Centre for
Economic
and Financial
Research
at
New Economic
School

October 2016

Subjective Contingencies and Limited Bayesian Updating

Stefania Minardi
Andrei Savochkin

Working Paper No 222

CEFIR / NES Working Paper series

Subjective Contingencies and Limited Bayesian Updating*

Stefania Minardi[†] and Andrei Savochkin[‡]

October 28th, 2016

We depart from Savage's (1954) common state space assumption and introduce a model that allows for a subjective understanding of uncertainty. Within the revealed preference paradigm, we uniquely identify the agent's subjective state space via her preferences conditional on incoming information. According to our representation, the agent's subjective contingencies are coarser than the analyst's states; she uses an additively separable utility with respect to her set of contingencies; and she adopts an updating rule that follows the Bayesian spirit but is limited by her perception of uncertainty. We illustrate our theory with an application to the Confirmatory Bias.

1 Introduction

1.1 Motivation and objectives

The assumptions of the standard expected utility theory of choice under uncertainty have been challenged from many angles to produce various models with much higher descriptive power. These models have been used successfully in applications that range from mechanism design in microeconomics to the equity premium puzzle in macroeconomics. In this paper, we focus on the assumption of the standard theory that is not stated as an axiom in Savage (1954) but, rather, is built into his framework — that the *same* state space is used both by the analyst to formulate the model and by the decision maker to evaluate the

*We thank Sarah Auster, Pierpaolo Battigalli, Larry Epstein, John Geanakoplos, Marco LiCalzi, Edi Karni, Massimo Marinacci, Sujoy Mukerji, Jawwad Noor, Pietro Ortoleva, Nicola Pavoni, Zvi Safra, Larry Samuelson, Tomasz Strzalecki, Roee Teper, and, especially, Paolo Ghirardato, Itzhak Gilboa, and Fabio Maccheroni for many comments and very useful discussions, as well as seminar audiences at Bocconi, Harvard, HEC Paris, New Economic School, NYU, Pittsburgh University, Yale, IDGP 2015 Workshop (Barcelona), and PRIN 2016 Workshop (Ca' Foscari). The authors gratefully acknowledge the financial support of the Investissements d'Avenir program (ANR-11-IDEX-0003/Labex Ecodec/ANR-11-LABX-0047).

†Economics and Decision Sciences Department, HEC Paris, minardi@hec.fr, http://www.hec.fr/minardi/

‡New Economic School, asavochkin@nes.ru, http://pages.nes.ru/asavochkin/

uncertain prospects at hand.[1]

Consider the real-world decision problem of an individual choosing assets for investment. For such a problem, the formulation of a detailed state space on which to base the decision presents a number of tradeoffs. One state space may consist of vectors of potential returns in, say, one year of all assets traded in the market. For the decision problem at hand, such a state space is exhaustive. However, it is cognitively quite demanding to operate with a state space that contains innumerable combinations of values of several thousand variables. At the other extreme, the decision maker may consider only a few variables, such as the overall state of the economy, the consumer confidence index, or the prices of natural resources. All combinations of possible realizations of these variables form a different set of contingencies that is obviously coarse. However, besides their operational efficiency, the elements of this space may be interpreted in natural language and, ultimately, may be much easier to assign probabilities to. Clearly, there are many more ways in which the decision maker can organize her process of reasoning about the surrounding uncertainty.

This paper develops a revealed preference theory that departs from the standard theory of choice under uncertainty by giving the decision maker the freedom to come up with her own subjective contingencies for thinking about uncertainty. Our main goal is to link the structure of these contingencies to observable choice behavior.

Note that, unlike the vast literature initiated by Kreps (1979), our theory of subjective contingencies does not attempt to capture the decision maker's "states of mind." Put differently, the uncertainty in the model is not about her future taste or risk aversion or any other aspect of her personality. Our objective is quite different, and the only source of uncertainty in the model is the outside world. Formally, our agent can be described by a single utility function that represents her tastes (as well as a single belief, if the modeler so wishes), while her "subjective states" represent her perception of the uncertainty about *nature*.

The contribution of the paper can be described as follows. First, we propose a way to address the situation in which it is desirable to disentangle the decision maker's and the analyst's views. In particular, the decision maker is allowed to select her own "state

[1] The shortcomings of the Savagean state space assumption are discussed by, among others, Dekel, Lipman, and Rustichini (1998, 2001), Gilboa (2009, p. 136–137), Ghirardato (2001), and Epstein, Marinacci, and Seo (2007).

space” to reason about uncertainty. As our theory demonstrates, such situations can be modeled within the classic Savagean framework with the standard extension to allow for partial resolution of uncertainty and updating. The decision maker’s subjective perception of uncertainty — the set of “states” or contingencies with which she operates — can be identified from choice data, and, hence, does not need to be fixed exogenously or known by the analyst in advance. The identification comes from a novel source — the analysis of the decision maker’s conditional preferences that incorporate the information that some event has occurred. Moreover, our approach is applicable regardless of whether preferences conform to the expected utility model or display some ambiguity-sensitive pattern. Second, we propose a tractable representation of the decision maker’s preferences that is consistent with the existing models in the literature, as well as a novel updating rule to account for incoming information. The proposed updating rule is based on the concept of consistency between the arrived information and a subjective contingency of the decision maker. It is behaviorally different from the standard Bayesian updating but adheres to its spirit within the limits of the decision maker’s perception of events. Third, on the more technical side, we formulate the well-known property of dynamic consistency from a subjective perspective so that the model allows the decision maker to view uncertainty differently from the analyst. In addition, we introduce a few novel axioms that put discipline into the model: they reflect the idea that the decision maker has a certain level of understanding of the relationships among events and does not suffer from unrelated cognitive limitations or misconceptions, such as the conjunction fallacy.

If the decision maker’s subjective “state space” is effectively equivalent to the analyst’s state space, then the decision maker understands all events fully. If she does not understand some events fully, the decision maker’s dynamic choice exhibits patterns that the analyst may regard as errors — classified as inclusion and exclusion errors. As an illustration of the behavioral implications, we present a simple example that shows that subjective understanding of uncertainty can lead to the Confirmatory Bias, a well-known phenomenon that the literature describes as mainly a manifestation of bounded rationality.

One might find it conceivable that the agent’s coarse “state space” (and, hence, the type of behavior discussed) could be related to unawareness. While this is definitely a possibility, we note that the agent’s selection of a particular set of subjective contingencies — her subjective "state space" — also may be dictated by cognitive constraints; be a result of some optimization that takes into account costs of reasoning through large state spaces; or even be a matter of taste, by which the decision maker deliberately decides to ignore some dimensions of the uncertainty that she is facing. Moreover, we are not interested in the reasons that have prompted the decision maker to use a particular set of subjective contingencies. From a revealed preference perspective, we limit our scope to developing a way to learn from the choice data what these subjective contingencies are. To this end, we provide a representation result and an updating rule; and we discuss some implications of our assumptions, such as inclusion and exclusion errors, that are independent of the underlying reasons for using a subjective state space — whether unawareness or some other factor.[2]

To illustrate our objectives, we consider a stylized problem of a decision maker who evaluates the possibility of investing in a firm's stock at the beginning of a year. By the end of the year, the firm will pay out the accumulated profit as a dividend and then close. For simplicity, suppose that there is no time discounting, and the dividend can take the value of only \$4, \$2, or \$1 per share. Thus, the investor faces a decision problem under uncertainty, whereas the analyst observes the investor's willingness to buy shares, as well as their derivatives, at various prices. In addition, the data also contain the investor's willingness to buy the instruments, conditional on mid-year financial reports stating whether the interim profit is low or high. For the sake of discussion, suppose that the investor ex ante believes that the dividend of \$4 is possible — e.g., she is willing to pay a positive price for a European call option with the strike price of \$3. Also, suppose that, conditional on learning that the interim profit is low, she is willing to pay, at most, \$1 for the firm's share — that is, she thinks that the dividend can only be \$1. Now, let us bring the analyst into the picture. Suppose that he knows all the details about the firm's operation and describes the uncertainty about the payouts with the state space $\Omega = \{\omega_g, \omega_m, \omega_b\}$, where ω_g

[2] We also wish to note that building a formal bridge between our model and a theory of unawareness not only requires choosing a specific way of modeling unawareness, but also taking a stand about what the decision maker is unaware of. As will be clear from our setup, it cannot be that the decision maker is unaware of some individual states or events, as our theory requires the decision maker to form preferences conditional on all events and imposes a version of dynamic consistency. Hence, a successful link between our model and unawareness requires postulating that the source of the decision maker's potential unawareness are some other types of objects, or, more likely, a *relationship* between objects.

represents the scenario of a strong demand for the firm's product throughout the year that the firm correctly anticipates and responds to by expanding; ω_m represents the scenario of a strong demand to which the firm does not respond; and ω_b represents the scenario of a weak demand. Holding one share of the firm entitles the holder to the dividend payments of \$4 in state ω_g, \$2 in state ω_m, and \$1 in state ω_b. To the analyst, the firm's business model also implies that a low-profit mid-year report corresponds to the subset $\{\omega_g, \omega_b\}$ of the state space — the report means either that demand is weak or that the firm is sacrificing its profits temporarily in order to expand and reap higher profits by the end of the year. Hence, the analyst concludes that the investor's choices cannot be accommodated within the standard Savagean paradigm, as they are inconsistent with assigning any probability, zero or nonzero, to the state ω_g. Indeed, if the probability of ω_g (and, hence, receiving the dividend of \$4) were zero, then the investor would not be willing to pay for an option with the strike price of \$3. If the probability of ω_g were nonzero, then she would be willing to pay more than \$1 for the firm's share conditional on the event $\{\omega_g, \omega_b\}$.[3] Moreover, the analyst can infer that the investor probably perceives the uncertainty differently, in a way that is incompatible with his state space Ω. Given the analyst's knowledge, the event "the interim profit is low" is formally mapped to the subset $\{\omega_g, \omega_b\}$ of his state space and is interpreted differently by the decision maker: it seems that the decision maker does not realize that the interim profit can be low because the firm is expanding and that expansion can actually be a good signal for an investor. We will revisit this example after stating our main theorem. There, we will also illustrate the idea that our model has a certain discipline and cannot describe arbitrary departures from the standard expected utility paradigm. In particular, given the investor's choices described above, our model implies that she must believe that the dividend of \$4 is possible after some other event — for instance, the event $\{\omega_m\}$ of the high-profit mid-year report. At the level of interpretation, it might be that, in the simplified world of the decision maker, the higher the mid-year profit, the higher the expected end-of-year dividend.

As this example suggests, violations of Bayesian updating will arise in our model as the

[3] This contradiction is naturally affected by the analyst's selection of his state space Ω. The state space described above is succinct and exhaustive given what the analyst knows. However, his selection of the state space may be driven not solely by that — any enlargement of the state space implies an expansion of the space of mappings from states to outcomes and, therefore, is also constrained by what is observable or contractible in a particular environment.

observable manifestation that the decision maker understands the surrounding uncertainty and processes incoming information through a simpler set of subjective contingencies compared to the analyst's state space. The literature has recognized the potential link between deviations from Bayes Rule and the complexity (or, more generally, the limited understanding) of the world. In that vein, Manski (2004), for instance, emphasizes the importance of eliciting the agent's view of the world in order to better understand the effective updating mechanism that she adopts.[4]

1.2 Preview of the results

Our model is constructed for the benefit of the analyst whose objective is to test whether the decision maker's preferences satisfy certain consistency properties and represent her choices with a particular utility function. To this end, we take the perspective of an external observer and focus on the analyst's model of the agent's behavior. We posit that the analyst operates in the Savage framework: he brings forth a state space Ω (which, of course, need not be known to the decision maker) and expresses any action that the decision maker may choose by a vector that assigns an outcome to each state in Ω. Following the tradition, an *act* refers to the formal representation of an action as a state-contingent profile of outcomes. We assume that the analyst's state space is comprehensive, in the sense that the decision maker cannot discern eventualities that are indistinguishable from the analyst's view.[5]

The decision maker's choices are modeled through a collection of preferences over acts, ex ante and conditional on the information that some event has occurred. A key aspect of our model is that we study the decision maker's behavior in the ordinal course of updating her preferences in response to the partial resolution of uncertainty; that is, pieces of information do not come as revelations that can change the agent's understanding of uncertainty. For instance, in our financial example, learning that the interim profit is low does not induce

[4] "[...] The information that persons receive rarely maps cleanly into a textbook exercise in probability updating. [...] Expectations formation in real life requires persons to assimilate government announcements, media reports, personal observations, and other forms of information that may be generated in obscure ways. [...] Understanding expectations formation will also require intensive probing of persons to learn how they perceive their environments and how they process such new information as they may receive" (Manski, 2004, pp. 1368-69).

[5] This assumption seems very natural in applications and is, implicitly, dominant in the vast axiomatic literature on decision making under uncertainty (see, e.g., Gilboa and Marinacci (2013)).

the investor to learn about the strategy of sacrificing profits for the sake of growth if such a possibility did not occur to her before, even though she can conceive of the possibility of a \$4 dividend per se.

Our main results show that a set of mild axioms is equivalent to the decision maker acting as if (i) she has in mind a collection of subjective contingencies represented by subsets of Ω that capture her coarse understanding of uncertainty; and (ii) she uses a particular updating procedure to account for incoming information.

According to our first representation result, the decision maker ranks acts f and g at the ex ante stage as

$$f \succsim g \quad \Leftrightarrow \quad \sum_{i \in \mathcal{S}} V_i(f) \geq \sum_{i \in \mathcal{S}} V_i(g),$$

and, conditional on the information that an event $E \subset \Omega$ has occurred, as

$$f \succsim_E g \quad \Leftrightarrow \quad \sum_{i \in \mathcal{S} \text{ consistent with } E} V_i(f) \geq \sum_{i \in \mathcal{S} \text{ consistent with } E} V_i(g).$$

The formal content of the expression "$i \in \mathcal{S}$ consistent with E" in this pseudo-representation is stated in Theorem 1. The index set $\mathcal{S}$ enumerates the decision maker's subjective contingencies, which form a partition of Ω (or a subset of Ω). The functionals V_i for $i \in \mathcal{S}$ capture the assignment of a utility level to an act for each contingency. Then, the decision maker calculates the overall value of an act using an additively separable criterion.

The link between ex ante and conditional preferences plays a key role in the model: to compare acts conditionally, the decision maker uses the same collection of functionals V_i but, similar to the standard Bayesian agent, considers only those contingencies that are consistent with the received information and discards the rest. Consistency of a contingency with an event is assessed subjectively and is the channel through which the agent reveals her coarse understanding of uncertainty. The reason is that our agent processes information by using her own subjective contingencies, which are less expressive than the analyst's state space. Consequently, her inferences from a given piece of information might be different from those she would make if she had the analyst's state space in mind. To provide some intuition, if the agent understands the uncertainty fully (we define this term formally later), and if, ex ante, she treats all states in Ω as possible, then $\mathcal{S}$ can be identified with Ω; in this case, a contingency $i \in \mathcal{S}$ is consistent with an event E if and only if $i \in E$. On the contrary, if the agent has a coarse understanding of uncertainty, then some (or all) contingencies in $\mathcal{S}$ can be identified with (non-singleton) subsets of Ω. In turn, such a limited understanding of uncertainty affects the way preferences are updated by prompting the agent to make inferences that the analyst may view as errors. In particular, the agent may treat a state as possible, conditional on some event E, even if such a state does not belong to E (*inclusion error*); and, conversely, she may treat a state as ex ante non-null but impossible, conditional on some event E, even if it does belong to E (*exclusion error*).

A noteworthy feature of our representation results is that, unlike Savagean states, each subjective contingency does not describe one complete resolution of uncertainty. In turn, the collection $\{V_i\}_{i\in\mathcal{S}}$ of functionals plays precisely the role of specifying the agent's evaluation of an act on each contingency i. Our first representation result does not impose any strong structure on these functionals and focuses, instead, on the key implications of allowing for a subjective view of the world. Our aim is to provide a general utility representation that serves as a unifying model for various specifications of (i) the mapping that determines the utility level associated with each act on each contingency; and (ii) the procedure establishing the consistency of each contingency with a realized event. Relatedly, we remain agnostic on whether or not the agent is aware of her coarse understanding of uncertainty. Nevertheless, the flexibility of our representation allows more-specific interpretations and leaves room to study whether the agent is pessimistic or optimistic toward resolution of uncertainty that she does not fully understand.[6]

Our second representation result is a special case of the general model and takes a particularly tractable form. Given any $f, g \in \mathcal{F}$,

$$f \succsim g \quad \Leftrightarrow \quad \sum_{i\in\mathcal{S}} u(f^*(i))\,\mu_i \geq \sum_{i\in\mathcal{S}} u(g^*(i))\,\mu_i,$$

where each $f \in \mathcal{F}$ is represented by a subjective act $f^* : \mathcal{S} \to X$, mapping subjective contingencies into outcomes, defined as $f^*(i) = \sum_{\omega\in\Omega} p_i(\omega) f(\omega)$ for all $i \in \mathcal{S}$. Conditional on an event E,

$$f \succsim_E g \quad \Leftrightarrow \quad \sum_{i\in\mathcal{S}(E)} u(f^*(i))\,\mu_{i|E} \geq \sum_{i\in\mathcal{S}(E)} u(g^*(i))\,\mu_{i|E},$$

where $\mu_{i|E} = \frac{\mu_i}{\sum_{j\in\mathcal{S}(E)} \mu_j}$ and $\mathcal{S}(E) := \{i \in \mathcal{S} : \sum_{\omega\in E} p_i(\omega) \geq \alpha_i\}$.

[6] For instance, in the special case in which the functions V_i take the maxmin form, one may argue that the decision maker behaves cautiously because she is aware of her coarse understanding of uncertainty.

In this representation, u denotes a von Neumann-Morgenstern utility function; μ_i is the subjective probability of the contingency i; and $p_i(\omega)$ is the relative weight of the state ω in the contingency i. The key part is the formula to determine the contingencies that are consistent with E (i.e., the set $\mathcal{S}(E)$): when checking for consistency, the agent "disregards" states that have weights below a certain threshold. Note that, although the agent's ex ante behavior is consistent with the Subjective Expected Utility model, she perceives uncertainty in a coarse way. Therefore, this specification can be thought of as a minimal departure from the standard paradigm that allows for a subjective understanding of uncertainty.

The two main theorems of the paper provide axiomatic characterizations of these preference representations by focusing on the way information is processed. Note that we impose a dynamic consistency property, which maintains that the decision maker is capable of contingent planning, and arriving information never comes as a surprise. Thus, she is not prompted to revise her plans. Hence, this axiom formally reflects our objective of studying the decision maker's understanding of uncertainty "here and now," without referring to the possibility of discovering new eventualities that she did not consider at an earlier stage and of acquiring a more sophisticated vision as uncertainty unfolds.

Finally, we perform a comparative statics analysis along multiple dimensions. First, we provide a behavioral characterization of the situations in which one agent understands uncertainty better than another. Second, we identify choice patterns that can be interpreted as updating errors (relative to the analyst's state space) and characterize the notion of being more (or less) prone to inclusion and exclusion errors.

The rest of the paper is organized as follows. Next, we discuss the related literature. Section 2 introduces the setup. Section 3 presents our general model: the axiomatic foundations and the representation theorem, as well as uniqueness and comparative statics results. Section 4 specializes to a more structured model akin to subjective expected utility. Example 2 in Subsection 4.2 uses this special case in an application to the Confirmatory Bias.

1.3 Related Literature

The most relevant works for our study are Mukerji (1997) and Ghirardato (2001), as well as Ahn and Ergin (2010).

Mukerji (1997) and Ghirardato (2001) study a decision maker who perceives the state space imperfectly. Mukerji (1997) introduces two kinds of state spaces: In his notation, Ω represents a state space on which the decision maker is capable of forming beliefs, and Θ is the space of "payoff-relevant" states on which the acts are defined. The connection between them — an implication mapping $\Gamma : \Omega \rightrightarrows \Theta$ — represents the decision maker's knowledge. Then, Mukerji discusses the relationship between beliefs on Ω and induced beliefs on Θ, and argues that the conservative attitude of the decision maker should lead to Choquet expected utility preferences over acts. Ghirardato (2001) considers preferences that are defined over acts that map coarse states into *sets* of outcomes — i.e., acts are multivalued correspondences. In this way, Ghirardato develops a model in which a sophisticated decision maker can cope with the coarseness of her perception of uncertainty: the model has a parameter that represents the agent's pessimism regarding the possible outcomes of the multivalued acts, and the overall value of an act is computed as a Choquet integral. These two works and our paper are complementary with respect to their objectives. Mukerji (1997) and Ghirardato (2001) propose very insightful models that are applicable once the analyst has fixed the way the decision maker understands uncertainty — the Γ mapping in Mukerji (1997) or the representation of natural actions of the decision maker via multivalued acts defined on some coarse state space in Ghirardato (2001). In our work, the stage of computing the value of a particular act is captured through general functions V_i, whose structure is open to further research (and via a simple expected utility-like weighted sum in the case of the second representation). Instead, we perform a revealed preference-type analysis that focuses on identifying the decision maker's subjective understanding of uncertainty — her coarse "state space" — in the first place.

Ahn and Ergin (2010) incorporate framing effects into a model of decision making under uncertainty. They are interested in comparing the agent's preferences across different frames, and their theory identifies events that are immune to framing effects — therefore transparent — and events that can be overlooked. Ahn and Ergin's (2010) setting assumes that the analyst directly observes contingencies that are relevant to the decision maker in evaluating an act; in their setting, descriptions of acts take the form of a list of contingency-outcome pairs. We, on the other hand, are interested in *learning* what the decision maker's contingencies are, and Ahn and Ergin's concept of immune events and our concept of fully understood events are not comparable. If one views decision frames more broadly, not just as partitions of the state space, then the content of this paper can be thought of as a model of the decision maker in one particular frame. If some events are described differently (without changing their meaning from the analyst's perspective), then our decision maker's conditional preferences may be different. Since we focus more on the decision maker's understanding of events, not of acts, it is possible to think of an extension for our theory that explicitly introduces a representation of events (e.g., via propositional statements)[7] and regard them as frames. Then, one can use our approach to elicit the decision maker's understanding of uncertainty for each frame and study the ways in which this understanding varies with different frames.[8]

The concept of a subjective state space goes back, at least, to Kreps (1979, 1992). In a framework of preferences over menus of alternatives, Kreps suggests that a preference for flexibility reflects uncertainty about the individual's future preferences. In his setting, the subjective states correspond exactly to possible future preferences. By refining Kreps' menu approach, Dekel, Lipman, and Rustichini (2001) deliver conditions for the uniqueness of the state space, whereas Epstein, Marinacci, and Seo (2007) relate an individual's awareness of her coarse state space to ambiguity-sensitive behavior. Nehring (1999) and Epstein and Seo (2009) provide further alternative extensions of Kreps' work. Ahn and Sarver (2013) propose a joint representation of ex ante preference for flexibility, as in Dekel et al. (2001), and ex post random choice behavior, as in Gul and Pesendorfer (2006). The connection between our work and the above papers lies in our shared concern with identifying a subjective "state space." However, our theory does not fall within the bounds of this literature because of both the methodology and the objectives. Our framework is not based on menus, and our subjective "states" do not consist of the agent's utility functions (or sets of functions). We also have a different objective, as we are interested primarily in the subjective view of the states of nature in its relationship with information processing.

The notion of subjective acts plays an important role in the work of Kochov (2015). In a setting of preferences over consumption streams, his main goal is to identify events that can be thought of as unforeseen. The gist of his notion is that an event is interpreted as

[7]See, e.g., Mukerji (1997, §3) or Lipman (1999).

[8]It is also conceivable that describing events differently may prompt the decision maker to think about possibilities that did not occur to her before and to change her awareness. Hence, one can also think about an extension that adds changing awareness to our model.

unforeseen if the agent is unable to realize that her payoff is, in fact, independent of the occurrence of this event because of a specific intertemporal structure of the stream she is facing. Kochov also proposes a transformation of acts into their subjective form that makes them measurable with respect to the algebra generated by foreseen events. Despite some terminological overlap, the main research question of our work is orthogonal to Kochov's: in particular, we focus on the choice behavior that manifests even in the absence of the time dimension. Hence, one can think of a decision maker who operates with a "state space" that is different from the analyst's, but does not exhibit choice patterns discussed in his work and for whom all events are foreseen, or the opposite, or exhibits both traits.

A noteworthy feature of our model is that the agent reveals her limited understanding of uncertainty by displaying choices that are inconsistent with Bayes' Rule. In the axiomatic decision theory literature, Epstein (2006) provides a foundation for non-Bayesian updating by positing that individuals may be tempted to change their ex ante beliefs about future states after observing the realization of a signal at an interim stage. Ortoleva (2012) generalizes Bayes' Rule to study the decision maker's reactions following the realization of events that were thought to have small (or zero) probability. These works accommodate various well-known findings, such as overreaction or underreaction to arriving information. In our model, departures from Bayes' Rule are direct consequences of the agent's coarse understanding of uncertainty.

2 Setup

Let X denote the set of *outcomes* — formally, a nonempty convex subset of a separable metric space (for instance, X could be the set of all lotteries on a set of elementary prizes). We assume that X is commonly understood by the analyst and the decision maker. From the analyst's perspective, the uncertainty is captured by a finite state space denoted by Ω. A *state-contingent act* is a function $f : \Omega \to X$; the set of all state-contingent acts is denoted by $\mathcal{F}$. Our objects of interest are ex ante and conditional preferences on $\mathcal{F}$.[9]

The binary relation $\succsim_E$ on $\mathcal{F}$ describes the agent's preferences conditional on the realiza-

[9]These are the semantics used by the analyst to describe the agent's observed behavior. One could take a syntactic approach and derive the set $\mathcal{F}$ from an abstract set of actions available to the decision maker.

tion of an event $E \subseteq \Omega$; similarly, $\succsim_\Omega$ describes the ex ante preferences before receiving any information.[10] By a preference relation we understand a binary relation that is reflexive, transitive, and complete. The ex ante preference relation $\succsim_\Omega$ is denoted simply by $\succsim$. As argued in the Introduction, our analysis relies on the assumption that the occurrence of an event does not affect the agent's understanding of uncertainty. That is, we assume that the agent acquires this information in a way that does not make her realize possible connections among facts that she did not understand before or otherwise change her perception of the world.

We also observe that defining the decision maker's preferences on $\mathcal{F}$ is not without loss of generality. Indeed, it implies that the agent must regard as interchangeable any two actions that, in the analyst's space, are represented by identical state-contingent acts. Hence, this modeling choice rules out the possibility that the agent's view of uncertainty is finer than that of the analyst. Albeit interesting, such a situation is beyond the scope of this paper.

3 General Representation of Preferences With Subjective Contingencies

3.1 Axiomatic foundations

We next turn to the set of axioms that form the foundations of our model. Since the analyst (or the modeler) and the decision maker have different perspectives on uncertainty, we have to decide whose point of view to adopt in formulating the axioms. This choice affects the way axioms are stated (and the language to be used) and interpreted. When formulated from the analyst's viewpoint, they describe properties of the choice behavior that he should expect from a decision maker who may understand the uncertainty and its resolution differently, but who is fully rational otherwise. In particular, if the decision maker were to read the axioms herself, she would not have to understand them. By contrast, when formulated from the decision maker's viewpoint, axioms may represent principles that appear reasonable from her perspective and according to which she would like to act to be consistent in her choices. This paper takes the analyst's viewpoint in laying the axiomatic

[10] Studying conditional preferences in a Bayesian setting goes back, at least, to Myerson (1986a,b).

foundation. As a result, the axioms are testable without the analyst knowing anything about the decision maker's perspective in advance.

Basic assumptions

We start with three technical conditions. The first is that the agent's preferences are continuous. Second, we require that the space of outcomes admits elements that are best and worst — i.e., the space of outcomes is order-bounded. This assumption does not play an important role in our results, but simplifies the transition from outcomes to utility levels. Third, we postulate that the uncertainty faced by the decision maker is nontrivial — i.e., there exist at least three distinct informative events.

Axiom A1 (Continuity)**.** *For any $h \in \mathcal{F}$, the sets $\{f \in \mathcal{F} : f \succsim h\}$, $\{f \in \mathcal{F} : h \succsim f\}$, $\{f \in \mathcal{F} : f \succsim_E h\}$, and $\{f \in \mathcal{F} : h \succsim_E f\}$ for all $E \subset \Omega$ are closed.*

Axiom A2 (Best and Worst Outcomes)**.** *There exist $x_*, x^* \in X$ such that $x^* \succsim f$ and $f \succsim x_*$ for all $f \in \mathcal{F}$.*

Axiom A3 (Nontriviality)**.** *There exist $A, B, C \subset \Omega$ such that $\succsim_A$, $\succsim_B$, and $\succsim_C$ are nondegenerate and differ from each other and from $\succsim$.*[11]

To conclude the list of basic axioms, we make two assumptions about consistency between ex ante and conditional preferences.

As in many models of choice under uncertainty, we seek separation between "beliefs" and "tastes," and we assume that the decision maker's tastes are not affected by information that partially resolves the uncertainty. Therefore, the ranking of constant outcomes remains unchanged for all nondegenerate conditional preferences.

Axiom A4 (Outcome Preference Consistency)**.** *For all $E \subset \Omega$ and $x, y \in X$, if $\succsim_E$ is nondegenerate, then $x \succsim y \Leftrightarrow x \succsim_E y$.*

Next, we introduce a subjective version of the dynamic consistency property. As discussed earlier, this is going to be one of the key axioms in our analysis.

To state the axiom, we recall the standard definition of a null event: for any preference relation $\succsim$, an event E is called $\succsim$-null if, for any acts f and g that differ only on E, we have $f \sim g$. Let us also introduce one piece of notation.

[11] A preference relation $\succsim$ is said to be degenerate if $f \sim g$ holds for all $f, g \in \mathcal{F}$.

Notation 1. For any $E \subseteq \Omega$ and $\omega \in \Omega$, we write $P_E(\omega)$ to denote the statement that ω is not $\succsim_E$-null; i.e., there exist $f \in \mathcal{F}$ and $x, y \in X$ such that $x\{\omega\}f \succ_E y\{\omega\}f$.[12]

Hence, $P_E(\omega)$ signifies that the state ω is revealed to remain subjectively possible after the decision maker learns that E has occurred.

The dynamic consistency property can be stated as follows.

Axiom A5 (Subjective Dynamic Consistency)**.** *For any $E \subset \Omega$ such that $\succsim_E$ is nondegenerate, and any $f, g \in \mathcal{F}$ such that $f(\omega) = g(\omega)$ for all $\omega \in \Omega$ such that $\neg P_E(\omega)$, we have $f \succsim_E g \Leftrightarrow f \succsim g$.*

The rationale behind this axiom is based on the standard argument. Suppose that the decision maker faces two acts that are identical in states that are ruled out by an event E. Then, ex ante, she should understand that her choice between these two acts matters only if E occurs. Accordingly, the axiom asserts that her ex ante choice and her choice conditional on E are the same.

This property provides an important link between information and choice that has far-reaching implications for modeling the type of behavior that we are interested in. Indeed, it implies that a message that an event E has occurred conveys nothing besides this mere fact and, therefore, does not affect the agent's understanding of uncertainty. Thus, Subjective Dynamic Consistency indicates that our agent is forward-looking, can contemplate future events at the ex ante stage, and, after learning that the event E has occurred, is always willing to carry out the parts of contingent plans that were optimal ex ante.

Note, also, that the statement of the axiom departs slightly from the commonly used ones because it takes a subjective perspective on the agent's reasoning. Instead of considering acts f and g, which coincide outside of an event E "objectively" — for all $\omega \notin E$, — our formulation applies to acts that coincide for ω's that are *revealed* to be ruled out by E (in the sense that $P_E(\omega)$ does not hold). In this way, we do not require the agent to share the analyst's view on whether two given acts may yield different payoffs only if E occurs, and, at the same time, we are able to keep the statement of the axiom fully testable.

[12]As usual, for any $x \in X, \omega \in \Omega$, and $f \in \mathcal{F}$, we denote by $x\{\omega\}f$ the act in $\mathcal{F}$ yielding the outcome x in the state ω, and $f(\omega')$ in all states $\omega' \neq \omega$.

Information-processing axioms

The second group of axioms is concerned with how the decision maker processes information about events that occur.

Consistent with our objective of modeling a decision maker who has no difficulty understanding the descriptions of events but may not understand fully how they resolve the payoff-relevant uncertainty, we assume that larger events imply wider sets of possibilities.

Axiom A6 (Monotonicity of the Possibility Predicate). *For any $A \subseteq B \subseteq \Omega$ and any $\omega \in \Omega$, $P_A(\omega) \Rightarrow P_B(\omega)$.*

Next, we assume that if two events are revealed to be equivalent in terms of what is subjectively possible if they occur, then the conditional preferences on these two events have to be identical.

Axiom A7 (Equivalence of Events). *Suppose that $E, E' \subseteq \Omega$ are related so that*

(i) for any $\omega \in E' \backslash E$, $P_E(\omega)$ holds, and

(ii) for any $\omega \in E \backslash E'$, $P_E(\omega)$ does not hold.

Then, $\succsim_{E'} = \succsim_E$.

In words, if two events are such that (i) states that E' contains in addition to E are already deemed possible after E, and (ii) states that E' misses relative to E are already deemed to be impossible after E, then these two events generate identical conditional preferences. Compared to the previous axiom, this axiom maintains a stronger connection between two events and the respective conditional preferences — they must rank all acts in exactly the same way.

To assess the type of behavior allowed by this axiom, it is useful to examine its content in the standard case of subjective expected utility preferences, full understanding of uncertainty, and Bayesian updating. In this case, ex ante and conditional preferences are represented by

$$\begin{cases} V(f) = \sum_{\omega \in \Omega} u(f(\omega))\, p(\omega), \\ V(f \mid E) = \sum_{\omega \in E} u(f(\omega))\, p(\omega) \quad \text{for } E \subset \Omega, \end{cases} \tag{1}$$

where p is a probability measure on Ω, and u is a utility function over outcomes. Hence, in Part (ii), if $\omega \in E$ but $P_E(\omega)$ does not hold, then $p(\omega) = 0$. In this case, the decision maker's preferences conditional on E' are the same as those conditional on E. Part (i) in the standard case becomes vacuous, because the combination of conditions $\omega \notin E$ and $P_E(\omega)$ in the antecedent is impossible.

It remains to impose minimal requirements on the decision maker's understanding of intersections and unions of events.

Axiom A8 (Understanding of Relationships Between Events)**.** *Suppose that $A \subseteq \Omega$ is such that $P_A(\omega) \Leftrightarrow \omega \in A$ for all $\omega \in \Omega$, and $B \subseteq \Omega$. Then,*

(i) If there exists $\omega_0 \in \Omega$ such that $P_A(\omega_0)$ and $P_B(\omega_0)$ hold, then $P_{A\cap B}(\omega_0)$ must also hold.

(ii) If there exists $\omega_0 \in \Omega$ such that $P_{A\cup B}(\omega_0)$ holds but $P_A(\omega_0)$ does not, then $P_B(\omega_0)$ must hold.

The axiom is intuitive. Yet, it is not implied by anything else that we have postulated, and imposes a certain discipline on the decision maker's understanding of what is possible in various situations. The axiom also contains one non-trivial qualification that A is such that $P_A(\omega)$ holds if and only if $\omega \in A$. It captures what we will later call a full understanding of the event A and, hence, restricts the applicability of the axiom to situations in which the analyst observes that the decision maker completely understands at least one of the two events in question.

Finally, we impose one more assumption regarding the decision maker's reaction to events that are deemed impossible. For expositional simplicity, we have already assumed that if the decision maker is asked to rank acts conditional on an event that, ex ante, she viewed as impossible, then she reports indifference between any two acts — i.e., her preferences become degenerate. This assumption is implicitly built into the statements of some of the earlier axioms — for instance, in the form of the qualification "for any E such that $\succsim_E$ is nondegenerate."[13] Now, we postulate that the decision maker regards as impossible events whose descriptions are *objectively* self-contradicting and that are represented by the empty set.

Axiom A9 (Indifference Upon Impossible Events)**.** *For all $f, g \in \mathcal{F}$, $f \sim_\varnothing g$.*[14]

[13]For our theory at large, it is not strictly necessary to make assumptions about the decision maker's preferences after occurrence of events that were regarded as impossible at the ex ante stage. An earlier version of this paper proceeded without any such assumptions — at the expense of making the setup and the representation theorems a bit heavier.

[14]This assumption is also not strictly necessary. Moreover, unlike previous axioms, this one can be tested only by asking the decision maker hypothetical questions, but not by observing her choice conditional on

3.2 The representation

The following definition will be useful for stating our representation result.

Definition 2. We say that a set $S \subseteq \Omega$ is the *support* of a function $F : \mathcal{F} \to \mathbb{R}$ if

(i) $F(f) = F(g)$ for all $f, g \in \mathcal{F}$ such that $f|_S = g|_S$; and

(ii) for each $\omega \in S$, there exist $f \in \mathcal{F}$ and $x, y \in X$ such that $F(x\{\omega\}f) \neq F(y\{\omega\}f)$.

Now, we are ready to state our main theorem, which follows in an if-and-only-if manner from the axioms stated earlier.

Theorem 1. *An ex ante preference relation $\succsim$ and a collection $\{\succsim_E\}_{E \subset \Omega}$ of conditional preference relations jointly satisfy Axioms (A1)–(A9) if and only if there exist*

- *a set of indices (subjective states) $\mathcal{S} = \{1, \ldots, n\}$ for some $n \in \mathbb{N}$ with $n \geq 3$;*
- *a collection $\Pi = \{C_1, \ldots, C_n\}$ of nonempty and mutually disjoint subsets of Ω;*
- *a collection of nonconstant continuous functions $V_i : \mathcal{F} \to \mathbb{R}$ for $i \in \mathcal{S}$, with a compact range, and such that C_i is the support of V_i for all $i \in \mathcal{S}$; and*
- *a collection of functions $\sigma_i : 2^{C_i} \to \{0, 1\}$ for $i \in \mathcal{S}$, such that, for all $i \in \mathcal{S}$, $\sigma_i(\varnothing) = 0$, $\sigma_i(C_i) = 1$, and $\sigma_i(A) \geq \sigma_i(B)$ for all $A, B \subseteq C_i$ satisfying $A \supseteq B$*

such that:

(i) $\succsim$ is represented by

$$V(f) = \sum_{i \in \mathcal{S}} V_i(f) \quad \forall f \in \mathcal{F}; \tag{2}$$

(ii) For each $E \subset \Omega$, $\succsim_E$ is represented by

$$V(f \mid E) = \sum_{i \in \mathcal{S} : \sigma_i(E \cap C_i) = 1} V_i(f) \quad \forall f \in \mathcal{F}; \tag{3}$$

(iii) Functions $(V_i)_{i \in \mathcal{S}}$ agree on ranking sure outcomes

$$V_i(x) \geq V_i(y) \Leftrightarrow V_j(x) \geq V_j(y) \quad \forall i, j \in \mathcal{S} \text{ and } x, y \in X. \tag{4}$$

The main message of the theorem can be summarized as follows.

1. The decision maker's view of uncertainty is captured by a set $\mathcal{S}$ of contingencies. Each contingency $i \in \mathcal{S}$ identifies a nonempty subset C_i of Ω, a function $V_i : \mathcal{F} \to \mathbb{R}$

the actual realization of the event. Dropping this axiom, however, requires extending the statements of subsequent theorems, and we do not regard an additional generalization in that direction as worthwhile.

specifying the utility value of an act on contingency i, and a function $\sigma_i : 2^{C_i} \to \{0,1\}$ determining whether contingency i is consistent with an event E.

2. Conditional on the realization of an event E, the decision maker identifies a subset of contingencies that she perceives to be consistent with the received information, and she updates her preferences in a Bayesian-like manner — this is captured by the relationship between Equations (3) and (2). Observe that, in the simplest case in which $E = C_i$, we have $V(f \mid C_i) = V_i(f)$ for $i \in \mathcal{S}$, and, therefore, Representation (3) implies that $V(f) = \sum_{i \in \mathcal{S}} V(f \mid C_i)$.
3. When evaluating acts, the decision maker's preferences are additively separable with respect to her subjective contingencies — this is captured by Equation (2).

As emphasized in the Introduction, the key feature of the contingencies is that they are *subjective* — i.e., the decision maker is free to come up with her own set of contingencies that reflect her understanding of what is possible. Her choice behavior reveals such contingencies. In particular, our proof of the theorem in the Appendix is based on a constructive procedure; hence, the analyst can, in principle, follow its steps to recover the decision maker's subjective contingencies by asking her to rank sufficiently many acts. Anticipating the result in Subsection 3.3, we also note that the set of contingencies is identified uniquely.

Another important property of the subjective contingencies is that they correspond to *sets* of the analyst's states — the cells of Π — and, as such, are coarse. This coarseness has no notable implications for the agent's ex ante preferences in isolation. Rather, it manifests in the way preferences are updated; for instance, it may result in overlooking or misclassifying certain scenarios, as illustrated in the investment example from the Introduction. However, within the bounds of her understanding, the agent does her best to process information. Namely, she determines which contingencies remain a possibility and which are ruled out; then, similar to the standard Bayesian updating recollected in (1), she computes the conditional value of an act by dropping the additive terms referring to contingencies that are no longer relevant.

The novelties in our updating rule are related to the fact that the decision maker's contingencies are coarse. If an event $E \subseteq \Omega$ is measurable with respect to the algebra generated by the partition Π, then the updating rule implied by Theorem 1 takes a rather conventional form: a contingency C_i disappears from the representation of conditional preferences if $C_i \cap E = \varnothing$, and stays otherwise. However, if E is not measurable, then the conditional preference relation $\succsim_E$ coincides with the conditional preference relation $\succsim_{\tilde{E}}$, where $\tilde{E}$ is a measurable event that is uniquely determined by E.[15] Specifically, if C_i is entirely contained in E, then it is also contained in $\tilde{E}$; symmetrically, if C_i does not intersect E, then it does not intersect $\tilde{E}$. If $\varnothing \neq C_i \cap E \subset C_i$, then both situations $C_i \cap \tilde{E} = C_i$ and $C_i \cap \tilde{E} = \varnothing$ are possible and depend on the value of σ_i. In particular, one can envision an agent for whom $C_i \cap \tilde{E} = C_i$ most of the time (or always), or for whom $C_i \cap \tilde{E} = \varnothing$ most of the time (or always). Section 4 also presents a refined model in which the question about whether a contingency C_i is consistent with the event E is answered depending on the total mass of points in $C_i \cap E$.

The additive separability of representations (2)–(3) is noteworthy because this kind of structure is a precursor to defining the agent's *beliefs* about contingencies. The representation clearly determines what the decision maker regards as possible or impossible. In particular, the states outside of the union of the cells of Π are null and have no impact on the evaluation of acts. (In loose terms, the possibilities that are not covered by Π are viewed by the decision maker as having zero probability.) However, due to its generality, Theorem 1 is silent about the relative likelihood of contingencies in $\mathcal{S}$. One simple reason is that our functions V_i for $i \in \mathcal{S}$ are not normalized. However, the likelihoods of the contingencies can be pinned down by imposing more structure on the ex ante preferences, and one particular refinement of the model of Theorem 1 is provided in Section 4. We also note that the standard Bayesian expected utility maximization is a special case of our model, as elaborated in Subsection 3.4.

Examples

Next, we revisit the investment example from the Introduction to illustrate how our representation works and how the described choice behavior translates into a concrete specification of the objects listed in the statement of the theorem.

Example 1. *Recall that in the discussed example, the analyst's state space is* $\Omega = \{\omega_g, \omega_m, \omega_b\}$, *where* ω_g *represents the situation of strong demand and expansion, in which a firm pays the*

[15] This is the sense in which the decision maker's updating process is limited by her understanding of the world.

dividend of \$4 per share; ω_m *represents the situation of strong demand but no expansion, in which the firm pays \$2; and* ω_b *represents the situation of weak demand, in which the firm pays \$1. The decision maker is observed to rank* $(1,0,0) \succ (0,0,0)$ *and* $(4,2,1) \sim_{E^{low}} (0,0,1)$ *for* $E^{low} = \{\omega_g, \omega_b\}$*. As implied by intermediate results that we establish in the course of proving the theorem, the collection of all non-null states of the conditional preference relation* $\succsim_E$ *for any* $E \subseteq \Omega$ *is a union of elements of the partition* Π*. Hence, it must be that* $\{\omega_b\}$ *is an element of* Π*. Clearly,* $\{\omega_g\}$ *is not an element of* Π*, and, hence, if all states are ex ante non-null, it must be that* $\{\omega_g, \omega_m\}$ *is an element of* Π*. It follows that our story is consistent with the decision maker having two subjective "states":* $\Pi = \{C_1, C_2\}$ *where* $C_1 = \{\omega_b\}$ *and* $C_2 = \{\omega_g, \omega_m\}$*. Then,* σ_1 *is defined trivially by* $\sigma_1(\varnothing) = 0$ *and* $\sigma_1(\{\omega_b\}) = 1$*. The functions* V_1 *and* V_2 *should be defined such that* $V_2(1,0,0) > V_2(0,0,0)$*, their supports are* C_1 *and* C_2*, respectively, and they agree as in (4).*[16]

Regarding σ_2*, the decision maker's choices mentioned above are insufficient to pin it down uniquely. One example of the specification of* σ_2 *that does the job is* $\sigma_2(\varnothing) = 0$*,* $\sigma_2(\{\omega_g\}) = 0$*,* $\sigma_2(\{\omega_m\}) = 1$*,* $\sigma_2(\{\omega_g, \omega_m\}) = 1$*. Note that, under this specification, it must be that* $(1,0,0) \succ_{E^{high}} (0,0,0)$*, where* $E^{high} = \{\omega_m\}$ *is the event of the mid-year report of high interim profit. We also see that* $\succsim_{E^{high}} = \succsim_{\tilde{E}}$*, where* $\tilde{E} = \{\omega_g, \omega_m\} = C_2$ *is the event that the demand is strong — in the proposed specification, the decision maker subjectively views the events "high interim profit" and "strong demand" as equivalent. On a separate note, observe that, regardless of the specification of* σ_2*, it must be that* $(1,0,0) \succ_{C_2} (0,0,0)$*.*

This example also illustrates another important aspect of the differences between the decision maker's and the analyst's understanding of the world. Notice that the specification $\sigma_2(\{\omega_g\}) = 0$ implies that the event $\{\omega_g\}$ ("strong demand but low interim profit") is subjectively impossible — the conditional preference relation $\succsim_{\{\omega_g\}}$ is degenerate. At the same time, the decision maker is willing to pay a positive price for the vector $(1,0,0)$ that promises a positive payoff in state ω_g only. This may look like an inconsistency in the decision maker's behavior. However, note that it is the analyst's state space that provides a direct link between the case of strong demand plus low interim profit and the dividend

[16] For instance, there can be a von Neumann-Morgenstern utility function u over monetary lotteries, $V_1(z_g, z_m, z_b) = u(z_b)$, and $V_2(z_g, z_m, z_b) = \tilde{V}(u(z_g), u(z_m))$, where $\tilde{V} : \mathbb{R}^2 \to \mathbb{R}$ is some strictly monotone function. $\tilde{V}$ can be linear, as in the expected utility theory, or nonlinear, as in the theory of choice under ambiguity.

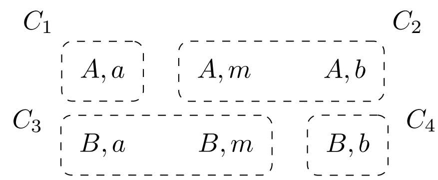

Figure 1.— Subjective contingencies in the Confirmatory Bias example

payment of \$4. If the decision maker understands the uncertainty differently and operates with a different set of contingencies, then this link ceases to be self-evident. This situation is typical in our model. While our decision maker may make inferences that look erroneous to the analyst, it is not a sign of mistakes on the logical or computational levels. Rather, those errors stem from the decision maker's cognitive state (e.g., what she knows about the firm's business) and may be viewed as parallel to Type-I and Type-II errors in mathematical statistics.

The next example illustrates the model in application to the widely documented Confirmatory Bias — the tendency to view arriving information through the lens of one's primary hypothesis, overweighting the news that supports it, and underweighting (or dismissing) the news that does not.[17]

Example 2. *Let A and B be two exhaustive and mutually exclusive hypotheses, and assume that the decision maker receives a signal $s \in \{a, m, b\}$ that provides some noisy information about the true hypothesis and has the following conditional probabilities:*

	$x = a$	$x = m$	$x = b$
$\Pr(s = x \mid A)$	0.5	0.3	0.2
$\Pr(s = x \mid B)$	0.2	0.3	0.5

Given those probabilities, signal $s = a$ is noisy evidence in favor of A; signal $s = b$ is noisy evidence in favor of B; and $s = m$ represents mixed evidence. Our aim is to model an agent who has a bias towards B in interpreting the signal.

[17]The primary hypothesis can be selected according to its likelihood or because the person hopes, or fears, that it is true. For an early psychological experiment, see, e.g., Wason (1960). In the classical experiment of Darley and Gross (1983), subjects viewed the same tape of a schoolgirl doing an achievement test but formed quite different opinions about her performance depending on the cues about her socioeconomic background. In economics, see, e.g., the model developed by Rabin and Schrag (1999).

Let the state space be $\Omega = \{A,B\} \times \{a,m,b\}$, *and the decision maker's subjective contingencies be represented by* $C_1 = \{A\} \times \{a\}$, $C_2 = \{A\} \times \{m,b\}$, $C_3 = \{B\} \times \{a,m\}$, *and* $C_4 = \{B\} \times \{b\}$, *as illustrated in Figure 1. In words, she subjectively sees only four possibilities:* C_1 *— the true hypothesis is* A, *and the evidence is* $s = a$ *(suggestive of* A*);* C_2 *— the true hypothesis is* A, *but the evidence is* $s \neq a$ *(not suggestive of* A*);* C_3 *— the true hypothesis is* B, *but the evidence is* $s \neq b$ *(not suggestive of* B*); and* C_4 *— the true hypothesis is* B, *and the evidence is* $s = b$ *(suggestive of* B*). Moreover, suppose that* $\sigma_2(\{A\} \times \{m\}) = \sigma_2(\{A\} \times \{b\}) = 0$ *and* $\sigma_3(\{B\} \times \{a\}) = \sigma_3(\{B\} \times \{m\}) = 1$, *whereas* σ_1 *and* σ_4 *are defined trivially. Note that the events "*A *holds" and "*B *holds" are represented by unions of some cells of the partition* $\{C_1, C_2, C_3, C_4\}$, *so, intuitively, our decision maker treats them in the same way as a Bayesian agent.*[18] *However, the same is not true for the event "the signal is* b*." Also note that the construction of* σ_2 *and* σ_3 *is not symmetric with respect to exchanging* A *and* B *— this asymmetry is what will create a bias towards* B.

For the sake of the example, assume that the functions V_i*'s are linear — that is, they have the Subjective Expected Utility form with a risk-neutral utility and some subjective probability* p *on* Ω.[19] *For convenience, we express* p *as* $p = \mu\, p_i$, *where* μ *is a probability measure over* $\mathcal{S}$ *and* p_i *is a probability measure over* C_i, *for all* $i \in \mathcal{S}$. *Hence,* $V_i(f) = \mu(C_i) \sum_{\omega \in C_i} f(\omega) p_i(\omega)$ *for all* $i \in \mathcal{S}$. *Assume that the ex ante probability that* B *holds is* 0.6. *In accordance with this and the specification of the signals, we let* μ *and* p_i *take the following values:* $\mu = (0.2, 0.2, 0.3, 0.3)$, $p_2 = (0.6, 0.4)$, *and* $p_3 = (0.4, 0.6)$.

Then, let f *be an act that pays* 1 *if* B *and* 0 *if* A, *and let* h *be an act that pays* 1 *always. Using* h *as the measuring rod, we can evaluate what the agent thinks about the likelihood of* B *after receiving some information. That is, we can compute the conditional values of* f *as follows:*

$$\begin{aligned}
V(f \mid s = a) &= \mu(C_3) &&= 0.6\,[\mu(C_1) + \mu(C_3)] &&= 0.6\,V(h \mid s = a),\\
V(f \mid s \in \{m,b\}) &= \mu(C_3) + \mu(C_4) &&= 0.75\,[\mu(C_2) + \mu(C_3) + \mu(C_4)] &&= 0.75\,V(h \mid s \in \{m,b\}),\\
V(f \mid s = b) &= \mu(C_4) && &&= 1\,V(h \mid s = b).
\end{aligned}$$

In all three cases, the decision maker's evaluation of the "bet" on B *is higher than what*

[18] We will discuss the relevant formal concept in Section 3.4.

[19] The case of the Subjective Expected Utility form is characterized in Theorem 6, whereas Theorem 1 focuses on the general concept of subjective uncertainty that is independent of the functional forms taken by V_i and σ_i.

it would be if conditional probabilities were computed by the standard Bayesian decision maker (0.375*,* ≈ 0.7059*,* ≈ 0.7895*, respectively). Moreover, if* $s = a$*, the decision maker's "posterior" regarding* B *equals her prior, despite the evidence against* B*; so, she acts as if she ignores this evidence. If* $s = b$*, our decision maker's bias takes the extreme form of viewing the signal as convincing evidence that her favorite hypothesis is true. Generally, her behavior can be interpreted as if she applies higher standards to evidence against* B *than to evidence against* A*, despite the symmetric conditional probabilities of the signals.*

The main driving force in the above observations is that the decision maker regards mixed or conflicting signals as consistent with B*, more than she does with* A*. The subjective treatment of the evidence is captured by the structure of functions* σ_2 *and* σ_3 *and is enabled by the subjective (and coarse) understanding of uncertainty in the first place.*

The ability of our model to generate the Confirmatory Bias is not a key finding of the paper. Rather, it is an illustration that our work addresses not only theoretical concerns about the Savagean common state space assumption as outlined in the Introduction, but can also accommodate behavioral patterns that are well-documented empirically. Presumably, there may well be other biases in information processing that our theory can rationalize and for which the representation of Theorem 1 provides a tractable way of modeling. As an illustration, consider the so-called Conservatism Bias, which is the tendency to give a lower weight to new information and to overweight prior beliefs.[20] The Conservatism Bias can be obtained in a setup similar to that in Example 2 if the agent's contingencies are $C_1 = \{A\} \times \{a\}$, $C_2 = \{A\} \times \{m, b\}$, $C_3 = \{B\} \times \{a\}$, $C_4 = \{B\} \times \{m, b\}$. Furthermore, our theory can be applied to phenomena such as the Winner's Curse.

3.3 Uniqueness

We turn to the question of uniqueness and show that the representation result in Theorem 1 has strong properties in this regard. This feature is particularly relevant since the question of uniqueness of subjective state spaces has drawn substantial attention in different contexts of the decision theory literature.

[20]This bias is discussed in Tversky and Kahneman (1974); for recent experimental evidence, see Gneezy, Hoffman, Lane, List, Livingston, and Seilier (2016).

Definition 3. We refer to a quadruple $(\mathcal{S}, (C_i)_{i\in\mathcal{S}}, (V_i)_{i\in\mathcal{S}}, (\sigma_i)_{i\in\mathcal{S}})$ that satisfies the conditions of Theorem 1 as a *representation with subjective contingencies.*

Proposition 2. *Suppose that* $(\mathcal{S}, (C_i)_{i\in\mathcal{S}}, (V_i)_{i\in\mathcal{S}}, (\sigma_i)_{i\in\mathcal{S}})$ *and* $(\mathcal{S}', (C_i')_{i\in\mathcal{S}'}, (V_i')_{i\in\mathcal{S}'}, (\sigma_i')_{i\in\mathcal{S}'})$ *are two representations with subjective contingencies. Then, they represent the same system* $(\succsim, \{\succsim_E\}_{E\subset\Omega})$ *of preference relations if and only if there exists a bijection* $\pi : \mathcal{S} \to \mathcal{S}'$ *such that*

(i) $C'_{\pi_i} = C_i$ *for all* $i \in \mathcal{S}$*;*

(ii) there exist $\alpha > 0$ *and* $\beta_i \in \mathbb{R}$ *for* $i \in \mathcal{S}$*, such that* $V'_{\pi_i} = \alpha V_i + \beta_i$ *for all* $i \in \mathcal{S}$*; and*

(iii) $\sigma'_{\pi_i} = \sigma_i$ *for all* $i \in \mathcal{S}$*.*

In words: First, the set of subjective contingencies and the corresponding partition of the state space are identified uniquely. Second, the model uniquely identifies the rules to determine contingencies that are consistent with an event. Finally, the model identifies, up to a joint positive affine transformation, the collection of evaluation functionals — the mappings from an act to the utility level that obtains from this act in each contingency.

3.4 Comparative statics

3.4.1 Understanding of uncertainty

Next, we discuss the ways in which our decision makers can be compared and the corresponding comparative statics in the representation of their preferences. The starting point of this exercise is a comparison in terms of their understanding of the nature of uncertainty — i.e., their perception of the state space.

Definition 4. A decision maker described by $(\succsim, \{\succsim_A\}_{A\subset\Omega})$ is said to *fully understand* an event $E \subseteq \Omega$ if

$$P_E(\omega) \quad \Leftrightarrow \quad P_\Omega(\omega) \text{ and } \omega \in E \qquad \text{for all } \omega \in \Omega.$$

Definition 5. One decision maker has a *finer understanding of uncertainty* than another if the first one fully understands all events $E \subseteq \Omega$ that the second one fully understands. In particular, two decision makers have *equal understanding of uncertainty* if the first one fully understands E if and only if the second one fully understands E, for any $E \subseteq \Omega$.

Now, we characterize our concept of full understanding of an event and the corresponding comparative notion.

Proposition 3. *Suppose that* $(\mathcal{S}, \Pi, (V_i)_{i\in\mathcal{S}}, (\sigma_i)_{i\in\mathcal{S}})$ *is a representation with subjective contingencies, and* $\succsim$ *is the ex ante preference relation that it represents. Then, the decision maker fully understands* $E \subseteq \Omega$ *if and only if* E *is the union of some elements of the collection* $\Pi \cup \mathcal{N}(\succsim)$*, where* $\mathcal{N}(\succsim) := \{\{\omega\} \mid \omega \in \Omega \text{ is } \succsim\text{-null}\}$*.*

Note that $\Pi \cup \mathcal{N}(\succsim)$ is a partition of Ω. Thus, an event is fully understood if and only if it is measurable with respect to the algebra generated by this partition.

Corollary 4. *Suppose that* $(\mathcal{S}^k, \Pi^k, (V_i^k)_{i\in\mathcal{S}^k}, (\sigma_i^k)_{i\in\mathcal{S}^k})$ *for* $k = 1, 2$ *are two representations with subjective contingencies, and* $\succsim^1$ *and* $\succsim^2$ *are the ex ante preference relations that they represent. Then, the first decision maker has a finer understanding of uncertainty if and only if* $\Pi^1 \cup \mathcal{N}(\succsim^1)$ *is a refinement of* $\Pi^2 \cup \mathcal{N}(\succsim^2)$*.*

A decision maker has the maximal understanding of uncertainty if her $\Pi \cup \mathcal{N}(\succsim)$ consists only of singleton sets. A full understanding of all events (i.e., maximal understanding of uncertainty) is intrinsic to standard economic agents who have expected utility preferences and perform Bayesian updating if they are viewed through the lens of our model. Indeed, suppose that the utility index of an expected utility maximizer is $u : X \to \mathbb{R}$, her prior is $p \in \Delta(\Omega)$, and Ω is enumerated as $\{\omega_1, \ldots, \omega_n, \ldots, \omega_{n+m}\}$ such that $p(\omega_l) > 0$ for all $l = 1, \ldots, n$ and $p(\{\omega_{n+1}, \ldots, \omega_{n+m}\}) = 0$. Then, we can set $\mathcal{S} = \{1, \ldots, n\}$, $\Pi = \{\{\omega_i\} \mid i \in \mathcal{S}\}$, $V_i(f) = u(f(\omega_i))\, p(\omega_i)$, and our representation (2)–(3) reduces to the standard procedure (1).

3.4.2 Updating errors

As noted in Section 3.1, our decision maker makes inferences from arriving information and updates her preferences so that individual states may be treated peculiarly. Indeed, conditional on an event E, some state $\omega \notin E$ may remain non-null, or, conversely, some ex-ante-non-null state $\omega' \in E$ may become null after updating. Thus, from the analyst's point of view, the agent makes inference errors. We refer to the first type of spurious inference as an *inclusion error* and to the second type as an *exclusion error*. These errors emerge because the agent has a limited understanding of uncertainty — in a similar fashion, a statistician makes type-I or type-II errors in hypothesis testing because he is not omniscient and makes inferences on the basis of finite and imperfect data.

The notation below and Definition 7 provide a formal statement of a comparative notion that captures situations in which one agent is more prone to exclusion (resp. inclusion) errors than another.

Notation 6. Let $(\succsim, \{\succsim_A\}_{A\subset\Omega})$ be a system of preference relations and $E \subseteq \Omega$. We define

$$M_-(E) := \{\omega \in E : P_\Omega(\omega) \text{ and } \neg P_E(\omega) \text{ hold}\},$$
$$M_+(E) := \{\omega \notin E : P_E(\omega) \text{ holds}\}.$$

Note that if sets $M_-(E)$ or $M_+(E)$ are nonempty for some event E, then the decision maker does not understand this event fully, and the two sets represent the states with respect to which she makes exclusion and inclusion errors, respectively.

Definition 7. Let two decision makers be described by $(\succsim^k, \{\succsim^k_A\}_{A\subset\Omega})$ for $k = 1, 2$.

(i) Decision Maker 1 is *more prone to exclusion errors* than Decision Maker 2 if, for any $E \subseteq \Omega$ such that $\succsim^1_E$ and $\succsim^2_E$ are nondegenerate, we have $M^2_-(E) \subseteq M^1_-(E)$.

(ii) Decision Maker 1 is *more prone to inclusion errors* than Decision Maker 2 if, for any $E \subseteq \Omega$ such that $\succsim^1_E$ and $\succsim^2_E$ are nondegenerate, we have $M^2_+(E) \subseteq M^1_+(E)$.

The subsequent discussion relies heavily on the following notion.

Notation 8. Let $(\mathcal{S}, (C_i)_{i\in\mathcal{S}}, (V_i)_{i\in\mathcal{S}}, (\sigma_i)_{i\in\mathcal{S}})$ be a representation with subjective contingencies. Then, we refer to

$$\mathcal{S}_c := \{i \in \mathcal{S} : |C_i| \geq 2\}$$

as the decision maker's *coarse subjective contingencies*.

The term "coarse" refers to the fact that the corresponding cells of the partition representing contingencies are coarser than singletons. This fact has a direct connection to imperfections in the agent's understanding of uncertainty. Indeed, any set outside of $\bigcup_{i\in\mathcal{S}_c} C_i$ must be fully understood, and, for any $E \subseteq \Omega$, the sets $M_-(E)$ and $M_+(E)$ must be contained in $\bigcup_{i\in\mathcal{S}_c} C_i$.

For the purpose of linking the comparative notions introduced earlier to the parameters of our representation — namely, the consistency maps $(\sigma_i)_{i\in\mathcal{S}}$ — we focus on the case in which two decision makers have equal understanding of uncertainty. Note that, as Corollary 4 implies, two such decision makers have the same (up to renumbering) sets of coarse subjective contingencies — that is, if $(\mathcal{S}^k, \Pi^k, (V_i^k)_{i\in\mathcal{S}^k}, (\sigma_i^k)_{i\in\mathcal{S}^k})$ for $k = 1, 2$ are the representations with subjective contingencies of the preferences of two agents with equal understanding of uncertainty, then there must exist a bijection $\pi : \mathcal{S}_c^1 \to \mathcal{S}_c^2$ such that $C_{\pi_i}^2 = C_i^1$ for all $i \in S_c^1$.

Now, we can state the comparative statics result.

Proposition 5. *Suppose that $(\mathcal{S}^k, (C_i^k)_{i\in\mathcal{S}^k}, (V_i^k)_{i\in\mathcal{S}^k}, (\sigma_i^k)_{i\in\mathcal{S}^k})$ for $k = 1, 2$ are two representations with subjective contingencies, and the decision makers that they represent have equal understanding of uncertainty. Furthermore, suppose that the set of coarse subjective contingencies for both of them is $\{1, \ldots, m\}$, and let $C_i := C_i^1 = C_i^2$ for $i = 1, \ldots, m$. Then, the following conditions are equivalent:*

(i) Decision Maker 1 is more prone to exclusion errors than Decision Maker 2;

(ii) Decision Maker 2 is more prone to inclusion errors than Decision Maker 1;

(iii) For all $i = 1, \ldots, m$, we have $\sigma_i^1(A) \le \sigma_i^2(A)$ for all $A \subseteq C_i$.

This result claims, first, that the tendency to make exclusion versus inclusion errors are two faces of the same trait: if one agent is more inclined than the other to make exclusion errors, then she must, at the same time, be less inclined to make inclusion errors. Second, in terms of the representation of their preferences, the fact that one agent is more prone to make inclusion errors (and less prone to make exclusion errors) than another is captured by the pointwise dominance relationship between their consistency maps that correspond to the agents' common coarse contingencies.

4 A More Structural Representation

While Theorem 1 captures the essential aspects of our theory, it leaves open the question about reasonable specifications of the evaluation functions V_i and consistency functions σ_i for contingencies $i \in \mathcal{S}$. In this section, we address this point by imposing a few additional axioms — Independence, Monotonicity, and one novel informational axiom — and deriving a more structural representation. This result preserves the gist of Theorem 1 about subjective states and updating and, at the same time, provides a familiar expected utility-like structure for functions V_i and a related weight-based specification for σ_i.

4.1 Axioms

We start by recalling well-known properties that we impose on the ex ante preferences.

Axiom A10 (Monotonicity). *If $f, g \in \mathcal{F}$ and $f(\omega) \succsim g(\omega)$ for all $\omega \in \Omega$, then $f \succsim g$.*

Axiom A11 (Independence). *If $f, g, h \in \mathcal{F}$ and $\alpha \in (0,1)$, then*

$$f \succsim g \quad \Leftrightarrow \quad \alpha f + (1-\alpha)h \succsim \alpha g + (1-\alpha)h.$$

Axioms (A10) and (A11) are the classic assumptions paving the way for the subjective expected utility theory of Anscombe and Aumann (1963).

In addition, we introduce an axiom that provides a link between the decision maker's attitude towards betting on an event and her ability to evaluate acts conditional on that event.

Axiom A12 (Admissibility of Valuable Events). *Let $x, y \in X$ such that $x \succ y$, $A, B \subseteq \Omega$ such that $\succsim_A$ is nondegenerate, and $P_A(\omega)$ for all $\omega \in B$. If $xBy \succsim xAy$, then $\succsim_B$ is nondegenerate.*

The essence of the axiom is that if $\succsim_A$ is nondegenerate and a "bet" on an event B is ex ante preferred to a bet on A, then B cannot be void and $\succsim_B$ should be nondegenerate, too. In the standard theory of Bayesian updating, in which admissibility of an event equals having positive probability, this implication always holds. In our theory, the axiom contains an additional qualification: we assume that any state in B is revealed to be possible after the event A has occurred, which implies that A and B are comparable in their informational content.

4.2 The representation

The extended list of axioms leads to the following representation.

Theorem 6. *A system $(\succsim, \{\succsim_A\}_{A\subset\Omega})$ of preference relations satisfies Axioms (A1)–(A12) if and only if there exist*

- *a set of indices (subjective states) $\mathcal{S} = \{1, \dots, n\}$ for some $n \in \mathbb{N}$ with $n \geq 3$;*
- *a collection $\Pi = \{C_1, \dots, C_n\}$ of nonempty and mutually disjoint subsets of Ω;*
- *a nonconstant, continuous, and affine function $u : X \to \mathbb{R}$ with a compact range;*

- *a probability measure* $\mu \in \Delta(\mathcal{S})$*;*
- *a collection of probability measures* $p_i \in \Delta(\Omega)$ *for* $i \in \mathcal{S}$ *such that* C_i *is the support of* p_i*; and*
- *a collection of numbers* $\alpha_i \in (0,1]$ *for* $i \in \mathcal{S}$

such that:

(i) For all $f, g \in \mathcal{F}$*,*

$$f \gtrsim g \quad \Leftrightarrow \quad \sum_{i \in \mathcal{S}} u(f^*(i))\,\mu_i \geq \sum_{i \in \mathcal{S}} u(g^*(i))\,\mu_i, \tag{5}$$

where, for each $f \in \mathcal{F}$*,* $f^* : \mathcal{S} \to X$ *is defined as*

$$f^*(i) = \sum_{\omega \in \Omega} p_i(\omega) f(\omega) \qquad \forall_{i \in \mathcal{S}}; \tag{6}$$

(ii) For all $f, g \in \mathcal{F}$ *and* $E \subset \Omega$*,*

$$f \gtrsim_E g \quad \Leftrightarrow \quad \sum_{i \in \mathcal{S}(E)} u(f^*(i))\,\mu_{i|E} \geq \sum_{i \in \mathcal{S}(E)} u(g^*(i))\,\mu_{i|E}, \tag{7}$$

where

$$\mu_{i|E} = \frac{\mu_i}{\sum_{j \in \mathcal{S}(E)} \mu_j} \tag{8}$$

and

$$\mathcal{S}(E) := \left\{ i \in \mathcal{S} : \sum_{\omega \in E} p_i(\omega) \geq \alpha_i \right\}. \tag{9}$$

In comparison to the analysis of Section 3, the first feature of this theorem consists of refining the general additively separable representation of Theorem 1 to make subjective contingencies have *probabilities* (captured by $\mu \in \Delta(S)$).

The second feature is related to the much more specific procedure of computing the value of an act on a single contingency. The decision maker can be viewed as operating with *subjective transformations* of acts, denoted above by f^* and g^*, that map contingencies to outcomes. These transformations are represented as weighted averages of the values of the original acts f and g. Hence, the decision maker compares acts by applying a subjective expected utility criterion to their subjective transformations.

For the purpose of ranking acts conditionally on some event, the decision maker discards the contingencies that are deemed irrelevant and then normalizes the probabilities of the remaining contingencies (i.e., re-calculates the probabilities by dividing them by their sum).

Hence, at the level of contingencies, the decision maker operates with probabilities in a perfectly Bayesian way. Within the scope of our general representation, the procedure captured by (5)–(8) constitutes the minimal departure from the standard paradigm.

The third feature lies in the specification of the contingencies that remain relevant after the realization of some event. Instead of relying on arbitrary monotone set functions, the procedure consists of computing the total weight of the intersection of the event and the set of states representing a contingency, and then using a threshold strategy to determine whether this contingency is consistent with the event. Moreover, the weights used for these calculations are the same as those used for computing the values of acts. As an interpretation, the decision maker captured by this procedure acts as if she is "overlooking" the low-weight states when incorporating arriving information into her decisions.

Next, we illustrate Theorem 6 by briefly revisiting our application to the Confirmatory Bias.

Example 3 (Continuation of Example 2)**.** *From Example 2, recall that the agent considers hypotheses A and B and receives a noisy signal $s \in \{a, m, b\}$ representing the correct hypothesis with conditional probabilities listed therein. Our goal is to model an agent who has a bias towards B in interpreting the signal. The ex ante probability that B holds is* 0.6*.*

As before, let the state space be $\Omega = \{A, B\} \times \{a, m, b\}$, the subjective contingencies be $C_1 = \{A\} \times \{a\}$, $C_2 = \{A\} \times \{m, b\}$, $C_3 = \{B\} \times \{a, m\}$, $C_4 = \{B\} \times \{b\}$, and the probability measures be $\mu = (0.2, 0.2, 0.3, 0.3)$, $p_2 = (0.6, 0.4)$, $p_3 = (0.4, 0.6)$. Example 2 has already made use of the SEU form for the functions V_i's. Now, the conditional evaluations of act f paying 1 *if B and* 0 *if A are computed simply as follows:*[21] $V(f \mid s = a) = \frac{\mu(C_3)}{\mu(C_1)+\mu(C_3)} = 0.6$*;* $V(f \mid s \in \{m, b\}) = \frac{\mu(C_3)+\mu(C_4)}{\mu(C_2)+\mu(C_3)+\mu(C_4)} = 0.75$*; and* $V(f \mid s = b) = \frac{\mu(C_4)}{\mu(C_4)} = 1$.

The novelty is that, by Theorem 6, we can also adopt a more specific updating mechanism based on the weights p_i and the thresholds α_i. So, let $\alpha_2 = 1$ and $\alpha_3 = 0.4$. In this calibration, the entire analysis of Example 2 holds. However, the strong bias for B in interpreting the signal is now related to a high value for α_2 and a relatively low one for α_3.

[21] In Example 2, we used act h as a measuring rod to simplify exposition (and to avoid re-normalization of conditional probabilities to 1). We no longer need this, given the more structured derivation of Theorem 6.

4.3 Uniqueness

Definition 9. We refer to a tuple $(\mathcal{S}, (C_i)_{i\in\mathcal{S}}, u, \mu, (p_i)_{i\in\mathcal{S}}, (\alpha_i)_{i\in\mathcal{S}})$ that satisfies the conditions of Theorem 6 as an *expected utility representation with subjective contingencies.*

The uniqueness properties of our expected utility representation with subjective contingencies are summarized by the following proposition.

Proposition 7. *Suppose that* $(\mathcal{S}, (C_i)_{i\in\mathcal{S}}, u, \mu, (p_i)_{i\in\mathcal{S}}, (\alpha_i)_{i\in\mathcal{S}})$ *and* $(\mathcal{S}', (C'_i)_{i\in\mathcal{S}'}, u', \mu', (p'_i)_{i\in\mathcal{S}'}, (\alpha'_i)_{i\in\mathcal{S}'})$ *are two expected utility representations with subjective contingencies. Then, they represent the same system* $(\succsim, \{\succsim_A\}_{A\subset\Omega})$ *of preference relations if and only if*

(i) there exist $k > 0$ *and* $b \in \mathbb{R}$ *such that* $u' = ku + b$*;*

(ii) there exists a bijection $\pi : \mathcal{S} \to \mathcal{S}'$ *such that, for all* $i \in \mathcal{S}$, $C'_{\pi_i} = C_i$, $\mu'_{\pi_i} = \mu_i$, $p'_{\pi_i} = p_i$*;*

(iii) $\alpha'_{\pi_i} \in (\alpha_i^{min}, \alpha_i^{max}]$*, where*

$$\begin{cases} \alpha_i^{min} := \max\left\{\sum_{\omega\in A} p_i(\omega) \mid A \subseteq C_i \textit{ and } \succsim_A \textit{ is degenerate}\right\} \textit{ and} \\ \alpha_i^{max} := \min\left\{\sum_{\omega\in A} p_i(\omega) \mid A \subseteq C_i \textit{ and } \succsim_A \textit{ is nondegenerate}\right\}. \end{cases} \tag{10}$$

As usual, the utility function is unique up to a positive affine transformation; most of the parameters related to subjective contingencies are unique up to relabeling of the contingencies; and the only nontrivial part of the statement concerns the thresholds $(\alpha_i)_{i\in\mathcal{S}}$. These thresholds can be identified only up to an interval for a simple reason — our state space is finite, and, hence, the total weights $\sum_{\omega\in A} p_i(\omega)$ for different $A \subseteq \Omega$ form a discrete set. We also note that the thresholds α_i of the representation can be chosen within each interval $(\alpha_i^{min}, \alpha_i^{max}]$ arbitrarily and independently across different contingencies, and the boundaries of these intervals can be determined uniquely from the agent's choice behavior.[22]

[22] To substantiate this claim, we give (without a proof) a behavioral definition of α^{min} and α^{max}. Suppose that $A \subset \Omega$ is such that $\omega \in A \Leftrightarrow P_A(\omega)$, A does not have a proper subset with the same property, and $B \subset \Omega$ is such that $A \cap B = \varnothing$. Then, A constitutes one of the agent's subjective contingencies, and the corresponding threshold boundaries can be computed as

$$\begin{cases} \alpha_A^{min} := \max\left\{\frac{\mathrm{CE}(x^*\,(E\cup B)\,x_*) - \mathrm{CE}(x^*\,B\,x_*)}{\mathrm{CE}(x^*\,A\,x_*)} \mid E \subseteq A, \succsim_{E\cup B} = \succsim_B\right\} \text{ and} \\ \alpha_A^{max} := \min\left\{\frac{\mathrm{CE}(x^*\,(E\cup B)\,x_*) - \mathrm{CE}(x^*\,B\,x_*)}{\mathrm{CE}(x^*\,A\,x_*)} \mid E \subseteq A, \succsim_{E\cup B} \neq \succsim_B\right\}, \end{cases}$$

where $\mathrm{CE} : \mathcal{F} \to [0,1]$ is the functional defined as $\mathrm{CE}(f) = \gamma \Leftrightarrow f \sim \gamma x^* + (1-\gamma)x_*$.

As a final remark, note that the limits on the identification of the collection $(\alpha_i)_{i\in\mathcal{S}}$ — up to an interval — do not preclude comparing these parameters across different agents. In particular, if two decision makers have the same ex ante preferences and equal understanding of uncertainty, then they may differ only in their thresholds $(\alpha_i)_{i\in\mathcal{S}}$, and the following observation holds:

Observation 8. *Suppose that two decision makers are described by $(\succsim, \{\succsim_A^k\}_{A\subset\Omega})$ for $k = 1, 2$, and their preferences admit expected utility representations with subjective contingencies $(\mathcal{S}, (C_i)_{i\in\mathcal{S}}, u, \mu, (p_i)_{i\in\mathcal{S}}, (\alpha_i^k)_{i\in\mathcal{S}})$ for $k = 1, 2$. Furthermore, suppose that $(\alpha_i^{k\,min})_{i\in\mathcal{S}}$ and $(\alpha_i^{k\,max})_{i\in\mathcal{S}}$ are defined by (10) for $k = 1, 2$. Then, for each $i \in \mathcal{S}$,*

$$(\alpha_i^{1\,min}, \alpha_i^{1\,max}] = (\alpha_i^{2\,min}, \alpha_i^{2\,max}] \text{ or } (\alpha_i^{1\,min}, \alpha_i^{1\,max}] \cap (\alpha_i^{2\,min}, \alpha_i^{2\,max}] = \varnothing.$$

The key implication of this observation is that if, for some $i \in \mathcal{S}$, the intervals for the thresholds α_i for these two agents are not identical, then one of the intervals must lie entirely to the left of the other one.

4.4 Comparative Statics

Since the expected utility representation with subjective contingencies is a special case of the general representation, the comparative notion of understanding of uncertainty and the corresponding characterization of Proposition 3 apply directly to the model presented in Section 4.2. Therefore, we focus here on the behavioral properties of the new parameters introduced there — namely, $(\alpha_i)_{i\in\mathcal{S}}$.

Consider two decision makers described by $(\succsim^k, \{\succsim_A^k\}_{A\subset\Omega})$ for $k = 1, 2$, with an equal understanding of uncertainty and such that $\succsim^1 = \succsim^2$. In terms of the representation, this implies that $\mathcal{S}^1 = \mathcal{S}^2, \Pi^1 = \Pi^2, \mu_i^1 = \mu_i^2, p_i^1 = p_i^2$ for all i, and u^1 is a positive affine transformation of u^2. Hence, such decision makers differ only in the way they update their preferences. In the next proposition, we apply the comparative notion of proneness to inclusion and exclusion errors (Definition 7) to our expected utility model with subjective contingencies.

Proposition 9. *Suppose that two decision makers are described by $(\succsim, \{\succsim_A^k\}_{A\subset\Omega})$ for $k = 1, 2$, and their preferences admit expected utility representations with subjective contingencies $(\mathcal{S}, (C_i)_{i\in\mathcal{S}}, u, \mu, (p_i)_{i\in\mathcal{S}}, (\alpha_i^k)_{i\in\mathcal{S}})$ for $k = 1, 2$. Then, the following conditions are equivalent:*

(i) Decision Maker 1 is more prone to exclusion errors than Decision Maker 2;

(ii) Decision Maker 2 is more prone to inclusion errors than Decision Maker 1;

(iii) For each $i \in \mathcal{S}$, $\alpha_i^{1\min} = \alpha_i^{2\min}$ or $\alpha_i^{1\min} \geq \alpha_i^{2\max}$.

As in Proposition 5, the tendency to make exclusion versus inclusion errors represents two faces of the same trait. The novelty here lies in Condition (iii), in which the original characterization in terms of pointwise dominance of the consistency maps is replaced with a more structured form of an ordering of intervals on the real line. Specifically, two decision makers are comparable in terms of their predisposition to make errors in their conditional behavior whenever, for each contingency i, the interval of possible values of α_i for one agent weakly dominates the corresponding one of the other agent.

Appendix

Throughout the entire Appendix, we use the mapping $Q : 2^\Omega \to 2^\Omega$ defined as $Q(E) := \{\omega \in \Omega : P_E(\omega) \text{ holds}\}$. This correspondence determines the states that are revealed to be possible after each event.

A Proofs of the results of Section 3

Lemma 10. *Assume that axioms (A1)–(A9) hold. Then:*

(i) For any $A \subseteq \Omega$, $\succsim_{Q(A)} = \succsim_A$, and $Q(Q(A)) = Q(A)$.

(ii) For any $S, T \subseteq \Omega$ such that $Q(S) = S$ and $Q(T) = T$, we have $Q(S \backslash T) = S \backslash T$.

Proof. Claim (i). The claim that $\succsim_{Q(A)} = \succsim_A$ follows from the Equivalence of Events axiom, and $Q(Q(A)) = Q(A)$ follows immediately from that.

Claim (ii). Let $S' := S\backslash T$. If $S' = \varnothing$, the claim holds by the definition of Q. Suppose that $S' \neq \varnothing$. Observe that $S \subseteq Q(S \cup T)$ by Monotonicity of the Possibility Predicate. Hence, $S' \subseteq Q(S \cup T)\backslash T = Q(S \cup T)\backslash Q(T) = Q(S' \cup T)\backslash Q(T) \subseteq Q(S')$, where the latter inclusion holds by Understanding of Relationships Between Events, Part (i), applied to sets T and S'.

Next, we claim that $Q(S') \cap T = \varnothing$. Indeed, if there exists $\omega_0 \in Q(S') \cap T$, then $P_\varnothing(\omega_0)$ should hold by Understanding of Relationships Between Events, Part (ii), applied to sets T and S', a contradiction to Indifference Upon Impossible Events. Since $Q(S') \subseteq S$ by Monotonicity of the Possibility Predicate, we conclude that $Q(S') \subseteq S'$, which completes the proof. $\square$

Lemma 11. *Assume that axioms (A1)–(A9) hold. For any $E \subseteq \Omega$ and $f, g \in \mathcal{F}$, if $Q(E) \subseteq \{\omega \in \Omega : f(\omega) = g(\omega)\}$, then $f \sim_E g$.*

Proof. Fix an arbitrary $E \subseteq \Omega$ and $f, g \in \mathcal{F}$ such that $Q(E) \subseteq \{\omega \in \Omega : f(\omega) = g(\omega)\}$. Let $\omega_1, \dots, \omega_m$ be an enumeration of the states of Ω. For each $i = 0, \dots, m$, let $h^i \in \mathcal{F}$ be defined as $h^i = g\,\{\omega_{i+1}, \dots, \omega_m\}\,f$, and note that $h^0 = g$ and $h^m = f$. We claim that $h^{i-1} \sim_E h^i$ for all $i = 1, \dots, m$, which will prove the claim of the lemma by the transitivity of $\sim_E$.

Indeed $h^{i-1}(\omega_j) = h^i(\omega_j)$ for all $j \neq i$, while $h^{i-1}(\omega_i) = g(\omega_i)$ and $h^i(\omega_i) = f(\omega_i)$. If $g(\omega_i) = f(\omega_i)$, then $h^{i-1} = h^i$, and the claim is proven. Otherwise, we have $\neg P_E(\omega_i)$, and, therefore, $h^{i-1} \sim_E h^i$ by the definition of $\neg P_E(\omega_i)$. $\square$

Lemma 12. *Assume that axioms (A1)–(A9) hold and suppose that $\{C_1, \dots, C_n\}$ for some $n \in \mathbb{N}$ is a partition of $Q(\Omega)$ such that $Q(C_i) = C_i$ for all $i = 1, \dots, n$. Then, for all $f, g \in \mathcal{F}$ such that $f \sim_{C_i} g$ for all $i = 1, \dots, n$, we have $f \sim g$.*

Proof. For each $i = 0, \dots, n$, let $h^i \in \mathcal{F}$ be defined as

$$h^i(\omega) := \begin{cases} g(\omega), & \text{if } \omega \in C_{i+1} \cup \dots C_n, \\ f(\omega), & \text{otherwise.} \end{cases}$$

We will prove by induction that $h^i \sim g$ for all $i = 0, \dots, n$, which will establish the claim of this lemma because $h^n = f$.

If $i = 0$, then $h^0 \sim g$ by Lemma 11 when Ω plays the role of E. Assume that $h^{i-1} \sim g$ for some $i = 1, \dots, n$. Then, observe that $h^{i-1} \sim_{C_i} g$ by Lemma 11, $g \sim_{C_i} f$ by the conditions of the lemma, and $f \sim_{C_i} h^i$ by Lemma 11. Therefore, by transitivity, we have $h^{i-1} \sim_{C_i} h^i$. Since that $h^{i-1}(\omega) = h^i(\omega)$ for all $\omega \notin C_i$, Dynamic Consistency implies that $h^{i-1} \sim h^i$, and, hence, $h^i \sim g$. $\square$

Date: 2016-10-28 15:45:21 Revision: 1.97.2.1

Lemma 13. *Assume that axioms (A1)–(A9) hold and suppose that* $\{C_1, \ldots, C_n\}$ *for some* $n \in \mathbb{N}$ *is a partition of* $Q(\Omega)$ *such that* $Q(C_i) = C_i$ *for all* $i = 1, \ldots, n$*, and* $U_i : \mathcal{F} \to \mathbb{R}$ *are continuous utility representations of* $\succsim_{C_i}$ *for* $i = 1, \ldots, n$*. Let* $U : \mathcal{F} \to \Gamma$ *be defined as* $U(f) = (U_1(f), \ldots, U_n(f))$ *for all* $f \in \mathcal{F}$*, and* $\mathcal{U} := U(\mathcal{F})$*. Finally, let* $v^0 \in \mathcal{U}$ *and* $(v^m)_{m \in \mathbb{N}} \in \mathcal{U}^{\mathbb{N}}$ *be such that* $v^m \to v^0$ *as* $m \to \infty$*. Then, there exist a subsequence* $(v^{m_k})_{k \in \mathbb{N}}$*,* $f^0 \in \mathcal{F}$*, and* $(f^k)_{k \in \mathbb{N}} \in \mathcal{F}^{\mathbb{N}}$ *such that* $U(f^k) = v^{m_k}$ *for all* $k \in \mathbb{N}$*,* $U(f^0) = v^0$*, and* $f^k \to f^0$ *as* $k \to \infty$*.*

Proof. For $i = 1, \ldots, n$, let $\varphi_i : [0,1] \to X$ be defined as $\varphi_i(\alpha) = \alpha x^* + (1-\alpha) x_*$, and note that it is continuous. By the Best and Worst Outcomes axiom, $U_i(\varphi_i(0)) \leq U_i(f) \leq U_i(\varphi_i(1))$ for all $f \in \mathcal{F}$ and $i = 1, \ldots, n$. Hence, the mappings $U_i \circ \varphi_i : [0,1] \to U_i(\mathcal{F})$ are surjections for $i = 1, \ldots, n$. For $i = 1, \ldots, n$, let $\psi_i : U_i(\mathcal{F}) \to [0,1]$ be defined such that $U_i \circ \varphi_i \circ \psi_i$ are the identity mappings, and $\psi : \mathcal{U} \to [0,1]^n$ be defined as $\psi(u) = (\psi_1(u_1), \ldots, \psi_n(u_n))$ for all $u \in \mathcal{U}$.

Next, let $\Phi : [0,1]^n \to \mathcal{F}$ be defined as $\Phi(\alpha)(\omega) = \varphi_i(\alpha_i)$ if $\omega \in C_i$ for some $i = 1, \ldots, n$ and $\Phi(\alpha)(\omega) = x_*$ if $\omega \notin Q(\Omega)$ for all $\alpha \in [0,1]^n$. We claim that $U \circ \Phi \circ \psi$ is the identity mapping on $\mathcal{U}$. Indeed, fix an arbitrary $u \in \mathcal{U}$, and let $f = \Phi(\psi(u))$. Observe that, for all $i = 1, \ldots, n$, we have $f(\omega) = \varphi_i(\psi_i(u_i))$ for all $\omega \in C_i$ by the construction of Φ. Therefore, for all $i = 1, \ldots, n$, we have $f \sim_{C_i} \varphi_i(\psi_i(u_i))$ by Lemma 11, and, since U_i is a utility representation of $\succsim_{C_i}$, it follows that $U_i(f) = U_i(\varphi_i(\psi_i(u_i))) = u_i$.

Now, in the sequence $(\psi(v^m))_{m \in \mathbb{N}}$ of elements of $[0,1]^n$ one can choose a subsequence $(\psi(v^{m_k}))_{k \in \mathbb{N}}$ that converges to some $\alpha \in [0,1]^n$. Let $f^k = \Phi(\psi(v^{m_k}))$ for $k \in \mathbb{N}$ and $f^0 = \Phi(\alpha)$. As shown in the previous step, $U(f^k) = v^{m_k}$ for all $k \in \mathbb{N}$. By the continuity of φ, $f^k \to f^0$ as $k \to \infty$. Moreover, by the continuity of U, $U(f^0) = \lim_{k \to \infty} U(f^k) = v^0$. $\square$

Lemma 14. *Suppose that* $(\mathcal{S}, (C_i)_{i \in \mathcal{S}}, (V_i)_{i \in \mathcal{S}}, (\sigma_i)_{i \in \mathcal{S}})$ *is a representation with subjective contingencies of the system* $(\succsim, \{\succsim_E\}_{E \subset \Omega})$*. Let* $E \subseteq \Omega$ *and* $\omega \in \Omega$*. Then, the following conditions are equivalent:*

(i) $P_E(\omega)$ *holds;*

(ii) there exists $j \in \mathcal{S}$ *such that* $\omega \in C_j$ *and* $\sigma_j(E \cap C_j) = 1$*, where* C_j *is a cell of the partition* Π.

Proof. Suppose that $P_E(\omega)$ holds but there is no $j \in \mathcal{S}$ such that $\omega \in C_j$ and $\sigma_j(E \cap C_j) = 1$. Then, for any $f \in \mathcal{F}$ and $x, y \in X$,

$$\begin{aligned} V(x\{\omega\}f \mid E) &= \sum_{i \in \mathcal{S}:\sigma_i(E \cap C_i)=1} V_i(x\{\omega\}f) \\ &= \sum_{i \in \mathcal{S}:\sigma_i(E \cap C_i)=1} V_i(f) && \text{(by the properties of } V_i\text{)} \\ &= V(y\{\omega\}f \mid E). \end{aligned}$$

Since $V(\cdot \mid E)$ represents $\succsim_E$, this contradicts to the fact that $P_E(\omega)$ holds.

Conversely, suppose that there exists $j \in \mathcal{S}$ such that $\omega \in C_j$ and $\sigma_j(E \cap C_j) = 1$. Next, due to the fact that C_j is the support of V_j, we can find $f \in \mathcal{F}$ and $x, y \in X$ such that $V_j(x\{\omega\}f) \neq V_j(y\{\omega\}f)$ and assume, without loss of generality, that $V_j(x\{\omega\}f) > V_j(y\{\omega\}f)$. Then,

$$\begin{aligned} V(x\{\omega\}f \mid E) &= \sum_{i \in \mathcal{S}:\sigma_i(E \cap C_i)=1} V_i(x\{\omega\}f) \\ &= V_j(x\{\omega\}f) + \sum_{i \in \mathcal{S}, i \neq j} V_i(x\{\omega\}f)\,\sigma_i(E \cap C_i) \\ &= V_j(x\{\omega\}f) + \sum_{i \in \mathcal{S}, i \neq j} V_i(f)\,\sigma_i(E \cap C_i) \\ &> V_j(y\{\omega\}f) + \sum_{i \in \mathcal{S}, i \neq j} V_i(f)\,\sigma_i(E \cap C_i) \\ &= V(y\{\omega\}f \mid E), \end{aligned}$$

and, therefore, $x\{\omega\}f \succ_E y\{\omega\}f$ because $V(\cdot \mid E)$ represents $\succsim_E$. $\square$

Proof of Theorem 1. *Only if* part. *Step 1. The algebra of contingencies.* Consider the collection of events

$$\mathcal{A} := \{Q(E) \mid E \subseteq \Omega\}.$$

Note that $\varnothing = Q(\varnothing) \in \mathcal{A}$ by Indifference Upon Impossible Events, and all sets in $\mathcal{A}$ are subsets of $Q(\Omega)$ by Monotonicity of the Possibility Predicate. By Lemma 10, we have $S \backslash T \in \mathcal{A}$ for any $S, T \in \mathcal{A}$. It follows that $S \cap T \in \mathcal{A}$ and $S \cup T \in \mathcal{A}$ for any $S, T \in \mathcal{A}$, and $\mathcal{A}$ is a Boolean algebra.

Step 2. The subjective state space. Since $\mathcal{A}$ is finite, we can let $\Pi = \{C_1, \ldots, C_n\}$ be the collection of the atoms of $\mathcal{A}$ — i.e., a collection of nonempty sets from $\mathcal{A}$ such that, for all

$D \in \mathcal{A}$, we have either $C_i \cap D = \varnothing$ or $C_i \cap D = C_i$ for all $i = 1, \ldots, n$. Clearly, Π is a partition of $Q(\Omega)$, and, for any $D \in \mathcal{A}$, there exists $k \in \mathbb{N}$ and a collection of indices $i_1, \ldots, i_k \in \{1, \ldots, n\}$ such that $D = \bigcup_{j=1}^{k} C_{i_j}$. Let $\mathcal{S} := \{1, \ldots, n\}$. Note that $n \geq 3$, because otherwise by Equivalence of Events there would be at most two distinct nondegenerate conditional preferences $\succsim_{C_1}$ and $\succsim_{C_2}$ that are also different from $\succsim$, a violation of Nondegeneracy.

Step 3. Conditional utilities. The space $\mathcal{F}$ is connected and separable. Then, for each $i \in \mathcal{S}$, we can apply Debreu Theorem to get a continuous utility function $U_i : \mathcal{F} \to \mathbb{R}$ that represents the preference relation $\succsim_{C_i}$. For each $i \in \mathcal{S}$, let $\Gamma_i := U_i(\mathcal{F})$ and note that Γ_i is a nondegenerate interval due to the Outcome Preference Consistency and Nontriviality axioms. Let $\Gamma := \prod_{i=1}^{n} \Gamma_i$.

Step 4. Additive separability. We define a binary relation $\trianglerighteq$ on Γ as $v \trianglerighteq w \Leftrightarrow f \succsim g$ for some $f, g \in \mathcal{F}$ such that $U_i(f) = v_i$ and $U_i(g) = w_i$ for all $i = 1, \ldots, n$. (Note that if $f \succsim g$ holds for some $f, g \in \mathcal{F}$ satisfying these conditions, then it holds for any such $f, g \in \mathcal{F}$ by Lemma 12.)

Now, we claim that $\trianglerighteq$ has the Coordinate Independence property (Wakker, 1989, Def. II.2.3): For any $v, w \in \Gamma$, $i \in \mathcal{S}$, $a, b \in \Gamma_i$,

$$(v_1, \ldots, v_{i-1}, a, v_{i+1}, \ldots, v_n) \trianglerighteq (w_1, \ldots, w_{i-1}, a, w_{i+1}, \ldots, w_n) \quad \Leftrightarrow$$
$$(v_1, \ldots, v_{i-1}, b, v_{i+1}, \ldots, v_n) \trianglerighteq (w_1, \ldots, w_{i-1}, b, w_{i+1}, \ldots, w_n).$$

Indeed, fix arbitrary $v, w \in \Gamma$, $i \in \mathcal{S}$, $a, b \in \Gamma_i$, and let $f, g \in \mathcal{F}$ be such that $U_j(f) = v_j$ and $U_j(g) = w_j$ for all $j = 1, \ldots, n$, $j \neq i$. For any $\zeta \in \Gamma_i$, let $z \in X$ be such that $U_i(z) = \zeta$, and observe that

$$\begin{aligned}
&(v_1, \ldots, v_{i-1}, \zeta, v_{i+1}, \ldots, v_n) \trianglerighteq (w_1, \ldots, w_{i-1}, \zeta, w_{i+1}, \ldots, w_n) \Leftrightarrow \\
&\quad z\, C_i\, f \succsim z\, C_i\, g \qquad\qquad \Leftrightarrow \quad \text{(by Dynamic Consistency)} \\
&\quad z\, C_i\, f \succsim_{Q(\Omega)\setminus C_i} z\, C_i\, g \quad \Leftrightarrow \quad \text{(by Lemma 11)} \\
&\qquad f \succsim_{Q(\Omega)\setminus C_i} g.
\end{aligned}$$

Note that the latter relationship does not depend on the value of ζ and, therefore, Coordinate Independence is proven.

Next, we prove that $\trianglerighteq$ is continuous — i.e., for any $w \in \Gamma$, the sets $\{v \in \Gamma : v \trianglerighteq w\}$ and $\{v \in \Gamma : w \trianglerighteq v\}$ are closed. Fix an arbitrary $w \in \Gamma$, and suppose that $v^0 \in \Gamma$ and a sequence

$(v^m)_{m\in\mathbb{N}}$ in Γ are such that $v^m \trianglerighteq w$ for all $m \in \mathbb{N}$ and $v^m \to v^0$ as $m \to \infty$. By Lemma 13, we can find a subsequence $(v^{m_k})_{k\in\mathbb{N}}$, $f^0 \in \mathcal{F}$, and a sequence of acts $(f^k)_{k\in\mathbb{N}}$ such that, for all $i = 1,\dots,n$, we have $U_i(f^k) = v_i^{m_k}$ for all $k \in \mathbb{N}$ and $U_i(f^0) = v_i^0$, and such that $f^k \to f^0$ as $k \to \infty$. Now, suppose that $g \in \mathcal{F}$ is such that $U_i(g) = w_i$ for all $i = 1,\dots,n$. Then, $v^k \trianglerighteq w$ implies that $f^k \succsim g$ by the construction of $\trianglerighteq$, which, by Continuity, implies that $f^0 \succsim g$, and, in turn, $v^0 \trianglerighteq w$. We conclude that the set $\{v \in \Gamma : v \trianglerighteq w\}$ is closed. Similarly, it can be shown that the set $\{v \in \Gamma : w \trianglerighteq v\}$ is closed, as well.

Finally, we claim that all coordinates in Γ are essential: Indeed, fix arbitrary $x, y \in X$ such that $x \succ y$. By Outcome Preference Consistency, we have $x \succ_{C_i} y$ and $U_i(x) > U_i(y)$ for all $i = 1,\dots,n$, Furthermore, for all $i = 1,\dots,n$, $x\, C_i\, y \sim_{C_i} x$ by Lemma 11, and, in turn, $x\, C_i\, y \succ_{C_i} y$ by transitivity, and, hence, $x\, C_i\, y \succ y$ by Dynamic Consistency. Then, $(U_1(y),\dots,U_{i-1}(y),U_i(x),U_{i+1}(y),\dots,U_n(y)) \triangleright (U_1(y),U_2(y),\dots,U_n(y))$ and, therefore, coordinate i is essential for all $i = 1,\dots,n$.

Having established all the above-listed properties of $\trianglerighteq$, we can conclude by Wakker (1989, Th. III.4.1) that there exist continuous functions $W_i : \Gamma_i \to \mathbb{R}$ for $i = 1,\dots,n$ such that, for all $v, w \in \Gamma$, $v \trianglerighteq w \Leftrightarrow \sum_{i=1}^n W_i(v_i) \geq \sum_{i=1}^n W_i(w_i)$.

Step 5. The ex-ante representation. For any $f, g \in \mathcal{F}$, we let $v := (U_1(f),\dots,U_n(f))$, $w := (U_1(g),\dots,U_n(g))$, and observe that

$$f \succsim g \quad \Leftrightarrow \quad v \trianglerighteq w \quad \Leftrightarrow \quad \sum_{i=1}^{n} W_i(U_i(f)) \geq \sum_{i=1}^{n} W_i(U_i(g)).$$

Let $V_i := W_i \circ U_i$ for all $i \in \mathcal{S}$. For each i, the function V_i is continuous (as a composition of continuous functions), is nonconstant, and has compact range by the Best and Worst Outcomes axiom. Finally, we prove that C_i is the support of V_i for any $i = 1,\dots,n$: Indeed, for any $f, g \in \mathcal{F}$ such that $f|_{C_i} = g|_{C_i}$, we have $f \sim_{C_i} g$ by Lemma 11; therefore, $U_i(f) = U_i(g)$ and, in turn, $V_i(f) = V_i(g)$. Moreover, for each $\omega \in C_i$, we have $P_{C_i}(\omega)$ by construction, and, therefore, there exist $f \in \mathcal{F}$ and $x, y \in X$ such that $x\,\{\omega\}\, f \succ_{C_i} y\,\{\omega\} f$, which implies that $V(x\,\{\omega\}\, f) > V(y\,\{\omega\} f)$.

Step 6. The conditional representation on $\mathcal{A}$. For all $E \in \mathcal{A}$, let

$$V(f \mid E) := \sum_{i\in\mathcal{S}:C_i\subseteq E} V_i(f) \quad \forall f \in \mathcal{F}.$$

We claim that $\succsim_E$ is represented by $V(\cdot \mid E)$ for all $E \in \mathcal{A}$. Indeed, for $E = \varnothing$ the claim holds trivially. Otherwise, consider an arbitrary $E \in \mathcal{A}\backslash\{\varnothing\}$, and notice that $\succsim_E$ is nondegenerate.

Fix an arbitrary $h \in \mathcal{F}$. Then, for any $f, g \in \mathcal{F}$, we have

$$
\begin{aligned}
& f \succsim_E g && \Leftrightarrow \quad \text{(by Lemma 11)} \\
& \quad fEh \succsim_E gEh && \Leftrightarrow \quad \text{(by Dynamic Consistency)} \\
& \quad fEh \succsim gEh && \Leftrightarrow \quad \text{(by the ex-ante representation)} \\
& \quad \textstyle\sum_{i \in \mathcal{S}} V_i(fEh) \geq \sum_{i \in \mathcal{S}} V_i(gEh) && \Leftrightarrow \quad \text{(because of the support of } V_i\text{)} \\
& \qquad\qquad \textstyle\sum_{i \in \mathcal{S}: C_i \subseteq E} V_i(f) \geq \sum_{i \in \mathcal{S}: C_i \subseteq E} V_i(g).
\end{aligned}
$$

Step 7. The functions σ_i and the conditional representation. For each $i \in \mathcal{S}$, let $\sigma_i \colon 2^{C_i} \to \{0, 1\}$ be a function defined by

$$\sigma_i(E) = 1 \quad \Leftrightarrow \quad C_i \subseteq Q(E). \tag{11}$$

These functions, clearly, satisfy the normalization properties: For all $i \in \mathcal{S}$, $\sigma_i(\varnothing) = 0$ and $\sigma_i(C_i) = 1$ because $Q(\varnothing) = \varnothing$ and $Q(C_i) = C_i$. Furthermore, the mapping $Q : 2^\Omega \to 2^\Omega$ is monotone (inclusion-wise) by Monotonicity of the Possibility Predicate, which implies that σ_i is monotone for all $i \in \mathcal{S}$.

For any $f, g \in \mathcal{F}$ and $E \subset \Omega$, we have

$$
\begin{aligned}
& f \succsim_E g && \Leftrightarrow \quad \text{(by Part (i) of Lemma 10)} \\
& \quad f \succsim_{Q(E)} g && \Leftrightarrow \quad \text{(by Step 6)} \\
& \quad \sum_{i \in \mathcal{S}: C_i \subseteq Q(E)} V_i(f) \geq \sum_{i \in \mathcal{S}: C_i \subseteq Q(E)} V_i(g) && \Leftrightarrow \quad \text{(by (11))} \\
& \qquad\qquad \sum_{i \in \mathcal{S}: \sigma_i(E \cap C_i) = 1} V_i(f) \geq \sum_{i \in \mathcal{S}: \sigma_i(E \cap C_i) = 1} V_i(g),
\end{aligned}
$$

and Representation (3) is proven.

If part. Assume that there exist a set $\mathcal{S} = \{1, \ldots, n\}$, a collection $\Pi = \{C_1, \ldots, C_n\}$ of subsets of Ω, a collection of functions $V_i : \mathcal{F} \to \mathbb{R}$, and a collection of functions $\sigma_i : 2^{C_i} \to \{0, 1\}$ for $i \in \mathcal{S}$ as described in Theorem 1 and such that statements (i)–(iii) hold.

Continuity. This follows easily from Representation (2) and the continuity of V_i for $i \in \mathcal{S}$.

Best and Worst Outcomes. Since the function V_1 has a compact range, there exist $x^*, x_* \in X$ such that $V_1(x^*) \geq V_1(f) \geq V_1(x_*)$ for all $f \in \mathcal{F}$. By Condition (iii), we have $V_i(x^*) \geq V_i(f) \geq V_i(x_*)$ for all $f \in \mathcal{F}$ and all $i \in \mathcal{S}$ and, therefore, by (2), we have $x^* \succsim f \succsim x_*$ for all $f \in \mathcal{F}$.

Nontriviality. Since $n \geq 3$, we can consider the sets C_1, C_2, C_3, and note that $\succsim_{C_1}$, $\succsim_{C_2}$, and $\succsim_{C_3}$ are nondegenerate because V_i are nonconstant for $i = 1, 2, 3$. It can be verified that $\succsim$, $\succsim_{C_1}$, $\succsim_{C_2}$, and $\succsim_{C_3}$ differ from each other by using (2)–(3) and checking how these preference relations rank the acts $x^* \, C_1 \, x_*$, $x^* \, C_2 \, x_*$, $x^* \, C_3 \, x_*$.

Outcome Preference Consistency. Fix an arbitrary $E \subset \Omega$, and suppose that $x, y \in X$ are such that $x \succsim y$. Then, $\sum_{i \in \mathcal{S}} V_i(x) \geq \sum_{i \in \mathcal{S}} V_i(y)$ implies that there exists $j \in \mathcal{S}$ such that $V_j(x) \geq V_j(y)$. By (iii), we have that $V_i(x) \geq V_i(y)$ for all $i \in \mathcal{S}$. Therefore, $\sum_{i \in \mathcal{S}: \sigma_i(E \cap C_i) = 1} V_i(x) \geq \sum_{i \in \mathcal{S}: \sigma_i(E \cap C_i) = 1} V_i(y)$ and $x \succsim_E y$. Conversely, suppose that $x, y \in X$ are such that $x \succsim_E y$. Since $\succsim_E$ is nondegenerate, there exists at least one $j \in \mathcal{S}$ such that $\sigma_j(E \cap C_j) = 1$. Then, there must be some $j \in \mathcal{S}$ such that $V_j(x) \geq V_j(y)$. By (iii), we have that $V_i(x) \geq V_i(y)$ for all $i \in \mathcal{S}$, and $\sum_{i \in \mathcal{S}} V_i(x) \geq \sum_{i \in \mathcal{S}} V_i(y)$, which implies that $x \succsim y$.

Subjective Dynamic Consistency. Suppose that $E \subset \Omega$ is such that $\succsim_E$ is nondegenerate, and $f, g \in \mathcal{F}$ are such that $f(\omega) = g(\omega)$ for all $\omega \in \Omega$ such that $\neg P_E(\omega)$. The equivalence $\sum_{i \in \mathcal{S}} V_i(f) \geq \sum_{i \in \mathcal{S}} V_i(g) \Leftrightarrow \sum_{i \in \mathcal{S}: \sigma_i(E \cap C_i) = 1} V_i(f) \geq \sum_{i \in \mathcal{S}: \sigma_i(E \cap C_i) = 1} V_i(g)$ obtains if we establish that $V_i(f) = V_i(g)$ for all $i \in \mathcal{S}$ such that $\sigma_i(E \cap C_i) = 0$. To prove the latter, fix any such $i \in \mathcal{S}$. For any $\omega \in C_i$, we have $\neg P_E(\omega)$ due to Lemma 14 and the fact that Π is a partition, and, hence, $f(\omega) = g(\omega)$ by assumption. Since C_i is the support of V_i, we obtain $V_i(f) = V_i(g)$, and the claim is proven.

Monotonicity of the Possibility Predicate. Suppose that $A \subseteq B \subseteq \Omega$ and $\omega \in \Omega$ is such that $P_A(\omega)$ holds. By Lemma 14, there exists $j \in \mathcal{S}$ such that $\omega \in C_j$ and $\sigma_j(A \cap C_j) = 1$. Since σ_j is monotone, $B \cap C_j \supseteq A \cap C_j$ implies $\sigma_j(B \cap C_j) \geq \sigma_j(A \cap C_j)$. Thus, $P_B(\omega)$ holds by Lemma 14.

Equivalence of Events. Suppose that $E, E' \subseteq \Omega$ are related so that: for any $\omega \in E' \backslash E$, $P_E(\omega)$ holds; and for any $\omega \in E \backslash E'$, $P_E(\omega)$ does not hold. By (3), $\succsim_E = \succsim_{E'}$ will be proven if we show that, for any $i \in \mathcal{S}$, $\sigma_i(E \cap C_i) = \sigma_i(E' \cap C_i)$. First, suppose that $\sigma_i(E \cap C_i) = 1$. Then, for all $\omega \in E \cap C_i$, we have $P_E(\omega)$ by Lemma 14, and, hence, it must be that $\omega \in E'$. Therefore, $E \cap C_i \subseteq E' \cap C_i$, and $\sigma_i(E' \cap C_i) = 1$ by the monotonicity of σ_i. Second, suppose that $\sigma_i(E \cap C_i) = 0$. Then, for all $\omega \in C_i \backslash E$, we have $\neg P_E(\omega)$ by Lemma 14, and, hence, it must be that $\omega \notin E'$. Therefore, $E \cap C_i \supseteq E' \cap C_i$, and $\sigma_i(E' \cap C_i) = 0$ by the monotonicity of σ_i.

Understanding of Relationships Between Events. Suppose that $A \subseteq \Omega$ is such that $P_A(\omega) \Leftrightarrow$

$\omega \in A$ for all $\omega \in \Omega$, and $B \subseteq \Omega$. Note that $Q(A) = A$, and, hence, A is a union of some cells of Π. First, suppose that $\omega_0 \in \Omega$ is such that $P_A(\omega_0)$ and $P_B(\omega_0)$ hold. Using Lemma 14, we find $i \in \mathcal{S}$ such that $\omega_0 \in C_i$, $\sigma_i(A \cap C_i) = 1$, and $\sigma_i(B \cap C_i) = 1$. Since $\sigma_i(\varnothing) = 0$, it must be that $A \cap C_i = C_i$, and, hence, $\sigma_i(A \cap B \cap C_i) = \sigma_i(B \cap C_i) = 1$. By Lemma 14, it follows that $P_{A \cap B}(\omega_0)$ holds. Second, suppose that $\omega_0 \in \Omega$ is such that $P_{A \cup B}(\omega_0)$ holds but $P_A(\omega_0)$ does not. Using Lemma 14, we find $i \in \mathcal{S}$ such that $\omega_0 \in C_i$, $\sigma_i((A \cup B) \cap C_i) = 1$, and $\sigma_i(A \cap C_i) = 0$. Since $\sigma_i(C_i) = 1$, it must be that $A \cap C_i = \varnothing$, and, hence, $\sigma_i(B \cap C_i) = \sigma_i((A \cup B) \cap C_i) = 1$. By Lemma 14, it follows that $P_B(\omega_0)$ holds.

Indifference Upon Impossible Events. The degeneracy of $\succsim_\varnothing$ follows from (3) and the property that $\sigma_i(\varnothing) = 0$ for all $i \in \mathcal{S}$. □

Proof of Proposition 2. The sufficiency of Conditions (i)–(iii) can be easily verified. We will prove the necessity. Let $(\mathcal{S}, (C_i)_{i \in \mathcal{S}}, (V_i)_{i \in \mathcal{S}}, (\sigma_i)_{i \in \mathcal{S}})$ and $(\mathcal{S}', (C_i')_{i \in \mathcal{S}'}, (V_i')_{i \in \mathcal{S}'}, (\sigma_i')_{i \in \mathcal{S}'})$ be two representations with subjective contingencies of the same system $(\succsim, \{\succsim_A\}_{A \subset \Omega})$.

Let the algebra $\mathcal{A}$ be defined as in Step 1 of the proof of Theorem 1). As follows from the representation of conditional preferences (3), both $\{C_i\}_{i \in \mathcal{S}}$ and $\{C_i'\}_{i \in \mathcal{S}'}$ constitute the collections of atoms of $\mathcal{A}$. Therefore, there must exist a bijection $\pi : \mathcal{S} \to \mathcal{S}'$ such that $C'_{\pi_i} = C_i$ for all $i \in \mathcal{S}$.

Condition (ii) holds directly by Wakker (1989, Obs. III.6.6′).

Now, for any $i \in \mathcal{S}$ and $E \subseteq \Omega$, it follows from Lemma 14 that

$$\sigma_i(E \cap C_i) = 1 \Leftrightarrow C_i \subseteq Q(E) \Leftrightarrow C'_{\pi_i} \subseteq Q(E) \Leftrightarrow \sigma'_{\pi_i}(E \cap C_i) = 1.$$

Thus, Condition (iii) holds, as well. □

Proof of Proposition 3. Fix an arbitrary $E \subseteq \Omega$, and let $E' \subseteq E$ be defined as $E' = \{\omega \in E : P_\Omega(\omega) \text{ holds}\}$.

We claim that if $P_\Omega(\omega)$ does not hold for some $\omega \in \Omega$, then $P_E(\omega)$ does not hold, either. Indeed, if $P_E(\omega)$ holds, then one can find $x, y \in X$ and $f \in \mathcal{F}$ such that $x\,\{\omega\}\,f \succ_E y\,\{\omega\}\,f$. Then, $\succsim_E$ is nondegenerate and, by the Subjective Dynamic Consistency axiom, it must be that $x\,\{\omega\}\,f \succ y\,\{\omega\}\,f$, a contradiction.

Then, by the Equivalence of Events axiom, $\succsim_E = \succsim_{E'}$. Hence, $P_E(\omega)$ holds if and only if $P_{E'}(\omega)$ holds. Consequently, by the definition of Q, E is fully understood if and only if $E' = Q(E')$.

Recall the algebra $\mathcal{A} = \{Q(A) \mid A \subseteq \Omega\}$ constructed in the proof of Theorem 1, the atoms of which are $C_1, \ldots, C_n$, the cells of Π. Then, $E' = Q(E')$ if and only if E' is a union of some elements of Π. We conclude that E is fully understood if and only if E is a union of some elements of $\Pi \cup \mathcal{N}(\succsim)$. $\square$

Proof of Proposition 5. We prove the proposition by showing that (i) $\Leftrightarrow$ (iii) and (ii) $\Leftrightarrow$ (iii).

(iii) $\Rightarrow$ *(i)* Suppose that $\sigma_i^1(S) \leq \sigma_i^2(S)$ for all $S \subseteq C_i$ and all $i = \{1, \ldots, m\}$. Fix an arbitrary $E \subseteq \Omega$ such that $\succsim_E^1$ and $\succsim_E^2$ are nondegenerate, and let $A := M_-^2(E)$, noting that $A \subseteq E$. Let $\mathcal{I} := \{i \in \{1, \ldots, m\} : A \cap C_i \neq \varnothing\}$. As follows from Lemma 14, we have $\sigma_i^2(E \cap C_i) = 0$ for all $i \in \mathcal{I}$. Then, $\sigma_i^1(A \cap C_i) = 0$ and, hence, $A \cap C_i \subseteq M_-^1(E)$ for all $i \in \mathcal{I}$ by the same Lemma 14. Since $A = \bigcup_{i \in \mathcal{I}}(A \cap C_i)$, we obtain $A \subseteq M_-^1(E)$.

(iii) $\Rightarrow$ *(ii)* Suppose that $\sigma_i^1(S) \leq \sigma_i^2(S)$ for all $S \subseteq C_i$ and all $i = \{1, \ldots, m\}$. Fix an arbitrary $E \subseteq \Omega$ such that $\succsim_E^1$ and $\succsim_E^2$ are nondegenerate, and let $A := M_+^1(E)$, noting that $A \cap E = \varnothing$. Let $\mathcal{I} := \{i \in \{1, \ldots, m\} : A \cap C_i \neq \varnothing\}$. As follows from Lemma 14, we have $\sigma_i^1(E \cap C_i) = 1$ for all $i \in \mathcal{I}$. Then, $\sigma_i^2(A \cap C_i) = 1$ and, hence, $A \cap C_i \subseteq M_+^2(E)$ for all $i \in \mathcal{I}$ by Lemma 14 again, which implies that $A \subseteq M_+^2(E)$.

(i) $\Rightarrow$ *(iii)* Suppose that $M_-^2(E) \subseteq M_-^1(E)$ for all $E \subseteq \Omega$ such that $\succsim_E^1$ and $\succsim_E^2$ are nondegenerate. Fix arbitrary $i \in \{1, \ldots, m\}$ and $A \subseteq C_i$. Our goal is to prove that $\sigma_i^2(A) = 0$ implies $\sigma_i^1(A) = 0$. Assume that $\sigma_i^2(A) = 0$, let $E := (\Omega \backslash C_i) \cup A$, and note that $\succsim_E$ must be nondegenerate because representations with subjective contingencies must have at least three cells. Observe that $A \subseteq M_-^2(E)$ by Lemma 14. By assumption, we have $A \subseteq M_-^1(E)$ and, therefore, $\sigma_i^1(A) = 0$ by Lemma 14 again.

(ii) $\Rightarrow$ *(iii)* Suppose that $M_+^1(E) \subseteq M_+^2(E)$ for all $E \subseteq \Omega$ such that $\succsim_E^1$ and $\succsim_E^2$ are nondegenerate. Fix arbitrary $i \in \{1, \ldots, m\}$ and $A \subseteq C_i$. Our goal is to prove that $\sigma_i^1(A) = 1$ implies $\sigma_i^2(A) = 1$. Assume that $\sigma_i^1(A) = 1$. If $A = C_i$ then $\sigma_i^2(A) = 1$ by normalization. Otherwise, let $B \in \Pi^2$ be such that $B \neq C_i$, $E := A \cup B$, and note that $\succsim_E$ must be nondegenerate. By Lemma 14, we have $C_i \backslash A \subseteq M_+^1(E)$ and, therefore, $C_i \backslash A \subseteq M_+^2(E)$ by assumption. Then, it must be that $\sigma_i^2(A) = 1$ by Lemma 14 again. $\square$

B Proofs of the Results of Section 4

Proof of Theorem 6. *Only if* part. *Step 1.* Since $\succsim$ satisfies the Anscombe-Aumann axioms, it admits a SEU representation via the map $f \mapsto \sum_{\omega\in\Omega} u(f(\omega))\, p(\omega)$ for some nonconstant affine function $u : X \to \mathbb{R}$ and a probability measure $p \in \Delta(\Omega)$ (see, e.g., Fishburn, 1970, Theorem 13.3). Moreover, u is continuous and has a compact range by the Continuity and Best and Worst Outcomes axioms, respectively.

Step 2. By Theorem 1, there exist $\mathcal{S} := \{1, \dots, n\}$, a collection $\Pi = \{C_1, \dots, C_n\}$ of nonempty disjoint subsets of Ω, and functions $V_i : \mathcal{F} \to \mathbb{R}$ and $\sigma_i : 2^{C_i} \to \{0, 1\}$ satisfying the conditions listed in the theorem such that functions $V' : \mathcal{F} \to \mathbb{R}$ and $V'(\cdot \mid E) : \mathcal{F} \to \mathbb{R}$ defined as

$$V'(f) = \sum_{i\in\mathcal{S}} V_i(f)$$

$$V'(f \mid E) = \sum_{i\in\mathcal{S}:\sigma_i(E\cap C_i)=1} V_i(f)$$

are utility representations of $\succsim$ and $\succsim_E$ for all $E \subset \Omega$, respectively. By the uniqueness of additively separable representations (see, e.g., Wakker, 1989, Obs. III.6.6′), there exist $k > 0$ and $b_i \in \mathbb{R}$ for all $i \in \mathcal{S}$ such that $V_i(f) = k \sum_{\omega\in C_i} u(f(\omega))\, p(\omega) + b_i$, and we also have $p\left(\cup_{i\in\mathcal{S}} C_i\right) = 1$.

Step 3. Let $\mu \in \Delta(S)$ be defined as $\mu_i = p(C_i)$ for each $i \in \mathcal{S}$, and note that, for all $i \in \mathcal{S}$, $\mu_i > 0$ because V_i is nonconstant. For each $i \in \mathcal{S}$, let $p_i : \Omega \to \mathbb{R}_+$ be defined as

$$p_i(\omega) := \begin{cases} p(\omega)/\mu_i, & \text{if } \omega \in C_i, \\ 0, & \text{otherwise.} \end{cases}$$

By the previous step, we can represent $\succsim_E$ as a positive affine transformation $V(\cdot \mid E)$ of $V'(\cdot \mid E)$ as

$$V(f \mid E) = \frac{\sum_{i\in\mathcal{S}(E)} \sum_{\omega\in\Omega} u(f(\omega))\, p_i(\omega)\, \mu_i}{\sum_{i\in\mathcal{S}(E)} \mu_i},$$

where $\mathcal{S}(E) := \{i \in \mathcal{S} : \sigma_i(E \cap C_i) = 1\}$. Given any $f \in \mathcal{F}$, define $f^* : \mathcal{S} \to X$ as $f^*(i) = \sum_{\omega\in\Omega} p_i(\omega) f(\omega)$ for all $i \in \mathcal{S}$. Then, we have that $V(f \mid E) = \sum_{i\in\mathcal{S}(E)} u(f^*(i))\, \mu_{i|E}$, where $\mu_{i|E} = \frac{\mu_i}{\sum_{j\in\mathcal{S}(E)} \mu_j}$, is also a utility representation of $\succsim_E$ for all $E \subset \Omega$.

Step 4. For each $i \in \mathcal{S}$, let $\alpha_i := \min\{p(E)/\mu_i \mid E \subseteq C_i \text{ and } \sigma_i(E) = 1\}$. We claim that, for any $i \in \mathcal{S}$ and $E \subseteq C_i$, if $p(E) \geq \alpha_i \mu_i$ then $\sigma_i(E) = 1$. Indeed, fix arbitrary $i \in \mathcal{S}$ and

$E \subseteq C_i$, and let $E_0 \subseteq C_i$ be such that $\sigma_i(E_0) = 1$ and $\alpha_i = p(E_0)/\mu_i$. Note that $P_{E_0}(\omega)$ holds for all $\omega \in C_i$ because C_i is the support of p_i. Hence, we can apply the Admissibility of Valuable Events axiom to outcomes x^* and x_* and events E_0 and E to obtain that $\succsim_E$ is nondegenerate. Since $E \subseteq C_i$, it follows from Lemma 14 that $\sigma_i(E) = 1$.

Recalling the definition of p_i for $i \in \mathcal{S}$, we conclude that $\sum_{\omega \in E} p_i(\omega) \geq \alpha_i \Leftrightarrow \sigma_i(E) = 1$ for all $E \subseteq C_i$ and $i \in \mathcal{S}$. Thus, $\mathcal{S}(E) = \{i \in \mathcal{S} : \sum_{\omega \in E} p_i(\omega) \geq \alpha_i\}$.

If part. Assume that there exist a set $\mathcal{S} = \{1, \ldots, n\}$, a collection $\Pi = \{C_1, \ldots, C_n\}$ of subsets of Ω, a utility index $u : X \to \mathbb{R}$, a probability measure $\mu \in \Delta(\mathcal{S})$, collections of probability measures $p_i : \Omega \to \mathbb{R}_+$ and numbers $\alpha_i \in (0,1]$ for $i \in \mathcal{S}$ as described in Theorem 6 and such that statements (i)–(ii) hold.

Clearly, functions $V_i : \mathcal{F} \to \mathbb{R}$ and $\sigma_i : 2^{C_i} \to \{0,1\}$ for $i \in \mathcal{S}$ defined as

$$V_i(f) = u(f^*(i))\,\mu_i \qquad \text{and} \qquad \sigma_i(E) = 1 \Leftrightarrow \sum_{\omega \in E} p_i(\omega) \geq \alpha_i$$

satisfy the conditions of Theorem 1, so Axioms (A1)–(A9) hold. Monotonicity and Independence axioms follow from the representation by standard arguments.

It remains to show that Admissibility of Valuable Events holds. Let $x, y \in X$ be such that $x \succ y$ and $A, B \subseteq \Omega$ such that $\succsim_A$ is nondegenerate, $B \subseteq Q(A)$, and $xBy \succsim xAy$. By contradiction, suppose that $\succsim_B$ is degenerate, so that $\sum_{\omega \in B} p_i(\omega) < \alpha_i$ for all $i \in \mathcal{S}$. Since $B \subseteq Q(A)$, we observe that if $B \cap C_i \neq \varnothing$ for some $i \in \mathcal{S}$, then, by Lemma 14 and the definition of σ_i, we have $\sum_{\omega \in A} p_i(\omega) \geq \alpha_i$ and, hence, $i \in \mathcal{S}(A)$. We also have $\mathcal{S}(A) \neq \varnothing$ because $\succsim_A$ is nondegenerate. Hence, we obtain:

$$\begin{aligned}
V(xAy) &= u(y) + (u(x) - u(y)) \sum_{i \in \mathcal{S}} \sum_{\omega \in A} p_i(\omega)\,\mu_i \\
&\geq u(y) + (u(x) - u(y)) \sum_{i \in \mathcal{S}(A)} \sum_{\omega \in A} p_i(\omega)\,\mu_i \\
&\geq u(y) + (u(x) - u(y)) \sum_{i \in \mathcal{S}(A)} \alpha_i \mu_i \\
&> u(y) + (u(x) - u(y)) \sum_{i \in \mathcal{S}(A)} \sum_{\omega \in B} p_i(\omega)\,\mu_i \\
&\geq u(y) + (u(x) - u(y)) \sum_{i \in \mathcal{S} : B \cap C_i \neq \varnothing} \sum_{\omega \in B} p_i(\omega)\,\mu_i = V(xBy),
\end{aligned}$$

a contradiction. $\square$

Proof of Proposition 7. The sufficiency of the conditions can be easily verified. We will prove the necessity. Let $(\mathcal{S}, (C_i)_{i\in\mathcal{S}}, u, \mu, (p_i)_{i\in\mathcal{S}}, (\alpha_i)_{i\in\mathcal{S}})$ and $(\mathcal{S}', (C'_i)_{i\in\mathcal{S}'}, u', \mu', (p'_i)_{i\in\mathcal{S}'}, (\alpha'_i)_{i\in\mathcal{S}'})$ be expected utility representations with subjective contingencies of the same system $(\succsim, \{\succsim_A\}_{A\subset\Omega})$. Since an expected utility representation with subjective contingencies is a special case of the representation in Theorem 1, Proposition 2 implies that there exists a bijection $\pi : \mathcal{S} \to \mathcal{S}'$ such that $C'_{\pi_i} = C_i$ for all $i \in \mathcal{S}$. The claimed relationships between u' and u, μ' and μ, and $(p'_i)_{i\in\mathcal{S}}$ and $(p_i)_{i\in\mathcal{S}}$ follow easily from the uniqueness of the subjective expected utility representation. It remains to prove the claim regarding the thresholds $(\alpha'_i)_{i\in\mathcal{S}'}$.

Let $V_i : \mathcal{F} \to \mathbb{R}$ and $\sigma_i : 2^{C_i} \to \{0,1\}$ for $i \in \mathcal{S}$ be defined as

$$V_i(f) = \sum_{\omega\in C_i} u(f(\omega))\, p_i(\omega)\, \mu_i \qquad \text{and} \qquad \sigma_i(A) = 1 \Leftrightarrow \sum_{\omega\in A} p_i(\omega) \geq \alpha'_{\pi_i}.$$

Clearly, $(S, (C_i)_{i\in\mathcal{S}}, (V_i)_{i\in\mathcal{S}}, (\sigma_i)_{i\in\mathcal{S}})$ is a representation with subjective contingencies of $(\succsim, \{\succsim_A\}_{A\subset\Omega})$. Therefore, for any $i \in \mathcal{S}$ and $A \subseteq C_i$, it must be that $\sum_{\omega\in A} p_i(\omega) \geq \alpha'_{\pi_i}$ if only if $\succsim_A$ is nondegenerate (because V_i is nonconstant). Indeed, the situations $\alpha'_{\pi_i} > \sum_{\omega\in A} p_i(\omega)$ for some $A \subseteq C_i$ such that $\succsim_A$ is nondegenerate, or $\alpha'_{\pi_i} \leq \sum_{\omega\in A} p_i(\omega)$ for some $A \subseteq C_i$ such that $\succsim_A$ is degenerate would result in a contradiction. This completes the proof that $\alpha'_{\pi_i} \leq \alpha_i^{max}$ and $\alpha'_{\pi_i} > \alpha_i^{min}$ for all $i \in \mathcal{S}$. □

Proof of Observation 8. Suppose that there exists $i \in \mathcal{S}$ such that $(\alpha_i^{1\,min}, \alpha_i^{1\,max}] \cap (\alpha_i^{2\,min}, \alpha_i^{2\,max}] \neq \varnothing$. Without loss of generality, assume that $\alpha_i^{1\,min} < \alpha_i^{2\,min} < \alpha_i^{1\,max}$. Let $A \subseteq C_i$ be such that $\sum_{\omega\in A} p_i(\omega) = \alpha_i^{2\,min}$. We have $\sum_{\omega\in A} p_i(\omega) > \alpha_i^{1\,min}$, so it must be that $\succsim_A^1$ is nondegenerate by the definition of α_i^{min}, and, therefore, $\sum_{\omega\in A} p_i(\omega) \geq \alpha_i^{1\,max}$ by the definition of α_i^{max}, a contradiction. □

Proof of Proposition 9. By Proposition 5, Decision Maker 1 is more prone to exclusion errors if and only if Decision Maker 2 is more prone to inclusion errors if and only if

$$\sum_{\omega\in A} p_i(\omega) \geq \alpha_i^1 \Rightarrow \sum_{\omega\in A} p_i(\omega) \geq \alpha_i^2 \qquad \forall_{A\subseteq C_i} \forall_{i\in\mathcal{S}:|C_i|\geq 2}. \tag{12}$$

Clearly, if $\alpha_i^1 \geq \alpha_i^2$ for all $i \in \mathcal{S}$, then (12) holds.

Conversely, suppose that (12) holds and assume, by contradiction, that $\alpha_i^{1\,min} < \alpha_i^{2\,min}$ for some $i \in \mathcal{S}$. Let $A \subseteq C_i$ be such that $\sum_{\omega\in A} p_i(\omega) = \alpha_i^{2\,min}$ and $\succsim_A^2$ is degenerate. As follows from Proposition 7 and Observation 8, it must be that $\alpha_i^{1\,min} < \alpha_i^1 \leq \alpha_i^{1\,max} \leq \alpha_i^{2\,min}$

and, therefore, $\sum_{\omega \in A} p_i(\omega) \geq \alpha_i^1$. By (12), we have $\sum_{\omega \in A} p_i(\omega) \geq \alpha_i^2$ and, therefore, $\succsim_A^2$ is nondegenerate, a contradiction. We conclude that $\alpha_i^{1\,min} \geq \alpha_i^{2\,min}$ for all $i \in \mathcal{S}$, which, given Observation 8, proves the proposition. □

References

Ahn, D. S. and H. Ergin (2010): "Framing Contingencies," *Econometrica*, 78, 655–695.

Ahn, D. S. and T. Sarver (2013): "Preference for Flexibility and Random Choice," *Econometrica*, 81, 341–361.

Anscombe, F. J. and R. J. Aumann (1963): "A Definition of Subjective Probability," *The Annals of Mathematical Statistics*, 34, 199–205.

Darley, J. M. and P. H. Gross (1983): "A hypothesis-confirming bias in labeling effects," *Journal of Personality and Social Psychology*, 44, 20–33.

Dekel, E., B. L. Lipman, and A. Rustichini (1998): "Recent developments in modeling unforeseen contingencies," *European Economic Review*, 42, 523–542.

——— (2001): "Representing preferences with a unique subjective state space," *Econometrica*, 69, 891–934.

Epstein, L. G. (2006): "An Axiomatic Model of Non-Bayesian Updating," *The Review of Economic Studies*, 73, 413–436.

Epstein, L. G., M. Marinacci, and K. Seo (2007): "Coarse contingencies and ambiguity," *Theoretical Economics*, 2, 355–394.

Epstein, L. G. and K. Seo (2009): "Subjective states: A more robust model," *Games and Economic Behavior*, 67, 408–427.

Fishburn, P. C. (1970): *Utility Theory for Decision Making*, New York: Wiley.

Ghirardato, P. (2001): "Coping with Ignorance: Unforeseen Contingencies and Non-Additive Uncertainty," *Economic Theory*, 17, 247–276.

Gilboa, I. (2009): *Theory of Decision under Uncertainty*, Cambridge University Press.

Gilboa, I. and M. Marinacci (2013): "Ambiguity and the Bayesian Paradigm," in *Advances in Economics and Econometrics: Theory and Applications*, ed. by D. Acemoglu, M. Arellano, and E. Dekel, Tenth World Congress of the Econometric Society, Cambridge University Press.

Gneezy, U., M. Hoffman, M. A. Lane, J. A. List, J. A. Livingston, and M. J. Seilier (2016): "Overconfidence or Other Violations of Bayesian Rationality? An Experiment on Belief Updating," Mimeo.

Gul, F. and W. Pesendorfer (2006): "Random Expected Utility," *Econometrica*, 74, 121–146.

Kochov, A. (2015): "A Behavioral Definition of Unforeseen Contingencies," Mimeo, University of Rochester.

Kreps, D. M. (1979): "A representation theorem for 'preference for flexibility'," *Econometrica*, 47, 565–

577.

——— (1992): "Static choice in the presence of unforeseen contingencies," in *Economic analysis of markets and games*, ed. by P. Dasgupta, D. Gale, O. Hart, and E. Maskin, MIT Press Cambridge, 258–281.

Lipman, B. L. (1999): "Decision theory without logical omniscience: Toward an axiomatic framework for bounded rationality," *The Review of Economic Studies*, 66, 339–361.

Manski, C. F. (2004): "Measuring Expectations," *Econometrica*, 72, 1329–1376.

Mukerji, S. (1997): "Understanding the nonadditive probability decision model," *Economic Theory*, 9, 23–46.

Myerson, R. B. (1986a): "Axiomatic Foundations of Bayesian Decision Theory," Center for Mathematical Studies in Economics and Management Science Discussion Paper 671, Northwestern University.

——— (1986b): "Multistage Games with Communication," *Econometrica*, 54, 323–358.

Nehring, K. (1999): "Preference for Flexibility in a Savage Framework," *Econometrica*, 67, 101–119.

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

Rabin, M. and J. L. Schrag (1999): "First Impressions Matter: A Model of Confirmatory Bias," *Quarterly Journal of Economics*, 114, 37–82.

Savage, L. J. (1954): *The Foundations of Statistics*, John Wiley & Sons, New York.

Tversky, A. and D. Kahneman (1974): "Judgment under Uncertainty: Heuristics and Biases," *Science*, 185, 1124–1131.

Wakker, P. P. (1989): *Additive Representations of Preferences: A New Foundation of Decision Analysis*, Kluwer Academic Publishers.

Wason, P. C. (1960): "On the failure to eliminate hypotheses in a conceptual task," *The Quarterly Journal of Experimental Psychology*, 12, 129–140.