possibility of discoveries that would require expansions of the state space. In a major break with our earlier work, this paper extends the analytical framework to incorporate decision makers' awareness of their potential ignorance, and the anticipation that actions may reveal consequences that were unspecified in the original formulation of the decision problem. The resulting state space is partitioned into a set of fully describable states and a set of states that are only partially describable or nondescribable. By contrast, in our earlier work the state space consisted solely of fully describable states. We discuss this issue in further details in section 2.1, following the construction of the state space, and again in section 3.1, where we illustrate how the predictions of choice behavior of the two models might differ.

This work also departs from the analytical framework we employed before in a different respect. Specifically, in Karni and Vierø (2013, 2015) the state space, constructed from finite sets of feasible acts and consequences, is finite, and the choice set consisted of conceivable Anscombe–Aumann (1963) acts, (that is, mappings from the state space to the set of lotteries on the feasible consequences). This formulation is based on the tacit assumption that the decision maker can conceive of acts whose state-contingent payoffs are lotteries on the set of feasible consequences. While analytically convenient, this construction is not entirely satisfactory. If lotteries on feasible consequences instead of the feasible consequences themselves are used to construct the state space then, by construction, the state space is infinite. This would complicate the analysis. To avoid the aforementioned inconsistency and, at the same time, to maintain the finiteness of the state space, in this work we redefine conceivable acts to be functions from states to feasible consequences. We then assume that decision makers can imagine choosing among conceivable acts randomly. Hence, the choice space is the set of probability distributions over the conceivable acts, dubbed mixed conceivable acts.

Within the new analytical framework we develop an axiomatic model of choice under uncertainty and analyze the behavioral implications of a decision maker's awareness of her unawareness. The sense that there might be consequences, lurking in the background, of which one is unaware may inspire fear or excitement, thereby affecting individual choice behavior. Our model assigns utility to the unknown consequences, thereby capturing the decision maker's attitude toward the discovery of unknown, or indescribable, consequences and the emotions it evokes. For instance, if the predominant emotion evoked by the unknown is fear, then confidence that one is unlikely to encounter unknown consequences would beget boldness of action while the lack of it would induce more prudent behavior.

To represent the attitude toward unawareness, we need to enrich the framework of Karni and Vierø (2013, 2015). In particular, because we require that bets should be possible to settle once uncertainty resolves, decision makers cannot meaningfully form preferences over acts that assign indescribable consequences to fully describable states. Therefore, to represent the attitudes toward indescribable consequences, we expand the set of conceivable acts to include acts that assign, to partially describable states only, consequences that will be discovered if these states obtain. The resulting model is a generalization of subjective expected utility with an extra parameter that captures the decision maker's "utility of the unknown." Comparing two decision makers, the one with the higher value of this utility of the unknown exhibits excitement, or optimism, toward the unknown, relative to the decision maker with the lower value, which reflects fear, or pessimism. The representation thus allows us to explicitly and formally express this attitude toward the unknown.

Another main thrust of this work is the analysis of the evolution of the decision maker's beliefs about her ignorance in the wake of the discovery of new consequences. We show that, with respect to such discoveries, the model of reverse Bayesian updating of Karni and Vierø (2013) is a special case of the present one. Furthermore, depending on the nature of the discoveries, the sense of ignorance, or the 'residual' unawareness, may shrink, grow, or remain unchanged. For instance, as unsuspected regions of the Earth or the solar system were discovered, fewer regions remained to be explored, and the sense of ignorance diminished. By contrast, some scientific discoveries, such as atoms or the structure of the DNA, resolved certain outstanding issues in physics and biology and, at the same time, opened up new vistas. These discoveries enhanced the sense that our ignorance is, in fact, greater than what was previously believed. Our model is designed to accommodate all the aforementioned possibilities of evolution of the sense of ignorance.

On a more mundane level, decision makers are routinely confronted with the need to make decisions in specific situations. For example, a decision maker about to embark on a trip must choose a means of transportation to get from here to there, or, following a diagnosis of illness, a decision maker must decide which treatment to seek. It is natural to approach such decisions by identifying the relevant courses of action and the outcomes that these actions may produce. It might happen, however, that due to lack of imagination or insufficient attention, the chosen course of action results in an outcome that the decision maker has failed to consider. Therefore, when facing a specific decision, a decision maker worries that she might fail to take into account all the relevant outcomes. The awareness that an outcome that should have been considered is, inadvertently, neglected, bears resemblance to awareness of unawareness and it similarly affects individual choice behavior. We discuss this similarity between awareness of unawareness and "small worlds" in further detail in the concluding remarks.

In the next section we present the analytical framework. In section 3, we present a subjective expected utility theory that captures the anticipated discovery of indescribable consequences. In section 4, we introduce additional axioms linking distinct levels of unawareness and a representation theorem that captures the evolution of a decision maker's beliefs in the wake of new discoveries. In section 5, we discuss a number of points, including small worlds, the evolution of beliefs, the behavioral manifestations of awareness of unawareness, the implications of applying our approach to defining "unknown unknowns" to the standard subjective expected utility models, and the related literature. The proofs are collected in the Appendix.

2. The analytical framework

In Karni and Vierø (2013, 2015), we modeled and analyzed the evolution of a decision maker's beliefs when her universe, formalized as a state space, expands in the wake of discoveries of new actions and/or consequences.[3] In this work, our investigation focuses on the effects of anticipating the discovery of unexpected consequences on a decision maker's choice behavior, and on the evolution of her beliefs and her sense of ignorance following such discoveries. In view of the differences in both the nature of the discoveries and the evolution of the state space, we leave the investigation of the anticipation of discovery of new feasible actions for future work.

The prospect of discovering consequences which the decision maker is unaware of and the sentiments, such as fear or excitement, that it evokes, presumably affects her choice behavior. Our first goal is to obtain a representation of preferences that assigns utility to unspecified consequences that may not even exist. The utility of unimaginable consequences represents the decision maker's emotions evoked by the prospect of their discovery.

2.1. Conceivable states and the objects of choice

Let A be a finite, nonempty, set of *basic actions* with generic element a, and C be a finite, nonempty, set of *feasible consequences* with generic element c. The elements of C are consequences that the decision maker is aware of. The key innovation compared to Karni and Vierø (2013, 2015) is that the decision maker may also entertain the idea that there might be consequences, of which she is unaware, that are unimaginable. We define $x = \neg C$ to be the abstract "consequence" that has the interpretation "none of the above." There may, in fact, be one consequence, any finite number, or an infinite number of consequences that the decision maker is unaware of, or no such consequence at all. The abstract consequence x captures all of these possibilities.[4] Ex ante, the decision maker cannot know which of these is true.[5] Let $\hat{C} = C \cup \{x\}$, which we will refer to as the set of *extended consequences*, with generic element $\hat{c}$. Together these sets determine the *augmented conceivable state space*, defined as

$$\hat{C}^A := \{s : A \to \hat{C}\}.$$

[3] The state space evolves differently depending on whether a new action or a new consequence is discovered. For details, see Karni and Vierø (2013).

[4] In section 4 we discuss the evolution of the decision maker's beliefs about the likelihood of the sets of partially-describable and nondescribable states. These beliefs are represented by a subjective probability distribution representing how likely the decision maker finds it that consequences she is currently unaware of will be discovered. One interpretation of this likelihood is that it reflects the decision maker's beliefs about the size of the set of consequences of which she is unaware.

[5] Machina (2003) mentions the possibility of capturing the anticipation of the unexpected by specifying a catch-all state, with a label like "none of the above." Unlike our approach, according to which "none of the above" refers to unspecified consequences, Machina applies the term to unspecified states.

That is, the augmented conceivable state space is the set of all functions from A to $\hat{C}$ and is, by definition, exhaustive.[6] The sets A and C also determine the subset of *fully describable conceivable states*, $C^A := \{s : A \to C\}$.

To illustrate, consider the following situation: there are two different medications, designed to treat the same health problem, that must be taken regularly. For simplicity suppose that each medication can lead to one of two known outcomes, success and failure (e.g., reducing the cholesterol level below a target threshold). Suppose that one of the medications has been used for some time, while the second medication was just approved but tests show it to be more effective. Each of the medications might have long-term, unforeseen, side effects that will not be known for some time. We can describe the situation as follows: Corresponding to the two medications there are two basic actions, $A = \{a_1, a_2\}$, and corresponding to the two known possible outcomes there are two feasible consequences, $C = \{c_1, c_2\}$. The unknown possible side-effects are denoted by x. The resulting augmented conceivable state space consists of nine states as depicted in the following matrix:

$$\begin{array}{cccccccccc} A \setminus S & s_1 & s_2 & s_3 & s_4 & s_5 & s_6 & s_7 & s_8 & s_9 \\ a_1 & c_1 & c_2 & c_1 & c_2 & x & x & c_1 & c_2 & x \\ a_2 & c_1 & c_1 & c_2 & c_2 & c_1 & c_2 & x & x & x \end{array} \tag{1}$$

The subset of fully describable conceivable states in this example is $C^A = \{s_1, ..., s_4\}$.

The situation described in this example can be used to highlight the difference between the present model and the model of Karni and Vierø (2013). In our earlier work there is no analogue of x. Thus, the only consequences that the decision maker can conceive of are c_1 and c_2. Therefore, the state space that represents her conception of the world consists of the subset of fully describable states (in this example that is $C^A = \{s_1, ..., s_4\}$). In other words, despite her past experience, which includes discoveries of consequences of which she was unaware, the decision maker believes that her current conception of the universe is complete and fully describable. The model excludes the analysis of choice behavior of decision makers who, having repeatedly learned that their conception of the possible consequences of their actions was incomplete, anticipate future discoveries of inconceivable consequences. In section 3.1 below we pursue this comparison and show that the two models might yield opposite predictions.

Define the set of *conceivable acts*, F, to be the set of all the mappings from the augmented conceivable state space to the set of feasible consequences.[7] Formally,

$$F := \{f : \hat{C}^A \to C\}. \tag{2}$$

Because conceivable acts are functions whose domain is the state space, adding them to the list of acts does not require further expansion of the state space. In other words, once the state is known, all uncertainty regarding the outcome of a conceivable act is resolved and no new states are created. By contrast, if a new *basic action* is either designed or discovered then, by definition,

[6] This method of constructing the state space from the primitive sets of feasible acts and consequences appears in Schmeidler and Wakker (1987) and Karni and Schmeidler (1991). It was used in Karni and Vierø (2013, 2015). The idea was also discussed in the philosophical literature (see Stalnaker, 1972; Gibbard and Harper, 1978). The augmentation due to "none of the above" is specific to the present paper.

[7] Note that the definition of conceivable acts in the present paper differs from the definition of conceivable acts in Karni and Vierø (2013, 2015). In our previous work, conceivable acts were functions from conceivable states to lotteries over consequences (i.e., Anscombe–Aumann, 1963, acts). For the reasons discussed in the introduction, the approach taken here is more satisfactory.

it assigns all the consequences to each state in $\hat{C}^A$. Thus, each state in $\hat{C}^A$ becomes an event in the newly defined state space. Consider the example in which there are two basic actions and two feasible consequences. If a new basic action is discovered, the state (c_i, c_j) becomes the event $\{(c_i, c_j, c_1), (c_i, c_j, c_2), (c_i, c_j, x)\}$, $i, j = 1, 2$. That is, because the payoff of the new basic action in this prior state can be c_1, c_2, or x, a new basic action means that the prior state (c_i, c_j) no longer resolves the uncertainty. By contrast, the states (c_i, c_j), $i, j = 1, 2$, completely resolves the uncertainty regarding the payoff of the new conceivable acts.[8]

By definition, the payoffs of the conceivable acts are restricted to *feasible* consequences (that is, their range does not include x). Because we adhere to the revealed preference methodology, we require that, for a given level of awareness, acts must be meaningfully described and their consequences effectuated once the uncertainty is resolved. Thus, if one must be able to effectuate the consequences specified by conceivable acts once uncertainty is resolved, then the specification in (2) is the most general possible. Including the abstract consequence "none of the above," or x, in the range of the conceivable acts would create a conceptual problem in fully describable states (e.g., the states $s_1, ..., s_4$ in the example in matrix (1)). In these states, x remains abstract, so a conceivable act that pays off x cannot be settled in those states and is, therefore, meaningless. While the decision maker could potentially describe such acts (as we just did), it is too farfetched to suppose that she could express preferences over them. Since the range of the basic actions includes x, the set F of conceivable acts does not include the basic actions.

The argument in the preceding paragraph only applies if we restrict the set of acts to maps whose range is the same set of consequences in all states. Without this restriction, we can expand the set of acts that preferences can meaningfully be expressed over. In particular, in states whose partial or complete descriptions include x, this abstract consequence takes a concrete meaning ex post, and conceivable acts that pay off x in one or more of these states can be settled. In the above example, with the state space depicted in (1), an act that assigns a consequence, which is neither c_1 nor c_2 and which will be discovered in the event $\{s_5, ..., s_9\}$, to one or more of the states $s_5, \ldots, s_9$, is well-defined. In other words, the decision maker can promise to deliver a newly discovered consequence, whatever it may be, if such a consequence is discovered, and she will be able to keep her promise if such a discovery is made.

To explore the possibility of assigning utility to unknown consequences, x, we extend the set of acts by adding functions whose range includes x as a possible payoff in the imperfectly describable states $\hat{C}^A \setminus C^A$. Formally, we define the set of *extended conceivable acts* F^* as follows:

$$F^* := \{f^* : \hat{C}^A \to \hat{C} \mid f^{*-1}(x) \subseteq \hat{C}^A \setminus C^A\}. \tag{3}$$

By definition, the consequences of extended conceivable acts are restricted to elements of C in the fully describable states, but can be any element of $\hat{C}$, including x, in the imperfectly describable states. A schematic illustration in the context of the example in matrix (1) is given in Fig. 1. The range of the extended conceivable acts is C in $\{s_1, \ldots, s_4\}$, and $\hat{C}$ in $\{s_5, \ldots, s_9\}$.

Note that the set, A, of basic actions, and the set, F, of conceivable acts are (disjoint) subsets of the set of extended conceivable acts. Each basic action $a \in A$ is identified with the extended conceivable act $f^* \in F^*$ for which $f^*(s) = s(a)$, for all $s \in \hat{C}^A$. Note also that F^* does not include, among others, the constant act whose payoff is x. Given the decision maker's awareness, the set of extended conceivable acts F^* is the most that can be both meaningfully expressed and settled ex post.

[8] For more detailed discussion of the implications of discovering new basic actions, see Karni and Vierø (2013).

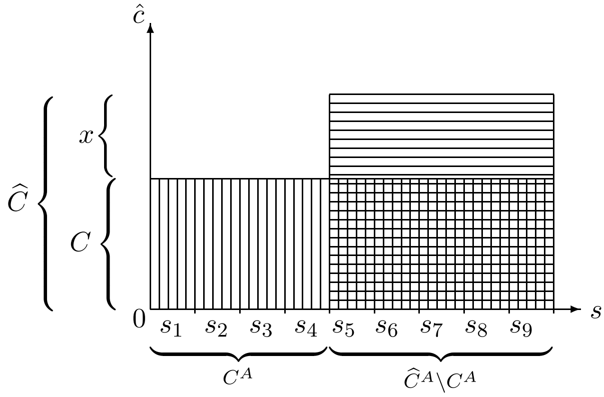

Fig. 1. Illustration of the set of extended conceivable acts.

We consider lotteries over extended conceivable acts. Formally, denote by $\Delta(F^*)$ the set of all probability distributions on F^*, and by $\Delta(F)$ its subset of all probability distributions on F. A generic element $\mu \in \Delta(F^*)$ selects an extended conceivable act in F^* according to the distribution μ. We refer to the elements of $\Delta(F^*)$ by the name *mixed extended conceivable acts*. The set $\Delta(F^*)$ of all such lotteries is the choice set. Decision makers are supposed to be able to form and express preferences over $\Delta(F^*)$.[9] It turns out that having the decision maker express preferences over the set of mixed extended conceivable acts is, in fact, sufficient to obtain a representation with a utility of unknown consequences. However, this representation requires a non-standard approach because the domain of preferences is "non-rectangular".

We abuse notation and denote by c also the constant act that assigns c to every state in $\hat{C}^A$, and by f the degenerate mixed extended conceivable act that assigns the unit probability mass to the conceivable act f. For all $\mu, \mu' \in \Delta(F^*)$ and $\alpha \in [0, 1]$, let $\alpha\mu + (1-\alpha)\mu' \in \Delta(F^*)$ be defined as pointwise mixtures on the support of the mixed conceivable acts (that is, $\left(\alpha\mu + (1-\alpha)\mu'\right)(f) = \alpha\mu(f) + (1-\alpha)\mu'(f)$, for all $f \in F^*$). Then $\Delta(F^*)$ is a convex set. Finally, for any $f, g \in F^*$ and $E \subset \hat{C}^A$, let $g_E f$ denote the act in F^* defined by $(g_E f)(s) = g(s)$, if $s \in E$, and $(g_E f)(s) = f(s)$ otherwise.

3. Subjective expected utility with unknown consequences

Consider next a decision maker whose choices are characterized by a (strict) preference relation $\succ$ on $\Delta(F^*)$. We assume that $\succ$ satisfies the well-known axioms of expected utility theory.

(A.1) (Preorder) The preference relation $\succ$ on $\Delta(F^*)$ is asymmetric and negatively transitive.[10]

(A.2) (Archimedean) For all $\mu, \mu', \mu'' \in \Delta(F^*)$, if $\mu \succ \mu'$ and $\mu' \succ \mu''$ then there are $\alpha, \beta \in (0, 1)$ such that $\alpha\mu + (1-\alpha)\mu'' \succ \mu'$ and $\mu' \succ \beta\mu + (1-\beta)\mu''$.

(A.3) (Independence) For all $\mu, \mu', \mu'' \in \Delta(F^*)$ and $\alpha \in (0, 1]$, $\mu \succ \mu'$ if and only if $\alpha\mu + (1-\alpha)\mu'' \succ \alpha\mu' + (1-\alpha)\mu''$.

[9] We suppose implicitly that decision makers are able to use devices to randomize their choices. Evidence suggesting that decision makers deliberately randomize their choice is provided in Agranov and Ortolova (2016).

[10] This implies that $\succ$ is irreflexive and transitive (see Kreps, 1988, proposition 2.3).

Define the weak preference relation, $\succcurlyeq$, to be the negation of the strict preference relation, (i.e., $\succcurlyeq = \neg(\prec)$), and the indifference relation, $\sim$, to be the symmetric part of $\succcurlyeq$. Then, $\succcurlyeq$ is a weak order (i.e., complete and transitive) satisfying the corresponding version of independence.

Because the choice set is the set of mixed extended conceivable acts, which is less structured than the set of Anscombe–Aumann (1963) acts, we need additional structure from the axioms to obtain an expected utility representation.[11] For this purpose, we consider the mapping $\varphi : \Delta(F^*) \to (\Delta(\hat{C}))^{\hat{C}^A}$, where, for all $s \in \hat{C}^A$, $\hat{c} \in \hat{C}$ and $\mu \in \Delta(F^*)$,

$$\varphi_s(\mu)(\hat{c}) := \sum_{\{f \in Supp(\mu) \mid f(s)=\hat{c}\}} \mu(f). \tag{4}$$

The mapping φ transforms each mixed extended conceivable act into an Anscombe–Aumann act. More specifically, for each $s \in \hat{C}^A$, the vector $\varphi_s(\mu) \in \Delta(\hat{C})$ is the lottery that $\varphi(\mu)$ assigns to the state s. It is important to note that the support of the lotteries in the resulting Anscombe–Aumann acts is a subset of C in the fully describable states C^A, and a subset of $\hat{C}$ in the imperfectly describable states $\hat{C}^A \setminus C^A$. Thus, the set of Anscombe–Aumann acts defined by the mapping in (4) inherits the non-rectangular shape of the set of extended conceivable acts.

Henceforth, for $\{\mu : Supp(\mu) \subseteq F\}$ we also denote by $\varphi_s(\mu)$ the mixed conceivable act that assigns the probability $\varphi_s(\mu)(c)$ to the constant conceivable act c. Under this convention, the set $\Delta(C)$ also denotes the subset of mixed conceivable acts whose supports are restricted to the constant conceivable acts (that is, $\Delta(C) \subset \Delta(F)$).

Whereas the mapping φ yields a unique Anscombe–Aumann act for each $\mu \in \Delta(F^*)$, in general, every Anscombe–Aumann act in the set derived from $\Delta(F^*)$ using the mapping φ corresponds to multiple mixed extended conceivable acts. Hence the need for an extra axiom.[12] The next axiom asserts that the decision maker is indifferent among mixed extended conceivable acts whose images under φ are the same (that is, the decision maker is indifferent between mixed extended conceivable acts that are transformed to the same Anscombe–Aumann act).

(A.4) (Extended Indifference) For all $\mu, \mu' \in \Delta(F^*)$, if $\varphi(\mu) = \varphi(\mu')$ then $\mu \sim \mu'$.

The next Lemma shows that preference relations restricted to $\Delta(F)$ satisfying (A.1)–(A.4) have expected utility (over conceivable acts) and additively separable (across states) representations. To state the Lemma, we invoke the following definition: A set of real-valued functions $\{W_s\}_{s \in \hat{C}^A}$ on C, representing a preference relation $\succ$ on $\Delta(F)$, is *unique up to cardinal unit-comparable transformation* if the set $\{\hat{W}_s\}_{s \in \hat{C}^A}$ on C also represents the same preference relation if and only if $\hat{W}_s = bW_s + d_s$, $b > 0$.

Lemma 1. *A preference relation $\succ$ on $\Delta(F)$ satisfies (A.1)–(A.4) if and only if there exist real-valued functions $\{W_s\}_{s \in \hat{C}^A}$ on C, unique up to cardinal unit-comparable transformation, such that, for all $\mu, \mu' \in \Delta(F)$,*

$$\mu \succ \mu' \Leftrightarrow \sum_{f \in F} \mu(f) \sum_{s \in \hat{C}^A} W_s(f(s)) > \sum_{f \in F} \mu'(f) \sum_{s \in \hat{C}^A} W_s(f(s)). \tag{5}$$

[11] Here we follow a procedure mentioned in Kreps (1988), Chapter 7. The next axiom is suggested there.

[12] Notice that everything could be done directly with this non-rectangular set of extended Anscombe–Aumann acts, but for the reasons discussed in the introduction, we find the starting point of mixed extended conceivable acts more satisfactory. If one were to start from the extended Anscombe–Aumann acts, one would not need axiom (A.4).

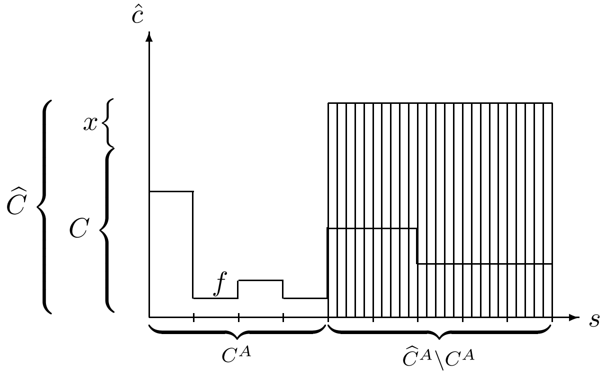

Fig. 2. Illustration of the set of conditional extended conceivable acts for a particular f.

Following Savage (1954), a state $s \in \hat{C}^A$ is said to be *null* if $\hat{c}_{\{s\}} f \sim \hat{c}'_{\{s\}} f$, for all $\hat{c}, \hat{c}' \in \hat{C}$, for all $f \in F^*$. A state is said to be *nonnull* if it is not null.

To state the next axiom we use the following notation: Let $\tilde{F} := \{\tilde{f} : \hat{C}^A \setminus C^A \to \hat{C}\}$ (that is, $\tilde{F}$ is the set of all functions from the set of imperfectly describable states to the set of *extended* consequences). Define sets of *conditional extended conceivable acts* as follows: For every $f \in F$, let

$$F_{\hat{C}^A \setminus C^A}(f) := \{\tilde{f}_{\hat{C}^A \setminus C^A} f \in F^* \mid \tilde{f} \in \tilde{F}\}$$

(that is, $F_{\hat{C}^A \setminus C^A}(f)$ is the set of all acts in F^* that are extensions of $f \in F$). A schematic illustration is given in Fig. 2. The range of f is C and the acts in $F_{\hat{C}^A \setminus C^A}(f)$ all agree with f on C^A and return any consequence in $\hat{C}$ in the states in $\hat{C}^A \setminus C^A$. Note that $\cup_{f \in F} F_{\hat{C}^A \setminus C^A}(f) = F^*$.

We denote by $\Delta\big(F_{\hat{C}^A \setminus C^A}(f)\big)$ the corresponding set of mixed conditional extended conceivable acts. For each $f \in F$ and $\hat{p} \in \Delta(\hat{C})$, let $\hat{p}_{\hat{C}^A \setminus C^A} f$ denote the distribution in $\Delta(F^*)$ that, for all $\hat{c} \in \hat{C}$, assigns the probability $\hat{p}(\hat{c})$ to the extended conceivable act $\hat{c}_{\hat{C}^A \setminus C^A} f$.

Given $f \in F^*$, let $F^*(f, s) := \{c_{\{s\}} f \in F^* \mid c \in C\}$ if $s \in C^A$ and $\tilde{F}(f, s) := \{\hat{c}_{\{s\}} f \in F^* \mid \hat{c} \in \hat{C}\}$ if $s \in \hat{C}^A \setminus C^A$. Denote by $\Delta(F^*(f, s))$ and $\Delta(\tilde{F}(f, s))$ the subsets of mixed extended conceivable acts whose supports are $F^*(f, s)$ and $\tilde{F}(f, s)$, respectively.

(A.5) (Monotonicity) For all $f \in F^*$, (a) For all nonnull $s \in C^A$, $\mu, \mu' \in \Delta(F^*(f, s))$, and $\varphi_s(\mu), \varphi_s(\mu') \in \Delta(C) \subset \Delta(F^*)$, it holds that $\mu \succ \mu'$ if and only if $\varphi_s(\mu) \succ \varphi_s(\mu')$. (b) For all $\mu, \mu' \in \Delta(\tilde{F}(f, s))$ and nonnull $s \in \hat{C}^A \setminus C^A$, it holds that $\mu \succ \mu'$ if and only if $\varphi_s(\mu)_{\hat{C}^A \setminus C^A} f \succ \varphi_s(\mu')_{\hat{C}^A \setminus C^A} f$.

In the monotonicity axiom the mixed conceivable acts, μ and μ', have as their supports conceivable acts whose payoffs differ in a single state, s. The Anscombe–Aumann acts, induced by μ and μ', agree in every state except s, in which they yield $\varphi_s(\mu)$ and $\varphi_s(\mu')$, respectively. The axiom states that the direction of preference between μ and μ' is the same as the direction of preference between the mixed conceivable acts that have distributions $\varphi_s(\mu)$ and $\varphi_s(\mu')$ over the constant conceivable acts. Similarly, the mixed extended conceivable acts $\mu, \mu' \in \Delta(\tilde{F}(f, s))$, have the same conditional distributions $\varphi_s(\mu)_{\hat{C}^A \setminus C^A} f$ and $\varphi_s(\mu')_{\hat{C}^A \setminus C^A} f$ over the conditional (on $\hat{C}^A \setminus C^A$) constant conceivable acts.[13] Thus, our monotonicity axiom is of the same spirit and plays the same role as the monotonicity axiom in the Anscombe–Aumann model. However, its expression is different because the decision maker's choice set consists of mixed extended conceivable acts.[14]

The next axiom requires that the decision maker is not indifferent among all mixed extended conceivable acts.

(A.6) (Nontriviality) The strict preference relation $\succ$ on $\Delta(F^*)$ is nonempty.

Note that (A.6) implies the existence of consequences, $c^*, c_* \in \hat{C}$, such that $c^* \succ c_*$.

Proposition 1. *Let $\succ$ be a preference relation on $\Delta(F^*)$, then the following two conditions are equivalent:*

(i) The preference relation $\succ$ satisfies (A.1)–(A.6).

(ii.a) There exist a real-valued, continuous, nonconstant, affine, function, U on $\Delta(C)$, and a probability measure, π on $\hat{C}^A$, such that, for all $\mu, \lambda \in \Delta(F)$,

$$\mu \succ \lambda \Leftrightarrow \sum_{s\in\hat{C}^A} \pi(s)U(\varphi_s(\mu)) > \sum_{s\in\hat{C}^A} \pi(s)U(\varphi_s(\lambda)). \tag{6}$$

(ii.b) For every $f \in F$, there exist a real-valued, non-constant, affine, function, U_f^* on $\Delta(\hat{C})$, and a probability measure, ϕ on $\hat{C}^A \setminus C^A$, such that, for all μ and λ in $\Delta(F_{\hat{C}^A\setminus C^A}(f))$,

$$\mu \succ \lambda \Leftrightarrow \sum_{s\in\hat{C}^A\setminus C^A} \phi(s)U_f^*(\varphi_s(\mu)) > \sum_{s\in\hat{C}^A\setminus C^A} \phi(s)U_f^*(\varphi_s(\lambda)). \tag{7}$$

Moreover, each of the functions U and U_f^* is unique up to positive linear transformations, the probability measures, π and ϕ, are unique, and $\pi(s) = \phi(s) = 0$ if and only if s is null.

The proof is in the appendix.

By the affinity of U, $U(\varphi_s(\mu)) = \Sigma_{c\in Supp(\varphi_s(\mu))}\varphi_s(\mu)(c)u(c)$, where u is a real-valued function on C. Similarly, for each $f \in F$, $U_f^*(\varphi_s(\mu) = \Sigma_{\hat{c}\in Supp(\varphi_s(\mu))}\varphi_s(\mu)(\hat{c})u_f(\hat{c})$, where u_f is a real-valued function on $\hat{C}$.

Since the sets $\Delta(F)$ and $\Delta\big(F_{\hat{C}^A\setminus C^A}(f)\big)$ intersect (see Figs. 1 and 2), the representations in (6) and (7) together imply that $U_f^*(p) = U(p)$, for all $f \in F$ and $p \in \Delta(C)$, and that $\phi(s) = \pi(s)/\pi(\hat{C}^A \setminus C^A)$, for all $s \in \hat{C}^A \setminus C^A$. However, the utility of the abstract consequence x, $U_f^*(x)$, may depend on the act f. The next axiom, separability, links the conditional representations in Proposition 1. The axiom requires that the ranking of mixed conditional extended conceivable acts whose supports are constant on the set of partially describable states, $\hat{C}^A \setminus C^A$, be independent of the conditioning act. This separability is not implied by the independence axiom because the payoff x is not defined on the subset of fully describable states.

[13] In part (b) of the axiom, we abuse notation slightly by letting $\varphi_s(\mu)_{\hat{C}^A\setminus C^A}f$ and $\varphi_s(\mu')_{\hat{C}^A\setminus C^A}f$ denote mixed extended conceivable acts whose support include extended conceivable acts that are constant at x on $\hat{C}^A \setminus C^A$.

[14] It is not straightforward to extend Axiom (A.5) to F^*, because F^* does not contain the constant act x.

(A.7) (Separability) For all $f, g \in F^*$ and $\hat{p}, \hat{q} \in \Delta(\hat{C})$, $\hat{q}_{\hat{C}^A \setminus C^A} f \succ \hat{p}_{\hat{C}^A \setminus C^A} f$ if and only if $\hat{q}_{\hat{C}^A \setminus C^A} g \succ \hat{p}_{\hat{C}^A \setminus C^A} g$.

In the next theorem, we use the separability axiom to combine the representations in (6) and (7). This allows us to obtain a general subjective expected utility representation that includes an assignment of utility to the abstract consequence x.

Theorem 1. *Let $\succ$ be a preference relation on $\Delta(F^*)$, then the following conditions are equivalent:*

(i) The preference relation satisfies axioms (A.1)–(A.7).
(ii) There exist real-valued, non-constant, affine, functions, U on $\Delta(C)$ and U^* on $\Delta(\hat{C})$, and a probability measure, π on $\hat{C}^A$, such that, for all $\mu, \lambda \in \Delta(F^*)$, $\mu \succ \lambda$ if and only if

$$\sum_{s \in C^A} \pi(s) U(\varphi_s(\mu)) + \sum_{s \in \hat{C}^A \setminus C^A} \pi(s) U^*(\varphi_s(\mu)) > \sum_{s \in C^A} \pi(s) U(\varphi_s(\lambda)) + \sum_{s \in \hat{C}^A \setminus C^A} \pi(s) U^*(\varphi_s(\lambda)). \tag{8}$$

Moreover, the functions U and U^* are unique up to positive linear transformations and they agree on $\Delta(C)$.[15] Also, the probability measure is unique, with $\pi(s) = 0$ if and only if s is null.

The proof is in the appendix. By the affinity of U^*, $U^*(\varphi_s(\mu)) = \Sigma_{\hat{c} \in Supp(\varphi_s(\mu))} \varphi_s(\mu)(\hat{c}) \times u^*(\hat{c})$, where u^* is a real-valued function on $\hat{C}$.[16]

As Theorem 1 shows, enriching the framework to include extended conceivable acts has allowed us to obtain an expected utility representation that assigns utility to unknown consequences. The representation consists of three elements: beliefs over states, a Bernoulli utility function over known consequences, and a parameter, $u^*(x)$, that captures the decision maker's utility of, or attitude toward, the unknown. This utility reflects whether the decision maker faces the unknown with fear, excitement, or indifference. Comparing two decision makers, the one with the higher value of $u^*(x)$ exhibits excitement toward the unknown relative to the decision maker with the lower value, who is more fearful toward the unknown. If the set of imperfectly describable states is null, the decision maker is a standard subjective expected utility maximizer.

3.1. An example

One advantage of our framework is that it distinguishes between states in which different basic actions result in new consequences, as illustrated in the matrix (1) in Section 2.1. Therefore this framework accommodates viewing different actions as being more or less likely to increase awareness. If unforeseeable consequences generate excitement, actions that are perceived as more likely to result in such consequences are expected to be preferred over similar

15 A different way of saying that U and U^* agree on $\Delta(C)$ is to say that U^* is an extension of U.

16 Recall that we use $\varphi_s(\mu)$ to also denote the mixed conceivable act that assigns the probability $\varphi_s(\mu)(c)$ to the constant conceivable act c. For expository convenience, we use this shorthand notation in the representation. To emphasize, the representation indeed evaluates mixed extended conceivable acts.

actions that are less likely to result in unforeseeable consequences. Consider, for example, the matrix (1). Suppose that the decision maker is confident that the action a_1 is unlikely to reveal an unforeseen consequence. Specifically, the medication a_1 has been in use for some time and has shown no side-effects. Suppose that the decision maker believes that if she chooses a_1 either the consequence c_1, success, or c_2 failure, will obtain. In other words, on the basis of past experience, the decision maker believes that if a_1 is implemented it is impossible that "neither c_1 nor c_2" (that is, x) will obtain. Formally, she considers the event $\{s_5, s_6, s_9\}$ to be null. By contrast, she considers x to be a real possibility if the new medication, a_2, is chosen. Thus, the event $\{s_7, s_8\}$ is assigned positive probability. By the representation (8),

$$a_1 \mapsto U(c_1)\left[\pi(s_1) + \pi(s_3) + \pi(s_7)\right] + U(c_2)\left[\pi(s_2) + \pi(s_4) + \pi(s_8)\right],$$

and

$$a_2 \mapsto U(c_1)\left[\pi(s_1) + \pi(s_2)\right] + U(c_2)\left[\pi(s_3) + \pi(s_4)\right] + U^*(x)\left[\pi(s_7) + \pi(s_8)\right].$$

Therefore, a choice of a_2 over a_1 yields a higher probability of encountering unforeseeable side-effects, or consequence, x. If $U(c_1) < U^*(x)$ and $\pi(s_3) \leq \pi(s_2) + \pi(s_8)$, then $a_2 \succ a_1$.[17]

To grasp the difference in the analysis if the same issue is addressed in the model of Karni and Vierø (2013), recall that in that model the state space consists solely of the fully describable states. Consequently, that model is silent on the distinction between the null event $\{s_5, s_6, s_9\}$ and the event $\{s_7, s_8\}$. Neither of these events, nor their union, can be expressed in that model. Hence, according to Karni and Vierø (2013), the utility associated with the two medications are:

$$a_1 \mapsto U(c_1)\left[\hat{\pi}(s_1) + \hat{\pi}(s_3)\right] + U(c_2)\left[\hat{\pi}(s_2) + \hat{\pi}(s_4)\right],$$

and

$$a_2 \mapsto U(c_1)\left[\hat{\pi}(s_1) + \hat{\pi}(s_2)\right] + U(c_2)\left[\hat{\pi}(s_3) + \hat{\pi}(s_4)\right],$$

where $\hat{\pi}$ is the subjective probability measure that figured in that work. It is possible, therefore, that our (2013) 'reverse Bayesianism' model would predict that a_1 is chosen over a_2 (that is, if $\hat{\pi}(s_3) > \hat{\pi}(s_2)$) while the present model predicts the opposite choice behavior. Furthermore, if $\pi(s_8) > 0$ then the opposite predictions may arise even when the two models are consistent in the sense of having the same likelihood ratios of s_3 and s_2 (that is, $\pi(s_2)/\pi(s_3) = \hat{\pi}(s_2)/\hat{\pi}(s_3)$).

4. Growing awareness and the evolution of beliefs

Thus far our attention was restricted to the axiomatic structure and representation of preference relations for a given level of awareness. We now turn to the study of the decision maker's growing awareness and the evolution of her beliefs in response to such expansions. The decision maker's awareness expands when she discovers a new consequence that was hidden behind a "veil of ignorance," that we referred to as "none of the above". As the analysis that follows makes clear, the characterization of the evolution of a decision maker's beliefs in the wake of her growing awareness does not require assigning utility to the abstract consequence "none of the above."

Henceforth, we use the subscript 0 to index the various sets under the prior level of awareness and the subscript 1 to index the various sets under the posterior level of awareness. Thus, C_0, x_0,

---

[17] Note that $U(c_1) > U(c_2)$. If $U(c_1) = U^*(x)$, then $a_2 \succcurlyeq a_1$ reduces to $[U(c_1) - U(c_2)][\pi(s_3) - \pi(s_2) - \pi(s_8)] \leq 0$, which is satisfied given the assumption about the probabilities. If $U(c_1) < U^*(x)$, we have that $a_2 \succ a_1$.

$\hat{C}_0$, C_0^A, $\hat{C}_0^A$, F_0, F_0^*, etc. refer to the respective sets under the prior level of awareness, with the analogous notation for the posterior sets.

As awareness grows, the state space evolves as follows. If a new consequence, $c' \notin C_0$, is discovered, the set of feasible consequences expands to $C_1 = C_0 \cup \{c'\}$. At the same time, the abstract consequence that has the interpretation "none of the above" becomes $x_1 = \neg C_1$, and the extended set of consequences becomes $\hat{C}_1 = C_1 \cup \{x_1\}$. The posterior conceivable state space is $\hat{C}_1^A$. In our illustrating example, if a new consequence c_3 is discovered, the augmented conceivable state space becomes

$$\begin{array}{ccccccccccccccccc} A/S & s_1 & s_2 & s_3 & s_4 & s_5' & s_5 & s_6' & s_6 & s_7' & s_7 & s_8' & s_8 & s_9' & s_9'' & s_9''' & s_9 \\ a_1 & c_1 & c_2 & c_1 & c_2 & c_3 & x_1 & c_3 & x_1 & c_1 & c_1 & c_2 & c_2 & c_3 & c_3 & x_1 & x_1 \\ a_2 & c_1 & c_1 & c_2 & c_2 & c_1 & c_1 & c_2 & c_2 & c_3 & x_1 & c_3 & x_1 & c_3 & x_1 & c_3 & x_1 \end{array} \tag{9}$$

The set of fully describable states also expands and is now $C_1^A = C_0^A \cup \{s_5', s_6', s_7', s_8', s_9'\}$. Thus, when a new feasible consequence is discovered, each of the prior fully describable states remains as before, while each of the prior imperfectly describable states is split into a fully describable state and one, or more, posterior imperfectly describable states. Hence, elements are added to the subset of fully describable states and, simultaneously, the number of imperfectly describable states increases.[18] As the decision maker's augmented conceivable state space expands, so does the set of conceivable acts, to $F_1 := \{f : \hat{C}_1^A \to C_1\}$, and the set of extended conceivable acts to $F_1^* := \{f^* : \hat{C}_1^A \to \hat{C}_1 \mid f^{*-1}(x_1) \subseteq \hat{C}_1^A \setminus C_1^A\}$. The corresponding set of mixed conceivable acts is $\Delta(F_1)$ and the set of mixed extended conceivable acts is $\Delta(F_1^*)$.

Because the set of conceivable acts is variable in our model, the preference relation must be redefined on the extended domain. Therefore, a decision maker is characterized by a collection of preference relations, one for each level of awareness over the corresponding set of mixed extended conceivable acts. We denote the strict preference relation on $\Delta(F_i^*)$ by $\succ_i$, $i = 0, 1$. In particular, the prior preference relation is denoted by $\succ_0$ on $\Delta(F_0^*)$ and the posterior preference relation by $\succ_1$ on $\Delta(F_1^*)$. We denote by φ^i the mapping given by (4) for awareness level i.

4.1. Reverse Bayesian updating

To link the preference relations across expanding sets of mixed extended conceivable acts, we invoke the relevant part of the invariant risk preferences axiom introduced in Karni and Vierø (2013), asserting the commonality of risk attitudes across levels of awareness.[19] Recall

[18] At first glance, the introduction of the abstract consequence x may seem to make the discovery of new consequences similar to the discovery of new actions as the two types discoveries can be expressed in terms of refinement of the original state-space. However, the two refinements are different. Unlike the discovery of new actions which refines the state space by associating to every state in the prior state space a set of states, one for each consequence, (see Karni and Vierø, 2013), the refinement of the prior space induced by the discovery of consequences of which the decision maker was unaware, is confined to the originally partially and non-describable states and takes the specific form as illustrated in (9). A formal definition of the refinement appears in the discussion preceding the introduction of axiom (A.10) below.

[19] The axiom appears in almost all works on unawareness and unforeseen contingencies. In particular, it is implicit in Maskin and Tirole (1999), Halpern and Rego (2014), Grant and Quiggin (2013), Heifetz et al. (2013a, 2013b), Auster (2013), and Schipper and Woo (2015). Ma and Schipper (2016) test the invariant risk preferences axiom experimentally and cannot reject it. Mengel, Tsakas, and Vostroknutov (2016) find that risk preferences are affected by exposure to an environment with imperfect knowledge of the state space. Ma and Schipper observe that Mengel et al.'s experimental design is likely to leave subjects suspecting that they are still unaware of some states after the initial exposure. Hence their result might capture changes in the subjects' attitude toward the unknown rather than in their risk preferences.

that $\Delta(C_0) \subset \Delta(F_0)$ also denotes the subset of mixed conceivable acts whose supports are the constant conceivable acts in F_0, and note that, for $C_1 \supset C_0$, we also have that $\Delta(C_0) \subset \Delta(F_1)$.

(A.8) (Invariant risk preferences) For $\succ_i$ on $\Delta(F_i)$, $i = 0, 1$, and for all $\mu, \mu' \in \Delta(C_0)$, it holds that $\mu \succ_0 \mu'$ if and only if $\mu \succ_1 \mu'$.

Assuming that $\succ_i$ on $\Delta(F_i)$, $i = 0, 1$, are non-trivial, there are $c_i^*, c_*^i \in C_i$ such that $c_i^* \succ_i c_*^i$, $i = 0, 1$. One implication of the invariant risk preferences axiom is that we may choose $c_0^* = c_1^* = c^*$ and $c_*^0 = c_*^1 = c_*$. Hence, for this particular purpose, we can simply write $c^* \succ_i c_*$, $i = 0, 1$.

The following two axioms depict additional links between the preference relations across different levels of awareness. The first axiom, dubbed Awareness Consistency I, asserts that the discovery of new consequences does not alter the decision maker's preferences conditional on the events that consist of a-priori fully describable states. Thus, such discoveries do not affect the part of her preferences that only concerns the initially fully describable and well-understood part of her universe. Formally, for every $h \in F_i$ and $E \subset \hat{C}_i^A$, let $F_i(h; E) := \{f_E h \in F_i \mid f \in F_i\}$, $i = 0, 1$. For all $E \subseteq C_0^A$, $\lambda \in \Delta(F_0(h; E))$, and $\lambda' \in \Delta(F_1(h'; E))$ define $\lambda = \lambda'$ on E if $\varphi_s^0(\lambda) = \varphi_s^1(\lambda')$, for all $s \in E$. That is, $\lambda = \lambda'$ on E if the mapping φ generates the same lottery for λ and λ' in all states in E.

(A.9) Awareness consistency I: For all $E \subseteq C_0^A$, for all $h \in F_0$, $\lambda, \mu \in \Delta(F_0(h; E))$, $h' \in F_1$, and $\lambda', \mu' \in \Delta(F_1(h'; E))$, such that, on E, $\lambda = \lambda'$ and $\mu = \mu'$, it holds that, $\lambda \sim_0 \mu$ if and only if $\lambda' \sim_1 \mu'$.

The second axiom, dubbed Awareness Consistency II, asserts that the discovery of new consequences does not alter the decision maker's preferences conditional on the events that consist of, a-priori, not fully describable states. In other words, the decision maker's ranking of a-priori measurable acts is independent of the detail with which she can describe the a-priori measurable sub-events. To state the axiom, we need additional notation and definitions. If $C_0 \subset C_1$ then, for each $s \in \hat{C}_0^A \setminus C_0^A$, there corresponds an event $E'(s) \subset \hat{C}_1^A \setminus C_0^A$, defined by

$$\begin{aligned} E'(s) = \{\hat{s} \in \hat{C}_1^A \setminus C_0^A \mid \forall a \in A,\ &\text{if } a(s) \in C_0, \text{ then } a(\hat{s}) = a(s), \\ &\text{and if } a(s) = x_0 \text{ then } a(\hat{s}) \in \{x_1\} \cup (C_1 \setminus C_0)\}. \end{aligned} \tag{10}$$

For each $E \subseteq \hat{C}_0^A \backslash C_0^A$ let $E'(E) := \cup_{s \in E} E'(s)$. A conceivable act $f' \in F_1$ is said to be *measurable with respect to* $\hat{C}_0^A$ if for all $c \in C_1$, $(f')^{-1}(c) = E'(E)$, for some $E \subseteq \hat{C}_0^A$. Let $F_1(\hat{C}_0^A)$ be the subset of the conceivable acts in F_1 that are measurable with respect to $\hat{C}_0^A$. There is a one-to-one correspondence between acts in F_0 and acts in $F_1(\hat{C}_0^A)$: For $f \in F_0$ and $f' \in F_1(\hat{C}_0^A)$, we write $f \simeq f'$ if $f'(s') = f(s)$, for all $s' \in E'(s)$ and $s \in \hat{C}_0^A$. For every $h' \in F_1(\hat{C}_0^A)$ and $E \subset \hat{C}_0^A$, let $F_1(\hat{C}_0^A; h'; E'(E)) := \{f_{E'(E)} h' \mid f \in F_1(\hat{C}_0^A)\}$ (that is, $F_1(\hat{C}_0^A; h'; E'(E))$ is the set of $\hat{C}_0^A$-measurable acts in F_1 that agree with h' outside of $E'(E)$). For $\lambda \in \Delta(F_0(h; E))$ and $\lambda' \in \Delta(F_1(\hat{C}_0^A; h'; E'(E)))$, we write $\lambda \simeq \lambda'$ if $\lambda(f) = \lambda'(f')$ when $f \simeq f'$.

(A.10) Awareness consistency II: For all $E \subseteq \hat{C}_0^A \backslash C_0^A$, $h \in F_0$, $\lambda, \mu \in \Delta(F_0(h; E))$, $h' \in F_1(\hat{C}_0^A)$ and $\lambda', \mu' \in \Delta(F_1(\hat{C}_0^A; h'; E'(E)))$, such that $\lambda \simeq \lambda'$ and $\mu \simeq \mu'$, it holds that $\lambda \sim_0 \mu$ if and only if $\lambda' \sim_1 \mu'$.

For preference relations satisfying the aforementioned axioms, Theorem 2 below asserts the existence and describes the uniqueness properties of a subjective expected utility representation for each level of awareness. More importantly, it describes the evolution of beliefs about the relative likelihoods of fully describable events and about the relative likelihoods of imperfectly describable events, in the wake of increasing awareness.

Theorem 2. *For each $C_0 \subset C_1$ and the corresponding preference relations $\succ_0$ on $\Delta(F_0)$ and $\succ_1$ on $\Delta(F_1)$, the following two conditions are equivalent:*

(i) The preference relations $\succ_0$ and $\succ_1$ satisfy (A.1)–(A.6) and, jointly, they satisfy (A.8)–(A.10).
(ii) There exist real-valued, continuous, nonconstant, affine, functions, U_0 on $\Delta(C_0)$ and U_1 on $\Delta(C_1)$, and probability measures, π_0 on $\hat{C}_0^A$ and π_1 on $\hat{C}_1^A$, such that for all $\mu, \lambda \in \Delta(F_0)$,

$$\mu \succ_0 \lambda \Leftrightarrow \sum_{s\in\hat{C}_0^A} \pi_0(s) U_0(\varphi_s^0(\mu)) > \sum_{s\in\hat{C}_0^A} \pi_0(s) U_0(\varphi_s^0(\lambda)), \tag{11}$$

and, for all $\mu', \lambda' \in \Delta(F_1)$,

$$\mu' \succ_1 \lambda' \Leftrightarrow \sum_{s\in\hat{C}_1^A} \pi_1(s) U_1(\varphi_s^1(\mu')) > \sum_{s\in\hat{C}_1^A} \pi_1(s) U_1(\varphi_s^1(\lambda')). \tag{12}$$

The functions U_0 and U_1 are unique up to positive linear transformations and for all $p \in \Delta(C_0)$, $U_0(p) = U_1(p)$. The probability measures π_0 and π_1 are unique and, for all $s, s' \in C_0^A$,

$$\frac{\pi_0(s)}{\pi_0(s')} = \frac{\pi_1(s)}{\pi_1(s')} \tag{13}$$

and, for all $s, s' \in \hat{C}_0^A \setminus C_0^A$,

$$\frac{\pi_0(s)}{\pi_0(s')} = \frac{\pi_1(E'(s))}{\pi_1(E'(s'))}. \tag{14}$$

By the affinity of U_i, $U_i(\varphi_s^i(\mu)) = \Sigma_{c\in Supp(\varphi_s^i(\mu))} \varphi_s^i(\mu)(c) u_i(c)$, where u_i is a real-valued function on C_i, for $i = 0, 1$. That $U_0(p) = U_1(p)$ for all $p \in \Delta(C_0)$ follows from axiom (A.8). Property (13) follows from axiom (A.9) and asserts that, in the wake of discoveries of new consequences, conditional on the initial set of fully describable states, the decision maker's subjective beliefs about the relative likelihoods of fully describable states remain unchanged. Property (14) follows from axiom (A.10) and asserts that the decision maker's subjective beliefs about the relative likelihood of events corresponding to a-priori partially describable states is the same as the relative likelihood of the corresponding states. Property (13) is reverse Bayesian updating following the discovery of a new consequence as in Karni and Vierø (2013). Thus, insofar as the discovery of new consequences is concerned, the model of Karni and Vierø (2013) is nested within the present one and corresponds to the special case when $\pi_i(C_i^A) = 1$ for all i. That is, in Karni and Vierø (2013), for any level of awareness the decision maker acts as if he assigns probability zero to future expansions of his awareness.

Remark. The main objective of Theorem 2 is the depiction of the evolution of the decision maker's beliefs. To attain this objective it is not necessary to consider the utility of the abstract consequences x_0 and x_1. Therefore, unlike in Theorem 1, in Theorem 2 the domains of the utility functions U_0 and U_1 are $\Delta(C_0)$ and $\Delta(C_1)$, respectively. It is a straightforward exercise to extend the representations in Theorem 2 to include utilities of the abstract consequences x_0 and x_1.

4.2. Decreasing and increasing sense of ignorance

A decision maker can respond to the discovery of a new consequence in one of three different ways: First, she could think that fewer consequences remain to be discovered. Second, the discovery of new consequences could reveal that the decision maker is more ignorant than she believed herself to be, and that more consequences than she suspected are waiting to be discovered. Third, she could consider that the current discovery has no effect on the prevalence of unknown consequences. Thus, the discovery of unforeseen consequences expands the decision maker's universe and, depending on their nature, may be accompanied by diminishing, growing, or unchanged sense of ignorance. These reactions have revealed preference manifestations that can be expressed axiomatically.

The next axiom captures the preferential expression of decreasing (increasing) sense of ignorance. In both cases, the axiom describes the decision maker's willingness to bet on, or against, discoveries of unforeseen consequences.

(A.11) (Decreasing (Increasing) Sense of Ignorance) For all $C_0 \subset C_1$, the corresponding sets of mixed conceivable acts $\Delta(F_0)$ and $\Delta(F_1)$, $\eta \in [0, 1]$, and $\lambda = \eta c_* + (1 - \eta)\, c^* \in \Delta(F_0)$, $\lambda' = \eta c_* + (1 - \eta)\, c^* \in \Delta(F_1)$, $\mu = c_{*C_0^A} c^* \in \Delta(F_0)$, and $\mu' = c_{*C_1^A} c^* \in \Delta(F_1)$, if $\lambda \sim_0 \mu$ then $\lambda' \succcurlyeq_1 (\preccurlyeq_1) \mu'$.

Note that this is a decreasing (increasing) sense of ignorance in the weak sense. It includes the cases of strictly decreasing (increasing) sense of ignorance, $\lambda' \succ_1 (\prec_1)\, \mu'$, and constant sense of ignorance, $\lambda' \sim_1 \mu'$, as special instances. The mixed conceivable acts λ and λ' only involve objective uncertainty, while μ and μ' are bets on discovering unforeseen consequences. A decision maker has a constant sense of ignorance if she is equally inclined to bet on something unforeseen arising (that is, on the realization of an imperfectly describable state) before and after the discovery of a new consequence. She has a strictly decreasing (increasing) sense of ignorance if she is less (more) inclined to bet on the realization of imperfectly describable states after the discovery.

Theorem 3 below quantifies decreasing (increasing) sense of ignorance by subjective probabilities. Specifically, if growing awareness is accompanied by decreasing (increasing) sense of ignorance, the subjective probability assigned to the 'residual' unawareness diminishes (grows).

Theorem 3. *For each pair $C_0 \subset C_1$ and the corresponding preference relations, $\succ_0$ on $\Delta(F_0)$ and $\succ_1$ on $\Delta(F_1)$, the following statements are equivalent:*

(i) The preference relations $\succ_0$ and $\succ_1$ satisfy (A.1)–(A.6) and, jointly, they satisfy (A.8)–(A.11).
(ii) There exists a representation as in Theorem 2 and, in addition,

$$\pi_0(\hat{C}_0^A \setminus C_0^A) \geq (\leq)\, \pi_1(\hat{C}_1^A \setminus C_1^A). \tag{15}$$

Inequality (15) includes the case of strictly decreasing (increasing) sense of ignorance, $\pi_0(\hat{C}_0^A \setminus C_0^A) > (<)\, \pi_1(\hat{C}_1^A \setminus C_1^A)$, and the case of constant ignorance, $\pi_0(\hat{C}_0^A \setminus C_0^A) = \pi_1(\hat{C}_1^A \setminus C_1^A)$, as special instances.

The model of Karni and Vierø (2013, 2015) is the special case of growing awareness in which the decision maker exhibits a constant sense of ignorance, assigning zero probability to discovery of new consequences. In those works, discoveries of unforeseen consequences are unanticipated.

5. Concluding remarks

5.1. Small worlds

The definitions of the state space and the set of conceivable acts, derived from the entire sets of basic actions and consequences, depict the *grand world*. When facing specific decisions, however, it is natural to suppose that the decision maker constructs the relevant choice set as follows: First, she identifies the *relevant* courses of action, or relevant basic actions, available (e.g., lists the means of transportation and routes to go from here to there, lists the available treatments of an illness). Second, she identifies the *relevant* consequences of the relevant basic actions (e.g., getting there late or not at all, allergic reaction to medication or bad outcome of surgery). Third, she constructs the relevant state space. For a given specific decision problem, let $A_r \subset A$ and $C_r \subset C$ denote, respectively, the relevant set of basic actions and consequences. Using these primitives, construct the *relevant* state space, $C_r^{A_r}$. The set of *relevant* conceivable acts is F_r (that is, the set of all mappings from $C_r^{A_r}$ to C_r). The set of all mixtures of these, $\Delta(F_r)$, constitutes the *relevant* choice set. Suppose that the decision maker's preferences on $\Delta(F_r)$ is the restriction of $\succ$ to $\Delta(F_r)$.

In this context, unawareness amounts to failure (e.g., due to lack of attention, forgetfulness) to consider some relevant consequences when constructing the choice set for the decision problem at hand. In other words, some consequences that the decision maker is aware of and that should have been included in the set of relevant consequences, are neglected.

Analogously to awareness of unawareness, the decision maker may anticipate that she may have neglected to include in her deliberation some relevant consequences. Define an abstract consequence $x_r = \neg C_r$ to represent neglected relevant consequences. Then, application of the analysis of section 2 yields the probability the decision maker assigns to the possibility of failing to include relevant consequences, and the utility of the concern (fear) that relevant consequences have been neglected.

5.2. The evolution of beliefs about describable events

Theorem 2 concerns the evolution of the relative likelihoods of fully describable (and also of the relative likelihoods of imperfectly describable) events, in the wake of discovery of new consequences, but is silent on the absolute likelihoods. By contrast, Theorem 3 concerns the evolution of the absolute likelihoods of the imperfectly describable event. Therefore, combining the results of the two theorems makes it possible to discuss the magnitude of the change in beliefs about the likelihoods of fully describable events. For instance, suppose that a new discovery is accompanied by a constant sense of unawareness. By Theorem 3, $\pi_0(\hat{C}_0^A \setminus C_0^A) = \pi_1(\hat{C}_1^A \setminus C_1^A)$. But $\Sigma_{s \in C_0^A} \pi_0(s) + \pi_0(\hat{C}_0^A \setminus C_0^A) = 1$ and $\Sigma_{s \in C_0^A} \pi_1(s) + \Sigma_{s \in C_1^A \setminus C_0^A} \pi_1(s) + \pi_1(\hat{C}_1^A \setminus C_1^A) = 1$. Hence, probability mass must be shifted from the set of originally fully describable states C_0^A to $C_1^A \setminus C_0^A$, proportionally (that is, the probabilities of all the states in C_0^A must be reduced equiproportionally). Similarly, an increasing sense of unawareness requires that probability mass must be shifted from C_0^A to $\hat{C}_1^A \setminus C_0^A$ proportionally, and that some of this probability must be shifted to $\hat{C}_1^A \setminus C_1^A$. Decreasing sense of unawareness implies that some probability mass of the event $\hat{C}_0^A \setminus C_0^A$ is shifted toward the newly describable event $C_1^A \setminus C_0^A$. In the latter instance, the effect of growing awareness on the subjective probability assigned to the set of originally fully describable states, C_0^A, is unpredictable.

5.3. Awareness of unawareness: behavioral manifestations

The theory advanced in this paper presumes that decision makers are aware of their unawareness. In other words, we suppose that decision makers are aware of the possible existence of indescribable consequences but have no clue as to what they might be. To elicit a decision maker's probability of the partially describable event we specify bets that mention the payoff in this event. This procedure is justified on the aforementioned presumption.

Our approach raises a methodological issue, namely, is there a way of testing the presumption that a decision maker is aware of his unawareness?[20] Put differently, how can an observer infer, from a decision maker's choice behavior, that she is aware of her unawareness? Below we describe possible patterns of choice that would indicate that the decision maker is indeed aware of being unaware.[21]

Partially specified bets: Consider the example in Section 3.1 of two basic actions and two feasible consequences. The augmented conceivable state space is depicted in matrix (1). Let $c_1 \succ c_2$ and consider the set of *partially specified bets*: $B := \{b : C^A \to C\}$. Note that the domain of these bets is the event, C^A, that consists of all the fully describable states, and the payoffs are feasible consequences.

Consider two bets: b_1 is a bet on the event that a_1 results in c_1, and b_2 is the bet that a_2 results in c_1. Formally, $b_1 = c_{1\{s_1,s_3\}}c_2$ and $b_2 = c_{1\{s_1,s_2\}}c_2$, where the states s_i, $i = 1, 2, 3$ are the states depicted in matrix (1). Note that b_1 specifies the payoffs as follows: Pay c_1 in the event $\{s_1, s_3\}$ and c_2 otherwise. Hence, in the description of b_1 there is no mention of any event other than $\{s_1, s_3\}$. In particular, there is no mention of the event $\{s_5, ..., s_9\}$ that consists of partially describable and non-describable states. Similarly, for b_2.[22]

The following choice patterns indicate that the decision maker is aware that the domain of the bets may include the partially describable states:

Pattern 1: $b_1 \sim b_2$ and $\neg(a_1 \sim a_2)$. The indifference, $b_1 \sim b_2$, means that the decision maker regards the events "a_1 pays c_1" and "a_2 pays c_1" as equally likely. If the decision maker is only aware of the fully describable states then b_1 and b_2 are replicas of a_1 and a_2, respectively. Therefore, a_1 and a_2 should be indifferent to one another. The fact that they are not is an indication that the decision maker considers possible events that are not fully describable.

Pattern 2: $b_i \succ b_j$, and $a_j \succcurlyeq a_i$, $i = 1, 2$, $j \neq i$. By the same reasoning as above, if the decision maker is only aware of the fully describable states, b_1 and b_2 are replicas of a_1 and a_2, respectively. Hence, $b_1 \succ b_2$ would imply that $a_1 \succ a_2$. The fact that it is not indicates that the decision maker considers possible events that are partially describable events.

[20] We thank Larry Epstein for raising this issue.

[21] The issue of detecting unawareness itself is investigated in Schipper (2013).

[22] The standard practice in decision theory is that bets specify the payoffs in every state. By contrast, partially specified bets are defined on the fully describable part of the state space. This partial specification is intended to allow the decision maker to complete the supports of the bets, if the unmentioned part of the state space exists in her mind, in any way she can imagine.

To illustrate this point consider again the example of the two medications described above. A decision maker may bet that the new, untried, medication, which proved more effective in clinical trials, is more likely to result in success than the old, tried, one. Yet, being worried by the prospect of unknown side-effects and believing that such effects are more likely if the newly approved medication is used, she chooses to take the tried medication. These choice patterns indicate not only that the decision maker is aware of partially describable events. They suggest that she regards the occurrence of such events more likely under one basic action than under another.

5.4. "Unknown unknown" consequences in subjective expected utility theory

The main objectives of this paper are to (a) provide a framework for the study of decision makers' anticipation that there may be consequences of which they are unaware, (b) quantify the decision makers' subjective beliefs regarding the likelihood of such event and how these beliefs change upon discoveries of consequences of which they were unaware, and (c) quantify decision makers' attitudes toward possible discovery of consequences of which they are unaware. A key idea invoked in this investigation is a "catch-all" consequence, defined as "none of the existing consequences," designed to capture the notion of "unknown unknown" consequences. A natural question is: Could this investigation be conducted by incorporating the "catch-all" consequence into a standard expected utility setting à la Anscombe and Aumann (1963) or Savage (1954)?

Consider the Anscombe–Aumann model. Let S be a finite set of states and denote by X a finite set of outcomes. Let $\hat{x}$ denote the "catch-all" outcome "none of the outcomes in X" (that is, $\hat{x} = \neg X$). Define $\hat{X} = X \cup \{\hat{x}\}$ and denote by $\Delta\hat{X}$ the set of probability distributions on $\hat{X}$. Extending Anscombe and Aumann's approach, the choice set consists of all mappings from S to $\Delta\hat{X}$, representing alternative courses of action. This formulation of the Anscombe–Aumann model requires that acts assign objective probabilities to the non-describable "catch-all" outcome $\hat{x}$. It also requires that the decision maker can consider a choice that includes the constant act that pays off $\hat{x}$ in every state (that is, receiving $\hat{x}$ for sure) when, by definition, this "unknown unknown" consequence cannot be described in a meaningful way.[23] Put differently, such formulation *assumes* the existence of $\hat{x}$ and that it is treated like any other consequence, thus defeating the purpose of inferring from the decision maker's behavior whether, and to what degree, he believes that consequences of which he is unaware exists. If one did admit acts whose range is $\Delta\hat{X}$, the subjective expected utility model would assign utility to $\hat{x}$, similarly to the result in Section 3 above. However, because the state space is fixed, the model cannot accommodate the assignment of subjective probability to the discovery of consequences of which the decision maker is unaware, or the analysis of the change of these beliefs upon the discovery of such consequences, including the changes in the decision maker's sense of unawareness. In other words, adding a "catch-all" consequence to the standard models will defeat some of the main objectives of this work. These observations are not restricted to the subjective expected utility models, rather they apply to all models based on analytical frameworks that take the state space to be a primitive concept.

[23] The fact that $\hat{x}$ can be described, abstractly, negatively as "none of the elements in X" does not lend it concrete meaning.

5.5. Related literature

The exploration of the issue of (un)awareness in the literature has pursued three different approaches; the epistemic approach, the game-theoretic or interactive decision making approach, and the choice-theoretic approach.

The epistemic approach is taken in Fagin and Halpern (1988), Dekel, Lipman, and Rustichini (1998), Modica and Rustichini (1999), Halpern (2001), Heifetz, Meier, and Schipper (2006, 2008), Li (2009), Hill (2010), Board and Chung (2011), Walker (2014) and Halpern and Rego (2009, 2013). Among these, Board and Chung (2011), Walker (2014) and Halpern and Rego (2009, 2013) consider awareness of unawareness. Halpern and Rego (2009) provide a logic that allows for an agent to explicitly know that there exists a fact of which he is unaware. They do so by introducing quantification over variables in their language. In Halpern and Rego (2009), it is impossible for an agent to consider it possible that he is aware of all formulas in the language and also consider it possible that he is not aware of all formulas. Halpern and Rego (2013) remedies this problem, such that an agent can be uncertain about whether he is aware of all formulas. The choice theoretic model we present is in line with the latter Halpern and Rego approach. Schipper (2015) provides an excellent overview of the epistemic literature as well as of the literature on awareness and unawareness more generally.

The game-theoretic, or interactive decision making, approach is taken in Halpern and Rego (2008, 2014), Heifetz, Meier, and Schipper (2013a, 2013b), Heinsalu (2014), and Grant and Quiggin (2013). The last develops a model of games with awareness in which inductive reasoning may cause an individual to entertain the possibility that her awareness is limited. Individuals thus have inductive support for propositions expressing their own unawareness. In this paper, we implicitly assume inductive reasoning to motivate considering awareness of unawareness.

The choice-theoretic approach to unawareness or related issues is taken in Li (2008), Ahn and Ergin (2010), Schipper (2013), Lehrer and Teper (2014), Kochov (2010), Walker and Dietz (2011), Alon (2015), and Grant and Quiggin (2015). The former four are discussed in detail in Karni and Vierø (2013). Walker and Dietz (2011), Kochov (2010), and Grant and Quiggin (2015) consider decision makers who are aware of their potential unawareness, and are thus the papers closest related to the present paper.

Walker and Dietz (2011) take a choice theoretic approach to static choice under "conscious unawareness." In their model, unawareness materializes in the form of coarse contingencies (that is, their state space does not resolve all uncertainty). Their representation is similar to Klibanoff, Marinacci, and Mukerji's (2005) smooth ambiguity model. The model of Walker and Dietz (2011) differs from ours in several respects: theirs is a static model and thus does not consider the issue of updating when awareness increases, their approach to modeling the state space differs from ours, and in their model a decision maker's beliefs are not represented by a single probability measure.

Kochov (2010) develops an axiomatic model of dynamic choice in which a decision maker knows that her perception of the environment may be incomplete. This implies that the decision maker's beliefs are represented by a set of priors, and that as the decision maker's perception of the universe becomes more precise, the priors are updated according to Bayes rule. Kochov's work differs from ours in the way the state space and its evolution are modeled, and in the representation of decision makers' beliefs.

Grant and Quiggin (2015) model unawareness by augmenting a standard Savage state space with a set of "surprise states". They also augment the set of standard consequences, which is an interval, by two unforeseen consequences, which are divided into two types: one favorable which is better than any standard consequence, and one unfavorable, which is worse than any standard consequence. Unforeseen consequences can only arise in surprise states. Their representation can be interpreted as if the decision maker follows a two-stage decision procedure, first categorizing each act as being subject to favorable, unfavorable, or no surprises, and second ranking acts. All acts subject to favorable surprises are, by assumption, better than all acts subject to no surprise, which are in turn better than all acts subject to unfavorable surprises. Within each category, acts are evaluated according to an expected uncertain utility (EUU) representation. Unlike Grant and Quiggin, we make no assumptions about the nature of unforeseen consequences, rather the utility of the unknown that we derive reflects how the decision maker ranks these consequences relative to the existing consequences. Thus, in our model, an act which the decision maker views as possibly leading to an unforeseen consequence need not be ranked extreme relative to acts that she views as not leading to something unforeseen. Also, in our model, beliefs over states determine how likely the decision maker views a particular act to reveal new consequences, while in Grant and Quiggin the decision maker cannot quantify the likelihood of surprise states.

Statistical theories of inductive inference have long wrestled with the problem of how to deal with the potential existence of unknown and unsuspected phenomena and how, once such phenomena occur, to incorporate the new knowledge into the corpus of the decision maker's prior beliefs. Zabell (1992) describes a particular instance of this issue, known as the sampling of species problem, involving repeated sampling which may result in an observation whose existence was not suspected (e.g., a new species).[24] "On the surface there would appear to be no way of incorporating such new information into our system of beliefs, other than starting from scratch and completely reassessing our subjective probabilities. Coherence of old and new makes no sense here: there are no old beliefs for the new to cohere with." (Zabell, 1992, p. 206). Zabell proceeds to detail a process, anticipated by De Morgan, that accommodates situations in which the possible species to be observed is not supposed to be known ahead of time.[25]

Despite the similarity of the objectives, and to some extent structure (think of repeated sampling as different acts and observed species as consequences) the solution for the sampling of species problem and the conclusion of our approach, dubbed 'reverse Bayesianism', are quite distinct. Perhaps the most important distinction is the specification of the prior. In the solution to the sampling of species problem, the prior is induced by exchangeability applied to the distinguished class of random partitions. In other words, it is implied by the stochastic structure of the problem and, as a result, loses its subjective flavor. For instance, the De Morgan rule creates an additional category: "new species not yet observed" and assigns it the probability $(N+t+1)^{-1}$, where N is the number of observations and t the number of known species.[26] By contrast, in 'reverse Bayesianism' the prior is a representation of the decision maker's subjective beliefs, which includes an assignment of subjective probability to the event of observing an indescribable consequence. Moreover, unlike our model of 'reverse Bayesianism', the solution to the sampling of species problem neither requires, nor does it yield, a utility valuation of the newly observed species or of the anticipated, yet indescribable, species.

[24] We are grateful to Teddy Seidenfeld for calling our attention to Zabell's work.

[25] The process is based on the idea of exchangeability of random partitions and it yields a representation theorem, a distinguished class of random partitions, and a rule of succession, describing the updated beliefs following the discovery of new species.

[26] See Zabell (1992), p. 209.

6. Appendix

6.1. Proof of Lemma 1

(**Sufficiency**) Since $\Delta(F)$ is a convex set and $\succ$ satisfies (A.1)–(A.3), by the expected utility theorem, there exists a real-valued function, $V : F \to \mathbb{R}$, such that $\succ$ on $\Delta(F)$ is represented by expected utility: For all $\mu, \mu' \in \Delta(F)$,

$$\mu \succ \mu' \Leftrightarrow \sum_{f \in F} \mu(f) V(f) > \sum_{f \in F} \mu'(f) V(f) \tag{16}$$

Moreover, V is unique up to positive linear transformation.

To show that $V(f) = \sum\limits_{s \in \hat{C}^A} W_s(f(s))$, fix $f^* \in F$ and, for each $f \in F$ and $s \in \hat{C}^A$, let $f^s = f_{\{s\}} f^* \in F$ be defined by: $f^s(s) = f(s)$ and $f^s(t) = f^*(t)$ if $t \neq s$.

Let $|\hat{C}^A| = n$. Consider the mixed conceivable acts, $\mu \in \Delta(F)$ that assigns probability $1/n$ to f and probability $(n-1)/n$ to f^*, and $\mu' \in \Delta(F)$ that assigns probability $1/n$ to each f^s, $s \in \hat{C}^A$. Then, by the definition in (4), $\varphi(\mu) = \varphi(\mu')$. Thus, by (A.4), $\mu \sim \mu'$. By the representation in (16), the last indifference is equivalent to

$$\frac{1}{n} V(f) + \frac{n-1}{n} V(f^*) = \frac{1}{n} \sum_{s \in \hat{C}^A} V(f^s). \tag{17}$$

For each $s \in \hat{C}^A$, define $W_s(\cdot) : C \to \mathbb{R}$ as follows[27]:

$$W_s(c) = V(c_{\{s\}} f^*) - \frac{n-1}{n} V\left(f^*\right).$$

Thus, for $f \in \hat{F}$, $W_s(f(s)) = V(f^s) - \frac{n-1}{n} V(f^*)$. This implies that $\sum_{s \in \hat{C}^A} W_s(f(s)) = \sum_{s \in \hat{C}^A} V(f^s) - (n-1) V(f^*)$. Multiplying by $1/n$ on both sides together with (17) implies that

$$V(f) = \sum_{s \in \hat{C}^A} W_s(f(s)). \tag{18}$$

Plugging (18) into (16), we get

$$\mu \succ \mu' \Leftrightarrow \sum_{f \in F} \mu(f) \sum_{s \in \hat{C}^A} W_s(f(s)) > \sum_{f \in F} \mu'(f) \sum_{s \in \hat{C}^A} W_s(f(s)). \tag{19}$$

(**Necessity**) This is immediate.

(**Uniqueness**) The uniqueness of $\{W_s\}_{s \in \hat{C}^A}$ follows from that of V. To see this, define $\hat{W}_s(\cdot) = b W_s(\cdot) + d_s$, $b > 0$, for all $s \in \hat{C}^A$. By definition, for all $s \in \hat{C}^A$ and $c \in C$, $\hat{W}_s(c) = b\left[V(c_{\{s\}} f^*) - \frac{n-1}{n} V(f^*)\right] + d_s$. Hence,

$$\sum_{s \in \hat{C}^A} \hat{W}_s(f(s)) = b \sum_{s \in \hat{C}^A} W_s(f(s)) + \sum_{s \in \hat{C}^A} d_s = bV(f) + d,$$

[27] Recall that c denotes both the outcome c and the constant act whose payoff is c in every state.

where $d = \sum_{s \in \hat{C}^A} d_s$. Since V is unique up to positive linear transformation, $\hat{V} = bV + d$ represents the same preferences as V. Hence, $\{\hat{W}_s\}_{s \in \hat{C}^A}$ represents the same preferences as $\{W_s\}_{s \in \hat{C}^A}$. It is easy to show that $\hat{W}_s(c) = \hat{V}(c_{\{s\}} f^*) - \frac{n-1}{n} \hat{V}(f^*)$. ♠

6.2. *Proof of Proposition 1*

(**Sufficiency**) By (A.1)–(A.4) and Lemma 1, we have that for all $\mu, \mu' \in \Delta(F)$,

$$\mu \succ \mu' \Leftrightarrow \sum_{f \in F} \mu(f) \sum_{s \in \hat{C}^A} W_s(f(s)) > \sum_{f \in F} \mu'(f) \sum_{s \in \hat{C}^A} W_s(f(s)).$$

By definition (4),

$$\sum_{f \in F} \mu(f) \sum_{s \in \hat{C}^A} W_s(f(s)) = \sum_{s \in \hat{C}^A} \sum_{c \in C} \varphi_s(\mu)(c) W_s(c).$$

Fix a non-null $s' \in \hat{C}^A$ (that such s' exists is an implication of (A.6)) and define, for $p \in \Delta(C)$, $U(p) = \sum_{c \in C} W_{s'}(c) p(c)$. By (A.5), for any $p, q \in \Delta(C)$,

$$\sum_{c \in C} W_{s'}(c) p(c) > \sum_{c \in C} W_{s'}(c) q(c) \Leftrightarrow \sum_{c \in C} W_s(c) p(c) > \sum_{c \in C} W_s(c) q(c)$$

for all non-null $s \in \hat{C}^A$.

Thus, standard arguments imply that

$$\mu \succ \mu' \Leftrightarrow \sum_{s \in \hat{C}^A} U(\varphi_s(\mu)) \pi(s) > \sum_{s \in \hat{C}^A} U\left(\varphi_s(\mu')\right) \pi(s),$$

where U is continuous, non-constant, affine, real-valued, and unique up to positive linear transformations, and the probability measure π is unique and $\pi(s) = 0$ if and only if s is null. This completes the proof of $(ii.a)$.

For all $f \in F$, a result analogous to Lemma 1 holds for $\succ$ on $\Delta(F_{\hat{C}^A \setminus C^A}(f))$. That is, for all $f \in F$, and for all $\mu, \mu' \in \Delta(F_{\hat{C}^A \setminus C^A}(f))$,

$$\mu \succ \mu' \Leftrightarrow \sum_{g \in F_{\hat{C}^A \setminus C^A}(f)} \mu(g) \sum_{s \in \hat{C}^A} W_s(g(s)) > \sum_{g \in F_{\hat{C}^A \setminus C^A}(f)} \mu'(g) \sum_{s \in \hat{C}^A} W_s(g(s)).$$

Furthermore, arguments analogous to those just used to prove (ii.a) serve to prove (ii.b). In particular, fix $f \in F$. Then (A.5) implies that, for any $\hat{p}, \hat{q} \in \Delta(\hat{C})$ and, for all non-null $s, s' \in \hat{C}^A \setminus C^A$

$$\sum_{\hat{c} \in \hat{C}} W_{s'}(\hat{c}) \hat{p}(\hat{c}) > \sum_{\hat{c} \in \hat{C}} W_{s'}(\hat{c}) \hat{q}(\hat{c}) \Leftrightarrow \sum_{\hat{c} \in \hat{C}} W_s(\hat{c}) \hat{p}(\hat{c}) > \sum_{\hat{c} \in \hat{C}} W_s(\hat{c}) \hat{q}(\hat{c}).$$

Hence, by (A.5), we have that for all $\mu, \mu' \in \Delta(F_{\hat{C}^A \setminus C^A}(f))$,

$$\mu \succ \mu' \Leftrightarrow \sum_{s \in \hat{C}^A \setminus C^A} U_f^*(\varphi_s(\mu)) \phi(s) > \sum_{s \in \hat{C}^A \setminus C^A} U_f^*\left(\varphi_s(\mu')\right) \phi(s),$$

where U_f^* is continuous, non-constant, affine, real-valued, and unique up to positive linear transformations, and the probability measure ϕ is unique, and $\phi(s) = 0$ if and only if s is null. This completes the proof of $(ii.b)$.

(**Necessity**) The proof that $(ii.a)$ and $(ii.b)$ imply (A.1)–(A.6) on the respective domains is straightforward. Since $\cup_{f\in F} F_{\hat{C}^A\setminus C^A}(f) = F^*$ and $F_{\hat{C}^A\setminus C^A}(f)) \cap F \neq \emptyset$ for all $f \in F$, the axioms necessarily hold on all of $\Delta(F^*)$.

(**Uniqueness**) Follows from standard arguments. ♠

6.3. *Proof of Theorem 1*

(**Sufficiency**) We give the part of the proof that does not follow directly from Proposition 1. The representations (6) and (7) imply that, for all $f \in F$ and for all $p, q \in \Delta(C)$, $p_{\hat{C}^A\setminus C^A} f \succ q_{\hat{C}^A\setminus C^A} f$ if and only if $U_f^*(p) > U_f^*(q)$. Hence, with appropriate normalization, for all $p \in \Delta(C)$, $U_f^*(p) = U(p)$, for all $f \in F$. Therefore, $U_f^*(p)$ is independent of f.

Suppose that $c^* \succ x \succ c_*$, let $\hat{p} = \alpha c^* + (1-\alpha) c_*$ be such that $\hat{p}_{\hat{C}^A\setminus C^A} f \sim x_{\hat{C}^A\setminus C^A} f$. By representation (7), this is equivalent to $U_f^*(\hat{p}) = U_f^*(x)$. Then, by axiom (A.7) and representation (7) we have that $U_g^*(\hat{p}) = U_g^*(x)$, for all $g \in F$. But $U_f^*(\hat{p}) = U_f^*(x)$ is equivalent to

$$U_f^*(x) = \alpha U\left(c^*\right) + (1-\alpha) U(c_*),$$

and $U_g^*(\hat{p}) = U_g^*(x)$ is equivalent to

$$U_g^*(x) = \alpha U\left(c^*\right) + (1-\alpha) U(c_*).$$

Hence, $U_f^*(x) = U_g^*(x) := u(x)$, for all $f, g \in F$.

Suppose instead that $x \succcurlyeq c^* \succ c_*$, and let $\hat{p} = \alpha x + (1-\alpha) c_*$ be such that $\hat{p}_{\hat{C}^A\setminus C^A} f \sim c^*_{\hat{C}^A\setminus C^A} f$. By representation (7), this is equivalent to $U_f^*(\hat{p}) = U_f^*(c^*)$. Then, by axiom (A.7) and representation (7) we have that $U_g^*(\hat{p}) = U_g^*(c^*)$ for all $g \in F$. But $U_f^*(\hat{p}) = U_f^*(c^*)$ is equivalent to

$$\alpha U_f^*(x) + (1-\alpha) U(c_*) = U(c^*),$$

and $U_g^*(\hat{p}) = U_g^*(x)$ is equivalent to

$$\alpha U_g^*(x) + (1-\alpha) U(c_*) = U(c^*).$$

Solving for $U_f^*(x)$ and $U_g^*(x)$ we get,

$$U_f^*(x) = U_g^*(x) = \frac{U(c^*) - U(c_*)}{\alpha} + U(c_*) := u(x)$$

for all $f, g \in F$.

Finally, if $c^* \succ c_* \succcurlyeq x$ let $\hat{p} = \alpha c^* + (1-\alpha) x$ such that $\hat{p}_{\hat{C}^A\setminus C^A} f \sim c_{*\hat{C}^A\setminus C^A} f$ then, by the same argument,

$$u_f^*(x) = u_g^*(x) = \frac{U(c_*) - \alpha U(c^*)}{1-\alpha} := u(x)$$

for all $f, g \in F$.

It follows that $U^*(\hat{p}) = \sum_{c\in C} \hat{p}(c) U(c) + \hat{p}(x) u(x)$, for all $\hat{p} \in \Delta(\hat{C})$.

The uniqueness of the subjective probabilities is implied by the uniqueness of the subjective probabilities in Proposition 1.[28]

(**Necessity**) The necessity of axioms (A.1)–(A.6) follows from Proposition 1. The necessity of (A.7) is immediate. ♠

[28] The uniqueness of π in conjunction with Proposition 1 imply that $\phi(s) = \pi(s)/\pi(\hat{C}^A \setminus C^A)$, for all $s \in \hat{C}^A \setminus C^A$.

6.4. Proof of Theorem 2

Necessity is immediate. We shall prove sufficiency.

Suppose that (i) holds. The representations (11) and (12) are implied by Proposition 1. By (11) and (12), the restriction of $\succ_0$ and $\succ_1$ to the mixed conceivable acts in $\Delta(C_0)$ whose support is the subset of constant conceivable acts in F_0, implies that, for any $p, q \in \Delta(C_0)$, $U_0(p) > U_0(q)$ if and only if $p \succ_0 q$ and that $U_1(p) > U_1(q)$ if and only if $p \succ_1 q$. By (A.8), $p \succ_0 q$ if and only if $p \succ_1 q$. Thus, by the uniqueness of the representations, U_0 and U_1 can be chosen so that $U_0 = U_1$ on $\Delta(C_0)$.

To prove (13), let $\lambda, \mu \in \Delta(F_0(h; E))$, and $\lambda', \mu' \in \Delta(F_1(h'; E))$ be as in (A.9). By definition of the functions φ^i, $\varphi_s^0(\lambda) = \varphi_s^0(\mu)$, for all $s \in \hat{C}_0^A \setminus E$ and $\varphi_s^1(\lambda') = \varphi_s^1(\mu')$, for all $s \in \hat{C}_1^A \setminus E$. Hence, $\lambda \sim_0 \mu$ if and only if

$$\sum_{s \in E} \pi_0(s) \left[U(\varphi_s^0(\lambda)) - U(\varphi_s^0(\mu)) \right] = 0 \tag{20}$$

and $\lambda' \sim_1 \mu'$ if and only if

$$\sum_{s \in E} \pi_1(s) \left[U(\varphi_s^1(\lambda')) - U(\varphi_s^1(\mu')) \right] = 0. \tag{21}$$

Axiom (A.9) together with (20) and (21) imply that

$$\sum_{s \in E} \frac{\pi_0(s)}{\pi_0(E)} \left[U(\varphi_s^0(\lambda)) - U(\varphi_s^0(\mu)) \right] = \sum_{s \in E} \frac{\pi_1(s)}{\pi_1(E)} \left[U(\varphi_s^1(\lambda')) - U(\varphi_s^1(\mu')) \right]. \tag{22}$$

By the hypothesis of axiom (A.9) (that is, $\lambda = \lambda'$ and $\mu = \mu'$ on E) and the definition of the functions φ^i, for all $s \in E$, $\varphi_s^0(\lambda) = \varphi_s^1(\lambda')$ and $\varphi_s^0(\mu) = \varphi_s^1(\mu')$. Hence, for all $s \in E$, $U(\varphi_s^0(\lambda)) = U(\varphi_s^1(\lambda'))$ and $U(\varphi_s^0(\mu)) = U(\varphi_s^1(\mu'))$. Thus, by (22),

$$\sum_{s \in E} \left[\frac{\pi_0(s)}{\pi_0(E)} - \frac{\pi_1(s)}{\pi_1(E)} \right] \left[U(\varphi_s^0(\lambda)) - U\left(\varphi_s^0(\mu)\right) \right] = 0. \tag{23}$$

The last equation implies (13).[29]

[29] To see this, let $E = \{s, s'\} \subset C_0^A$, such that, $U\left(\varphi_s^0(\lambda)\right) - U\left(\varphi_s^0(\mu)\right) = A > 0$, $U\left(\varphi_{s'}^0(\lambda)\right) - U\left(\varphi_{s'}^0(\mu)\right) := B < 0$, and $U\left(\varphi_{s''}^0(\lambda)\right) = U\left(\varphi_{s''}^0(\mu)\right)$, for all $s'' \in C_0^A \setminus E$. Then, by (23),

$$\left[\frac{\pi_0(s)}{\pi_0(s) + \pi_0(s')} - \frac{\pi_1(s)}{\pi_1(s) + \pi_1(s')} \right] A + \left[\frac{\pi_0(s')}{\pi_0(s) + \pi_0(s')} - \frac{\pi_1(s')}{\pi_1(s) + \pi_1(s')} \right] B = 0. \tag{24}$$

But

$$\frac{\pi_0(s')}{\pi_0(s) + \pi_0(s')} = 1 - \frac{\pi_0(s)}{\pi_0(s) + \pi_0(s')} \text{ and } \frac{\pi_1(s')}{\pi_1(s) + \pi_1(s')} = 1 - \frac{\pi_1(s)}{\pi_1(s) + \pi_1(s')}. \tag{25}$$

Substituting we obtain

$$\left[\frac{\pi_0(s')}{\pi_0(s) + \pi_0(s')} - \frac{\pi_1(s')}{\pi_1(s) + \pi_1(s')} \right] = - \left[\frac{\pi_0(s)}{\pi_0(s) + \pi_0(s')} - \frac{\pi_1(s)}{\pi_1(s) + \pi_1(s')} \right]. \tag{26}$$

Hence, (24) reduces to

To prove (14), let $\lambda, \mu \in \Delta(F_0(h; E))$ and $\lambda', \mu' \in \Delta(F_1(\hat{C}_0^A, h'; E'(E)))$ be as in (A.10). Then, $\varphi_s(\lambda) = \varphi_s(\mu)$, for all $s \in \hat{C}_0^A \setminus E$ and $\varphi_s(\lambda') = \varphi_s(\mu')$, for all $s \in \hat{C}_1^A \setminus E'(E)$. Hence, by Theorem 1, $\lambda \sim_0 \mu$ if and only if

$$\sum_{s \in E} \frac{\pi_0(s)}{\pi_0(E)} [U(\varphi_s(\lambda)) - U(\varphi_s(\mu))] = 0$$

and $\lambda' \sim_1 \mu'$ if and only if

$$\sum_{s \in E} \frac{\pi_1(E(s))}{\pi_1(E'(E))} \left[U(\varphi_s(\lambda')) - U\left(\varphi_s(\mu')\right)\right] = 0.$$

By the hypothesis of axiom (A.10), for all $s \in E$, $\varphi_{\hat{s}}\left(\lambda'\right) = \varphi_s(\lambda)$ and $\varphi_{\hat{s}}(\mu') = \varphi_s(\mu)$, for all $\hat{s} \in E(s)$. Hence, $U(\varphi_s(\lambda)) = U(\varphi_{\hat{s}}(\lambda'))$ and $U(\varphi_s(\mu)) = U(\varphi_{\hat{s}}(\mu'))$, for all $\hat{s} \in E(s)$. Thus,

$$\sum_{s \in E} \left[\frac{\pi_0(s)}{\pi_0(E)} - \frac{\pi_1(E(s))}{\pi_1(E'(E))}\right] [U(\varphi_s(\lambda)) - U(\varphi_s(\mu))] = 0.$$

By the same argument as in footnote 29, the last equation implies (14).

The uniqueness is an implication of the uniqueness in Proposition 1. ♠

6.5. *Proof of Theorem 3*

(**Sufficiency**) That the axioms imply existence of a representation as in Theorem 2 follows from the proof of Theorem 2. Let $\lambda, \mu \in \Delta(F_0)$ and $\lambda', \mu' \in \Delta(F_1)$ be as in Axiom (A.11). Suppose that $\mu \sim_0 \lambda$. But $\mu \sim_0 \lambda$ if and only if

$$c_{*C_0^A} c^* \sim_0 \eta c_* + (1 - \eta) c^*. \tag{30}$$

By the representation in (11) the last indifference holds if and only if

$$U_0(c_*) \pi_0(C_0^A) + U_0\left(c^*\right)\left(1 - \pi_0(C_0^A)\right) = U_0(c_*)\eta + U_0\left(c^*\right)(1 - \eta). \tag{31}$$

But, $U_0(c^*) > U_0(c_*)$. Hence, (31) holds if and only if

$$\eta = \pi_0(C_0^A). \tag{32}$$

By Axiom (A.11), $\mu \sim_0 \lambda$ implies that $\lambda' \succcurlyeq_1 \mu'$, which is equivalent to

$$\eta c_* + (1 - \eta) c^* \succcurlyeq_1 c_{*C_1^A} c^*. \tag{33}$$

$$\left[\frac{\pi_0(s)}{\pi_0(s) + \pi_0(s')} - \frac{\pi_1(s)}{\pi_1(s) + \pi_1(s')}\right](A - B) = 0. \tag{27}$$

But $A - B > 0$. Therefore, it must be that

$$\frac{\pi_0(s)}{\pi_1(s)} = \frac{\pi_0(s) + \pi_0(s')}{\pi_1(s) + \pi_1(s')}. \tag{28}$$

Thus,

$$\frac{\pi_0(s)}{\pi_1(s)} = \frac{\pi_0(s')}{\pi_1(s')}. \tag{29}$$

By the representation in (12), (33) holds if and only if

$$U_1(c_*)\,\eta + U_1\left(c^*\right)(1-\eta) \geq U_1(c_*)\,\pi_1(C_1^A) + U_1\left(c^*\right)\left(1-\pi_1(C_1^A)\right). \tag{34}$$

Hence, by the same argument as above, (34) holds if and only if

$$\pi_1(C_1^A) \geq \eta. \tag{35}$$

By (32) and (35) we have that

$$\pi_1(C_1^A) \geq \pi_0(C_0^A), \tag{36}$$

which is equivalent to $\pi_1(\hat{C}_1^A \setminus C_1^A) \leq \pi_0(\hat{C}_0^A \setminus C_0^A)$. The inequality in (36) is strict if and only if $\lambda' \succ_1 \mu'$ in Axiom (A.11), and holds with equality if and only if $\lambda' \sim_1 \mu'$ in Axiom (A.11).

(**Necessity**) The necessity of axioms (A.1)–(A.10) follows from the proof of Theorem 2. To show that (A.11) holds, let $\mu, \lambda \in \Delta(F_0)$ and $\mu', \lambda' \in \Delta(F_1)$ be as in (A.11). By (11), $\mu \sim_0 \lambda$ if and only if

$$U_0(c_*)\,\pi_0(C_0^A) + U_0\left(c^*\right)\left(1-\pi_0(C_0^A)\right) = U_0(c_*)\,\eta + U_0\left(c^*\right)(1-\eta). \tag{37}$$

Since $U_0(c^*) > U_0(c_*)$, (37) holds if and only if

$$\eta = \pi_0(C_0^A). \tag{38}$$

Suppose now that $\mu' \succ_1 \lambda'$. By (12), $\mu' \succ_1 \lambda'$ if and only if

$$U_1(c_*)\,\pi_1(C_1^A) + U_1\left(c^*\right)\left(1-\pi_1(C_1^A)\right) > U_1(c_*)\,\eta + U_1\left(c^*\right)(1-\eta).$$

Since $U_0(c^*) > U_0(c_*)$, this holds if and only

$$\pi_1(C_1^A) < \eta. \tag{39}$$

Now, expressions (38) and (39) imply that

$$\pi_0(C_0^A) > \pi_1(C_1^A). \tag{40}$$

However, by (15), $\pi_0(C_0^A) \leq \pi_1(C_1^A)$, which contradicts (40). ♠

References

Agranov, Marina, Ortoleva, Pietro, 2016. Stochastic choice and preferences for randomization. J. Polit. Econ. http://dx.doi.org/10.1086/689774 (forthcoming).

Ahn, David S., Ergin, Haluk, 2010. Framing contingencies. Econometrica 78, 655–695.

Alon, Shiri, 2015. Worst-case expected utility. J. Math. Econ. 60, 43–48.

Anscombe, Francis J., Aumann, Robert J., 1963. A definition of subjective probability. Ann. Math. Stat. 43, 199–205.

Auster, Sarah, 2013. Asymmetric awareness and moral hazard. Games Econ. Behav. 82, 503–521.

Board, Oliver, Chung, Kim-Sau, 2011. Object-Based Unawareness: Axioms. Working paper. University of Minnesota.

Dekel, Eddie, Lipman, Bart, Rustichini, Aldo, 1998. Standard state-space models preclude unawareness. Econometrica 66, 159–173.

Fagin, Ronald, Halpern, Joseph Y., 1988. Belief, awareness and limited reasoning. Artif. Intell. 34, 39–76.

Gibbard, Allan, Harper, W.L., 1978. Counterfactuals and two kinds of expected utility. In: Foundations and Applications of Decision Theory, vol. 1. Reidel Publishing Co., Dordrecht, pp. 125–162.

Grant, Simon, Quiggin, John, 2013. Inductive reasoning about unawareness. Econ. Theory 54, 717–755.

Grant, Simon, Quiggin, John, 2015. A preference model for choice subject to surprise. Theory Decis. 79, 167–180.

Halpern, Joseph Y., 2001. Alternative semantics for unawareness. Games Econ. Behav. 37, 321–339.

Halpern, Joseph Y., Rego, Leandro C., 2008. Interactive awareness revisited. Games Econ. Behav. 62, 232–262.
Halpern, Joseph Y., Rego, Leandro C., 2009. Reasoning about knowledge of unawareness. Games Econ. Behav. 67, 503–525.
Halpern, Joseph Y., Rego, Leandro C., 2013. Reasoning about knowledge of unawareness revisited. Math. Soc. Sci. 65, 73–84.
Halpern, Joseph Y., Rego, Leandro C., 2014. Extensive games with possibly unaware players. Math. Soc. Sci. 70, 42–58.
Heifetz, Aviad, Meier, Martin, Schipper, Burkhard C., 2006. Interactive unawareness. J. Econ. Theory 130, 78–94.
Heifetz, Aviad, Meier, Martin, Schipper, Burkhard C., 2008. A canonical model for interactive unawareness. Games Econ. Behav. 62, 304–324.
Heifetz, Aviad, Meier, Martin, Schipper, Burkhard C., 2013a. Unawareness, beliefs, and speculative trade. Games Econ. Behav. 77, 100–121.
Heifetz, Aviad, Meier, Martin, Schipper, Burkhard C., 2013b. Dynamic unawareness and rationalizable behavior. Games Econ. Behav. 81, 50–68.
Heinsalu, Sander, 2014. Universal type structures with unawareness. Games Econ. Behav. 83, 255–266.
Hill, Brian, 2010. Awareness dynamics. J. Philos. Log. 39, 113–137.
Karni, Edi, Schmeidler, David, 1991. Utility theory with uncertainty. In: Hildenbrand, Werner, Sonnenschein, Hugo (Eds.), Handbook of Mathematical Economics, vol. IV. Elsevier Science Publishers B.V.
Karni, Edi, Vierø, Marie-Louise, 2013. Reverse bayesianism: a choice-based theory of growing awareness. Am. Econ. Rev. 103, 2790–2810.
Karni, Edi, Vierø, Marie-Louise, 2015. Probabilistic sophistication and reverse bayesianism. J. Risk Uncertain. 50, 189–208.
Klibanoff, P., Marinacci, M., Mukerji, S., 2005. A smooth model of decision making under ambiguity. Econometrica 73, 1849–1892.
Kochov, Asen, 2010. A Model of Limited Foresight. Working paper. University of Rochester.
Kreps, David M., 1988. Notes on the Theory of Choice. Westview Press, Boulder.
Lehrer, Ehud, Teper, Roee, 2014. Extension rules or what would the sage do? Am. Econ. J. Microecon. 6, 5–22.
Li, Jing, 2008. A Note on Unawareness and Zero Probability. Working Paper. University of Pennsylvania.
Li, Jing, 2009. Information structures with unawareness. J. Econ. Theory 144, 977–993.
Ma, Wenjun, Schipper, Burkhard C., 2016. Does Exposure to Unawareness Affect Risk Preferences? A Preliminary Result. Working paper. U.C., Davis.
Machina, Mark J., 2003. States of the world and the state of decision theory. In: Meyer, Donald (Ed.), The Economics of Risk. W. E. Upjohn Institute for Employment Research.
Maskin, Eric, Tirole, Jean, 1999. Unforeseen contingencies and incomplete contracts. Rev. Econ. Stud. 66, 83–114.
Mengel, Friederike, Tsakas, Elias, Vostroknutov, Alexander, 2016. Past experience of uncertainty affects risk aversion. Exp. Econ. 19, 151–176.
Modica, Salvatore, Rustichini, Aldo, 1999. Unawareness and partitional information structures. Games Econ. Behav. 27, 256–298.
Savage, Leonard J., 1954. The Foundations of Statistics. Wiley, New York.
Schipper, Burkhard C., 2013. Awareness-dependent subjective expected utility. Int. J. Game Theory 42, 725–753.
Schipper, Burkhard C., 2015. Awareness. In: van Ditmarsch, Hans V., Halpern, Joseph Y., van der Hoek, Wiebe, Kooi, Barteld P. (Eds.), Handbook of Epistemic Logic. College Publications, London.
Schipper, Burkhard C., Woo, H.Y., 2015. Political Awareness, Microtargeting of Voters and Negative Electoral Campaigning. Working paper. The University of California, Davis.
Schmeidler, David, Wakker, Peter, 1987. Expected utility and mathematical expectation. In: Eatwell, John, Milgate, Murray, Newman, Peter (Eds.), The New Palgrave: A Dictionary of Economics. Macmillan Press.
Stalnaker, R.C., 1972. Letter to David Lewis. In: Harper, W.L., Stalnaker, R.C., Pearce, G. (Eds.), 1981. IFS Conditionals, Beliefs, Decision, Chance, and Time. Reidel Publishing Co., Dordrecht.
Walker, Oliver, 2014. Unawareness with “possible” possible worlds. Math. Soc. Sci. 70, 23–33.
Walker, Oliver, Dietz, Simon, 2011. A Representation Result for Choice under Conscious Unawareness. Working Paper. London School of Economics and Political Science.
Zabell, Sandy L., 1992. Predicting the unpredictable. Synthesis 90, 205–232.