

Do Bets Reveal Beliefs?

A Unified Perspective on State-Dependent Utility Issues

Jean Baccelli*

THEMA, Université de Cergy-Pontoise[†]

Abstract

This paper examines the preference-based approach to the identification of beliefs. It focuses on the main problem to which this approach is exposed, namely that of state-dependent utility. First, the problem is illustrated in full detail. Four types of state-dependent utility issues are distinguished. Second, a comprehensive strategy for identifying beliefs under state-dependent utility is presented and discussed. For the problem to be solved following this strategy, however, preferences need to extend beyond choices. We claim that this is a necessary feature of any complete solution to the problem of state-dependent utility. We also argue that this is the main conceptual lesson to draw from it. We show that this lesson is of interest to both economists and philosophers.

Keywords: beliefs, subjective probability, state-dependent utility, hypothetical preferences, revealed preference

The author thanks the Editor-in-Chief Wiebe van der Hoek, two referees, Mikael Cozic, Éric Danan, Vincent Éli, Itzhak Gilboa, Raphaël Giraud, Brian Hill, Marcus Pivato, Fanyin Zheng and especially Philippe Mongin for most helpful comments and discussions. The author is sole responsible for all remaining errors or omissions. Research for this paper was funded by the Université de Cergy-Pontoise, and before that by the École Normale Supérieure - Ulm (grants ANR-10-LABX-0087 IEC and ANR-10-IDEX-0001-02 PSL).

[†]THEMA, Université de Cergy-Pontoise, 33 boulevard du Port, 95011 Cergy-Pontoise, France (jean.baccelli [at] u-cergy.fr).

1 Introduction: the betting approach to the identification of beliefs

Suppose you want to identify the beliefs of an agent regarding the truth of a proposition or the realization of an event. This is sometimes necessary. For instance, if you are interested in opinion aggregation or in any topic studied in interactive epistemology, the identification of individual beliefs is a prerequisite.¹ This paper focuses on such identification. By assumption, it considers this identification as an end in itself and not just as a means to, say, predict the agent's behavior. While we want to stress from the outset that this is an irreducible assumption of our paper, we also want to highlight that this is not an artificial one. For instance, it might be that the agent is an expert and that you want to identify her beliefs not to predict her actions, but to help in your own decision-making.

There are at least two conceivable ways to identify beliefs. To start with, following a direct approach, one can simply ask the agent about the beliefs she holds. An example of this is when judgments of comparative likelihood² are accepted as primitive data. The standard economic approach, by contrast, is indirect. Its goal is to *infer beliefs from preference data*.³ Among other things, this is meant to ensure that the beliefs identified are the ones upon which the agent is ready to act in her own decision-making. Specifically, this approach is essentially inspired by betting behavior. Admittedly, assuming an agent prefers being wealthy to being poor, her willingness to

¹See e.g. Genest and Zidek, 1986 and Clemen and Winkler, 2007 for reviews of the literature on probabilistic opinion aggregation. See footnote 56 regarding the necessity of identifying beliefs when the goal is to aggregate preferences (rather than simply opinions) on uncertain outcomes. Regarding interactive epistemology, see in particular the *agreement theorem* literature, originating in Aumann, 1976 and reviewed e.g. in Bonanno and Nehring, 1997.

²This is the school of Koopman (1940), Good (1950) or Kraft, Pratt, and Seidenberg (1959). Although they all belong to the probabilistic tradition more specifically, we use the term *likelihood* to stay neutral regarding the structure of the beliefs to be identified.

³This is the school of Ramsey (1931), de Finetti (1937), Savage (1954/1972), and followers.

place money on one event rather than another *prima facie* suggests that she thinks the first is more likely to occur than the second. By betting, one tries not to change, but merely to benefit from the way one believes the world goes. For this reason, betting is but a minimal form of action which can be used, when some basic preferences are known, to infer beliefs. In this paper, any preference-based approach to the identification of beliefs will be called a version of *the betting approach*. The betting approach is exemplified in most concrete elicitation techniques,⁴ as well as in the theoretical results these techniques rely on. It is the topic of this paper.

From a decision-theoretic point of view, the betting approach raises specific questions. These questions do not pertain to the structure of the beliefs to be identified. Given a state space, the agent's beliefs may or may not take the form of a unique probability measure depending on, say, her take on Ellsberg's (1961) problem. Nor do those questions regard the decision rule she follows with respect to her beliefs. She may or may not be an expected utility maximizer depending on, say, her reaction to Allais' (1953) problem, when uncertainty reduces to risk. A preliminary issue⁵ to all those familiar questions is whether it is possible that her beliefs are revealed by her preferences, in other words, that her preferences are relied on as a measurement device. For a suggestive counterexample, imagine the following situation. A remote country is holding general elections, and you have incentives to anticipate their outcome correctly. Being unknowledgeable about politics, you seek the opinion of a local expert. She believes that this particular candidate will win. Unbeknownst to you, however, she as an individual voter is a passionate opponent of this candidate. She thinks of his election as an irreversible disaster for her country, so much in fact that she simply does

⁴See e.g. Schlag, Tremewan, and van der Weele, 2015 for a review of those techniques.

⁵Admittedly, there is at least one other issue which can be considered preliminary to the ones above. When preferences are interpreted as choices, the betting approach to the identification of beliefs is confronted with *the problem of moral hazard* (defined in Arrow, 1965, p. 142 and anticipated in Drèze, 1961, p. 78). Section 3.1 introduces this problem and discusses how it relates to that of state-dependent utility, the topic of this paper.

not care about any payoff she could collect in the event that he is elected. Therefore, her betting decisions with respect to this event are bound to lead you astray regarding what she believes, and thus what *you* should believe, about the result of the elections. As we will cash it out formally later on, the key issue is that her utility function (for money, in this example) depends on the prevailing state of nature (specifically, which candidate is elected) in a certain way. In the terminology of this paper, *the problem of state-dependent utility* refers to the fact that the preferences of a decision-maker can impede the identification of her beliefs, as illustrated by this case among other related ones. Noteworthy, state-dependent utility is not always a problem. In insurance economics,⁶ for instance, in line with ordinary thinking, it is common to assume that the decision-maker's utility function depends on her health status, and this is not the source of any particular problem. However, state-dependent utility stands in the way of identifying beliefs following the betting approach.⁷ This is the problem which this paper investigates.

This paper presents the problem of state-dependent utility in full detail, explains how it can be solved, and extracts from it a conceptual message of interest to both economists and philosophers. This last aspect is our main contribution. We show that the problem of state-dependent utility illustrates, so to speak from within, the limits of the *revealed preference* framework which officially prevails in economics. We explain that this should draw the attention of philosophers, especially those thinking about the appropriate way to investigate beliefs and concerned about the potential defects of

⁶The introduction of state-dependent utility in theoretical insurance economics is due to Arrow (see Arrow, 1953, 1973). See e.g. Finkelstein, Luttmer, and Notowidigdo, 2013 for recent empirical work on insurance demand under state-dependent utility.

⁷Relatedly, state-dependent utility is a source of complication for *pragmatic arguments* for probabilism (the claim according to which *rational* beliefs should take the form of a unique probability measure over a given set space), the expected utility rule, and the like. Some such arguments are known as the *Dutch book arguments* (see Hájek, 2008 for a review). It is beyond the scope of this paper to assess how these arguments fare in light of state-dependent utility issues (a topic Nau has investigated, see in particular 1995). Section 3.4, however, relates the conceptual discussion of state-dependent utility issues to the better-known philosophical discussions of pragmatic arguments for probabilism.

preference-based approaches. Our paper proceeds as follows. For the sake of simplicity and without loss of generality (this claim is justified in section 2.6), it focuses on *Bayesian agents*, whose beliefs take the form of a unique probability measure on a given state space and whose decisions are compatible with the expected utility rule. Faced with Ellsberg's and Allais' problems, they have non-paradoxical preferences. The first part of the paper is set in the framework of Savage's theorem. Four types of state-dependent utility issues are distinguished, one of which has not yet been discussed in the literature. Each of these issues is an obstacle for the basic betting approach, which we take Savage's result to embody. Several methodological comments are offered to explain why decision theorists rarely take all of these issues into account. The second part of the paper considers a theorem by Karni and Schmeidler, which illustrates the simplest way to revise the basic betting approach so as to achieve the identification of beliefs. Identification, nevertheless, comes at a price: for beliefs to be identified as Karni and Schmeidler explain, preferences need to extend beyond choices. We claim that such an enlargement is indeed *necessary* to solve the problem of state-dependent utility, which leads to the final discussion of revealed preference previewed above. Thus, technically as well as conceptually, our paper offers a unified perspective on state-dependent utility issues.

2 The problem of state-dependent utility

2.1 Theme: the case of Mr. Smith

Consider the following case, which is famous among decision theorists and can be used to give a comprehensive statement of the problem of state-dependent utility.⁸ Mr. Smith's wife, Mrs. Smith, is gravely ill. She has to undergo a crucial operation that will either kill her or cure her for good.

⁸It was sketched by Aumann in private correspondence with Savage. Thanks to Drèze, the letters have been made public in Savage and Aumann, 1987.

Suppose you want to identify Mr. Smith's beliefs, specifically, the ones he holds regarding the success or the failure of that operation. Imagine for instance he is a surgeon, and he is familiar with the medical procedure. His wife, for one, might be interested in knowing what he, as an expert, thinks of her chances of survival. She might even decide whether or not to undergo the operation based on his opinion on the matter. Assume she asks a decision theorist to apply the betting approach for her. To reveal Mr. Smith's beliefs, the decision theorist will let him place various sums of money either on the success or on the failure of the operation. Mr. Smith is a Bayesian agent: he attributes a certain subjective probability value to the first event and a complementary value to the other event. Yet Mr. Smith is also a husband: his life would be utterly changed, were his wife to die. We will show that this modest fact has far-reaching implications, namely, it determines the degree to which his beliefs can be identified following the betting approach.

To verify this, let us place the case of Mr. Smith in the framework of Savage's (1954/1972) theorem, the prime example of the betting approach in decision theory. Let S be the set of *states of nature*, define an *event* as any subset of this set, and let 2^S be the set of all possible events. Let X be the set of *consequences*. Let a decision-maker be characterized by a preference relation \succsim on $F := X^S$, the set of all possible *acts*, with \succ and \sim referring to strict preference and indifference. Savage's theorem gives sufficient conditions for \succsim to be representable according to the *subjective expected utility model*. Under those conditions, there exists a utility function $u : X \rightarrow \mathbb{R}$, unique up to a positive affine transformation, and a unique subjective probability function $\pi^\succsim : 2^S \rightarrow \mathbb{R}$ such that, for all acts f and g ,

the following equivalence holds:⁹

$$f \succsim g \Leftrightarrow \int_S u[f(s)] d\pi^\succsim(s) \geq \int_S u[g(s)] d\pi^\succsim(s). \quad (1)$$

Suppose then, more precisely, that the decision theorist tries to identify Mr. Smith's subjective probability function π^\succsim following Savage's approach. To specify Mr. Smith's preference relation according to his predicament, we shall focus on a concrete version of Savage's framework. Let X be a set of monetary results, $X = [0, 100]$, and S be the unit interval, $S = (0, 1]$. With $\gamma \in (0, 1]$ varying for the sake of convenience, let $E^* = (0, \gamma]$ correspond to the event that the operation succeeds and $\overline{E^*} = (\gamma, 1]$ to the event that it fails. Considering Mr. Smith as a Bayesian decision-maker, let him be endowed with a unique subjective probability measure on 2^S , denoted by π^* . For the sake of simplicity, let π^* be such¹⁰ that for any event $E = (a, b]$, $\pi^*(E) = b - a$. Considering Mr. Smith as a husband, let his utility function vary depending on the outcome of the operation, which gives a simple illustration of the state-dependent utility phenomena. In what follows, while always assuming that $u_{E^*}(x) = x$, we consider various forms for $u_{\overline{E^*}}(x)$ expressing various ways he might react to the death of his wife.¹¹ Four reactions are examined. They are four variations on one theme, given that our assumptions imply that Mr. Smith's preferences will always be representable by a function $v^* : F \rightarrow \mathbb{R}$ that can be decomposed as follows:

⁹Under those conditions, it is also the case that S is infinite, that π^\succsim is non-atomic and finitely additive, and that u is bounded (see e.g. Fishburn, 1970, chapter 14, and further technical clarifications in Wakker, 1993). These restrictions on u and on π^\succsim indicate that Savage's conditions do not *characterize* all preference relations representable according to the subjective expected utility model.

¹⁰Specifically, we suppose that π^* is a non-atomic finitely additive probability function on 2^S , the σ -algebra of all subsets of $S = (0, 1]$, such that the probability of any interval in $(0, 1]$ simply is its length. It follows from a classical result in Horn and Tarski, 1948 that such a function exists.

¹¹For the sake of concreteness, the case of Mr. Smith is presented with a dynamic twist, but this is inessential to the argument. For instance, you could be with him in the hospital waiting room and willing to identify his beliefs regarding the outcome of the operation on his wife at a time at which this operation has occurred. The analysis would be unchanged.

$$v^*(f) = \int_{s \in E^*} f(s) d\pi^*(s) + \int_{s \in \overline{E^*}} u_{\overline{E^*}}[f(s)] d\pi^*(s). \quad (2)$$

In each variation on $u_{\overline{E^*}}(x)$, by assumption, beliefs are also specified in π^* . The question is whether one can, following the betting approach, identify these beliefs which have been fixed by assumption. This amounts to *checking the reliability of preferences, taken as a measurement device for identifying beliefs*. Admittedly, this is done under the supposition that the framework above accurately describes Mr. Smith's situation. We are well aware that this is a debatable issue, but we postpone the discussion of description or re-description until all four variations have been examined.¹²

2.2 First variation: the cautious husband

Consider first the following reaction. Were his wife to die, Mr. Smith would become more cautious: without her support, he would abandon any ambitious project he would have undertaken otherwise.¹³ From an economic point of view, this can be interpreted as a change in risk attitude. Mr. Smith being a Bayesian decision-maker, his utility function must express this new risk attitude. Let us then assume that $u_{\overline{E^*}}(x) = \sqrt{x}$, for instance. For the sake of convenience, let us also suppose that $\gamma = \frac{1}{2}$, so that $\pi^*(E^*) = \pi^*(\overline{E^*}) = \frac{1}{2}$.

Among the four types of state-dependent utility issues presented in this paper, this one is the best known of all (see e.g. Karni and Schmeidler, 1993, p. 269 for a representative mention). It is also the closest to the form of state-dependent utility commonly considered in the insurance literature.

¹²See section 3.1. Our message at this point is essentially the following: before debating of re-description, let us first have a comprehensive look at *all* there would be to re-describe.

¹³Philosophers might think that this kind of example requires the assumption that, as he evaluates bets, Mr. Smith agrees with his future (more or less cautious) self. First, as we argued in footnote 11, the dynamic twist of the story is conceptually inessential. Second, notice that even if it mattered, no more agreement with one's future self would be required in state-*dependent* utility cases like the one above, than in any state-*independent* utility counterpart (as when Mr. Smith anticipates to stay just the same, were his wife to die). The same observation applies to the other variations on the case of Mr. Smith.

Because of such cautiousness, Mr. Smith’s beliefs cannot be identified following Savage’s approach. Indeed, in his model, it is crucial that bets indicate the same beliefs, whichever the stakes used. Let us express this requirement formally. Abusing notation, let x stand for the *constant act* which associates the consequence x to all states of nature. Let xEy stand for the act giving the consequence x if the true state of nature is in E and giving the consequence y otherwise. For the bets taken by the decision-maker to induce a well-defined *likelihood order* on the set of events, Savage’s fourth postulate requires that for all consequences x, y, x', y' such that $x \succ y$, $x' \succ y'$, and for all events E, E' , we have $xEy \succ xE'y$ if and only if $x'E'y' \succ x'E'y'$. This postulate is the only one of Savage’s postulates that is violated in our case.¹⁴ Despite their being generated by a perfectly well-defined underlying subjective probability function π^* , Mr. Smith’s preferences induce an inconsistent likelihood order on the set of events. Thus it would not be possible to proceed from this likelihood order to any numerical representation of it.

Savage’s fourth postulate is central to his influential approach to the identification of beliefs. It is also very closely related to the motivating interpretation of beliefs in terms of willingness to bet. Yet notice it cannot be considered as a compelling rationality condition. Certainly, it is a consistency condition. But the consistency in question concerns the measurement technique proposed and should not be extended too hastily to the subjective attitudes to be measured following this technique. As the present variation illustrates, the postulate requires *the invariance of risk attitudes*,¹⁵ whichever

¹⁴Consider (2), as fully specified with $u_{\overline{E^*}}(x) = \sqrt{x}$ and $E^* = (0, \frac{1}{2}]$. First, check that $x > y \Leftrightarrow x \succ y$. Then, taking e.g. $x = 100$, $y = 0 = y'$, $x' = 9$, $E = (0, \frac{1}{10}]$, $E' = (\frac{1}{2}, 1]$, check that $x \succ y$, $x' \succ y'$, $xEy \succ xE'y$ but $x'E'y' \prec x'E'y'$.

¹⁵In the Bayesian context we are focusing on here, “risk attitude” refers to the concavity properties of the utility function. From a mathematical point of view, when all the other postulates of Savage’s theorem are respected, Savage’s fourth postulate amounts to requiring that all (non-constant) conditional utility functions are related by a positive affine transformation. This appears more explicitly in a related theorem, due to Anscombe and Aumann (1963). This paper draws several comparisons between Savage’s theorem and the Anscombe-Aumann theorem. The framework of the latter result is adopted in section 3.

the true state of nature. As this phrasing highlights, this is obviously not a requirement of rationality. This implication¹⁶ of Savage’s fourth postulate should be highlighted more often than it is. A pioneer of the betting approach, de Finetti could identify the beliefs of Bayesian agents, provided they were risk-neutral.¹⁷ His successor Savage generalizes de Finetti’s approach, but only partially in this respect. His model can accommodate *any* risk attitude, including strict risk aversion or risk seeking. But it can only accommodate *one* risk attitude, independently of the true state of nature.

2.3 Second variation: the ascetic husband

Let us turn to another possible reaction. Were his wife to die, Mr. Smith would abandon all worldly concerns: he would devote the rest of his life to prayer and meditation. To capture this ascetic stance, let us suppose that $u_{\overline{E^*}}(x) = -x$, for instance. For the sake of convenience, let us also suppose that $\gamma = \frac{2}{3}$, so that $\pi^*(E^*) = \frac{2}{3}$ and $\pi^*(\overline{E^*}) = \frac{1}{3}$.

Such asceticism also stands in the way of the betting approach. As we explain in several steps, this is only partially analyzed in the current literature. It can be proved that Savage’s fourth postulate is respected here.¹⁸ As a result, Mr. Smith’s preferences induce a well-defined likelihood order on the set of events. Yet, the order induced is useless from a *probabilistic* point of view, because it suggests that Mr. Smith holds impossible events to be strictly more likely to happen than some possible events - namely, the

¹⁶This is but an implication of the postulate. The content of the postulate is far more general, in particular because like all of Savage’s postulates, it applies to arbitrary sets of consequences, non-numerical ones included (in which case, risk attitudes cannot be defined).

¹⁷See de Finetti, 1937 (for a modern exposition, see e.g. Gilboa, 2009, chapter 9).

¹⁸See the Appendix for a proof. Taken together, the cautious husband case and the ascetic husband case establish the logical independence of Savage’s third and fourth postulates. The framework of Anscombe and Aumann would prove less flexible than Savage’s on this topic. Endowed with the usual axioms, it would not let the difference between the cautious husband case and the ascetic husband case be fully expressed.

event that the operation fails.¹⁹ As indicated by the function π^* , he believes that this event can occur. But this fact cannot be accurately recognized by Savage’s take on the betting approach. Savage cannot proceed to any probabilistic numerical representation of the likelihood order discovered.

One aspect of this problem is well understood: Mr. Smith’s asceticism violates another of Savage’s postulates which excludes such troublesome cases. Let xEh stand for the act that gives the consequence x if the true state of nature is in E , and that is identical to some other act h otherwise. Define E as *null* for a decision-maker if for all consequences x and y , we have $xEh \sim yEh$, and define it as *non-null* otherwise. For the bets taken by the decision-maker to induce a likelihood order that is a *qualitative probability relation*, Savage requires that preferences conditional on non-null events are related to preferences over constant acts as simply as possible. Specifically, his third postulate requires in particular that, for all consequences x, y , and for any non-null event E , if $x \succ y$, then $xEh \succ yEh$. This postulate is the only one of Savage’s postulates that is violated in our case.²⁰ The following analogy may help in realizing how restrictive it is. In another context, it would say that if a social preference exists, then it has to be unanimous. Crucially, the implication above does *not* correspond to the more familiar and uncontroversial converse implication according to which if there is a unanimous preference, then it has to be the social preference too.

To the best of our knowledge, however, it has not yet been pointed out in the literature that a violation of Savage’s third postulate need not entail a violation of his fourth. This is what our example illustrates, unlike other

¹⁹Consider (2), as fully specified with $u_{\overline{E^*}}(x) = -x$ and $E^* = (0, \frac{2}{3}]$. Taking e.g. $x = 1$, $y = 0$, check that $x \succ y$ and that $x\emptyset y \succ x\overline{E^*}y$. This means that \emptyset is revealed to be held strictly more likely to happen than $\overline{E^*}$. This disqualifies the likelihood order induced by Mr. Smith’s preferences as a qualitative probability relation (as defined in Fishburn, 1970, p. 195). It is the most consensual property of qualitative probability relations, sometimes called “non-negativity”, which fails here to be satisfied.

²⁰Take e.g. $x = 1$, $y = 0$, $E = \overline{E^*}$. First, check that $x \succ y$, and that E is non-null. Then, taking $h = x$ for example, check that $xEh \prec yEh$, which violates the postulate.

more familiar counterexamples to Savage’s third postulate.²¹ More generally, our example suggests that the literature should abandon the terminology according to which Savage’s third and fourth postulates impose, respectively, “ordinal” and “cardinal” forms of state-independence upon the preferences of the decision-maker.²² Specifically, this terminology suggests that if the latter form of state-independence obtains, then the former does as well. As detailed above, this is disproved by the present variation on the case of Mr. Smith.

Presumably, Savage and others have checked the logical independence of the two postulates in question, but they may have thought that the mathematical observation underlying this subsection did not deserve to be reported. We disagree (for reasons that will appear in section 3.4), and we find their omission interesting in its own right. It is best explained, we believe, as follows. First, it is related to the importance of economic interpretations of decision-theoretic frameworks: from the standpoint of economic applications, the problems raised by preferences that would be decreasing in the amount of money received, as in our case of asceticism, are certainly not the most urgent to consider. Second, based on the motivating interpretation of beliefs in terms of willingness to bet, theorists probably reason as follows. There is little point in extracting a comparative likelihood relation if it is to deliver probabilistically unintelligible information, the argument goes, so that Savage’s fourth postulate should be considered only when his third is also satisfied. Yet, in both cases, privileged interpretations of the formalism are put forward, rather than logical arguments. This is illustrative of the way decision theorists often conduct axiomatic analysis: their mathematical work is oriented by pre-existing interpretations of their objects. The omission we are highlighting now is just one consequence of this influence. Digging further into the problem of state-dependent utility, we shall see that

²¹See for instance Savage’s own (meteorological) case in Savage, 1954/1972, p. 25.

²²For this terminology, see e.g. Wakker and Zank, 1999, p. 10, or Hill, 2010, p. 2045. It proves more relevant in the framework of Anscombe and Aumann, in which it originates.

it has other implications which are less apparent and even more significant.

2.4 Third variation: the devastated husband

Consider yet another reaction (which corresponds to the election case sketched in the introduction). Were his wife to die, Mr. Smith would be devastated: come what may, nothing would matter to him anymore. Let us suppose accordingly that, whichever the consequence x , $u_{\overline{E^*}}(x) = c$, with c some constant. For the sake of convenience, let us also suppose that $\gamma = \frac{1}{2}$, so that $\pi^*(E^*) = \pi^*(\overline{E^*}) = \frac{1}{2}$.

In this case, Mr. Smith's preferences satisfy all of Savage's postulates, and the betting approach does not run into any apparent obstacle. Yet, another problem arises, which is less familiar but not less challenging than the two previous ones (see e.g. Drèze, 1987, p. 67 or Karni, Schmeidler, and Vind, 1983, p. 1025 for discussions). While Mr. Smith's intimate conviction is that $\pi^*(E^*) = \pi^*(\overline{E^*}) > 0$, Savage must infer from his preferences that $\pi^{\succ}(E^*) > \pi^{\succ}(\overline{E^*}) = 0$. Indeed, in the present variation, the event that the operation fails qualifies as a null event and, in Savage's theorem, null events must be associated with a null subjective probability value.²³ By contrast, this variation illustrates the fact that, even under the constraint of all of Savage's postulates, this is not an *implication* of Mr. Smith's preferences but merely one interpretation which is *compatible* with them. Even when all of Savage's postulates are satisfied, it cannot be decided whether preferences conditional on null events should be decomposed as Savage suggests, as the product of a null subjective probability value and a non-constant utility function, or rather as the product of a strictly positive subjective probability value and a conditionally constant utility function. For terminological clarity,

²³Consider (2), as fully specified with $u_{\overline{E^*}}(x) = c$ and $E^* = (0, \frac{1}{2}]$. Taking e.g. $x = 1$ and $y = 0$, check that $x E^* y \succ x \overline{E^*} y$ and that $x \overline{E^*} y \sim x \emptyset y$. Noteworthily, in Savage's own presentation of his theorem, the claim that an event is null if and only if it is associated with a null subjective probability value is stated on a par with the general representation in (1).

let us refer to the former case as that of *cognitively null* events (a matter of probability, in the representation), and to the latter case as that *conatively null* events (a matter of utility, in the representation).²⁴ In effect, Savage claims that all null events are of the cognitive variety. Yet, in his framework, it cannot be decided whether a null event is of the cognitive or of the conative variety, because the same preferences are induced in both cases.

We shall refer to this problem as *the problem of null events*. Unlike the two previous issues, but similar to the one presented next, it does not pertain to the existence of a revealed qualitative probability relation. From now on in this section, such an existence is secured. As a result, it will always be possible to identify a subjective probability function following Savage's approach. Yet, it is still to be checked whether this subjective probability function really is the only one compatible with the underlying preferences that have been first given. Thus, what is fundamentally at issue is whether the qualitative probability relation induced by Mr. Smith's preferences should be considered, as Savage assumes, as a relevant guide to his beliefs.

The problem of null events is reminiscent of underdetermination problems which *non-triviality* suppositions, like Savage's fifth postulate,²⁵ are meant to rule out. A decision theorist cannot tell much of an agent's beliefs if the agent is indifferent between all available options: such preferences are compatible with any set of beliefs. Accordingly, agents are typically required to express one strict preference at least. But what is true of the *total* event is also true of any *partial* event. The only difference is that it would be overly restrictive²⁶ to require that, conditional on any partial event, the agent express a strict preference. This would simply rule out the case that any possible event is null and thus that any event other than the impossible ones, represented by

²⁴It is readily checked that the two concepts are logically independent.

²⁵This postulate requires that there exists consequences x, y such that $x \succ y$.

²⁶Besides, it would be impossible in Savage's model because his postulates imply that all *states*, viz. degenerate events, are null (this is why it is required that there are infinitely many of them, so that the non-triviality postulate above can nonetheless be satisfied).

the empty set, is associated with a null subjective probability value.

To acknowledge the structural underdetermination illustrated by this variation on the case of Mr. Smith, the uniqueness clause of Savage’s representation theorem must be revised.²⁷ Assume null events are the only ones to raise underdetermination issues regarding subjective probability values.²⁸ Then, in (1)²⁹, π^{\succsim} is replaceable by another probability function π'^{\succsim} if and only if for all non-null events E, E' , $\frac{\pi'^{\succsim}(E)}{\pi'^{\succsim}(E')} = \frac{\pi^{\succsim}(E)}{\pi^{\succsim}(E')}$. In other words, for any null event E'' , the probability value $\pi'^{\succsim}(E'')$ can be chosen arbitrarily. The representation of \succsim can always be preserved because, for any null event E'' , a conditionally constant utility function $u_{E''}(\cdot) = c_{E''}$ can also be introduced, which will induce the same preferences. Thus it appears that, even when all of his postulates are satisfied, Savage cannot claim to identify only one subjective probability function π^{\succsim} . Admittedly, the problem of null events has a special status among the various state-dependent utility issues. Yet null events are not the only ones to raise such underdetermination issues. In a nutshell: there is also a *problem of non-null events*, to which we turn now.

2.5 Fourth variation: the downhearted husband

Consider this last reaction. Were his wife to die, Mr. Smith would be downhearted: he would carry on with his life, but enjoy every bit of it less. Let us suppose accordingly that $u_{E^*}(x) = \frac{1}{2}x$. For the sake of convenience,

²⁷The kind of uniqueness clause to follow is not standard in discussions of Savage’s theorem or related results. In this context, the only explicit statement we are aware of is to be found in Wakker, 1987, p. 293. Albeit in a different analytical framework, it has also been considered in recent models of *awareness* (see Karni and Vierø, 2013, p. 2802). Essentially, this clause describes a renormalization of the probability values of non-null events. It amounts to making these values unique *as on a ratio scale*, rather than *absolutely* unique as they normally are.

²⁸We stress that this uniqueness clause is targeted at the problem of null events specifically. In effect, the “only if” direction of the proposition below will be questioned by the next variation on the case of Mr. Smith.

²⁹The new probability function π'^{\succsim} would contribute to representing the same underlying preferences, as stated in (1). But it would not represent the qualitative probability relation those preferences induce, following Savage’s construction. This is precisely this route to the identification of beliefs that is being questioned: the *existence* of an induced qualitative probability relation is one thing, its *relevance* is another.

let us also suppose that $\gamma = \frac{1}{2}$, so that $\pi^*(E^*) = \pi^*(\overline{E^*}) = \frac{1}{2}$.

Again, Mr. Smith's preferences satisfy all of Savage's postulates, and the betting approach does not run into any apparent obstacle. Yet, while Mr. Smith's expert opinion is that $\pi^*(E^*) = \pi^*(\overline{E^*})$, Savage must infer from his preferences that $\pi^{\succ}(E^*) > \pi^{\succ}(\overline{E^*})$.³⁰ As in the previous variation, the point is that even though all of Savage's postulates are satisfied, Mr. Smith's preferences can be represented by several combinations of probability and utility values. Undeniably, the particular aspect of the state-dependent utility problem focused on here is the less known of all (nevertheless, see Schervish, Seidenfeld, and Kadane, 1990, section 5 and Karni, 1996, section 2.1³¹ for discussions).

It may come as a surprise that, under the stringent constraint of Savage's postulates, non-null events raise underdetermination issues similar to the ones previously presented with respect to null events. The reason is that the uniqueness of π^{\succ} in (1) is not *absolute*, but *relative* to the utility function that is chosen. And while the respect of Savage's postulates implies that there exists a state-independent utility function $u : X \rightarrow \mathbb{R}$ such that the decision-maker's preferences can be represented according to the subjective expected utility model, it does not imply that u is the only admissible function here. If u is admissible, then *infinitely many* state-dependent utility

³⁰Consider (2), as fully specified with $u_{\overline{E^*}}(x) = \frac{1}{2}x$ and $E^* = (0, \frac{1}{2}]$. Taking any consequences x, y such that $x \succ y$, check that $x E^* y \succ x \overline{E^*} y$.

³¹Ever since this paper, Karni has tirelessly tried to make decision theorists aware of the issue and to solve it one way or another. This paper is the closest to ours in the literature. Unlike our paper, however, it focuses on this particular issue only, instead of discussing together all state-dependent utility issues.

functions $u' : X \times S \rightarrow \mathbb{R}$ are also admissible³² inasmuch as that, together with some jointly unique probability function π'^{\succsim} , they induce the same preferences. The respect of Savage's postulates implies that it is always *possible* to *normalize* state-dependent utilities to a state-independent function, not that it is *legitimate* to do so. The underlying preferences being the same in all cases, they cannot favor any specific decomposition over the others.

The devastated husband case and the downhearted husband case illustrate two sides of one issue, the first focusing on null events, the second on non-null events. It is sometimes presented as *the problem of state-dependent utility without state-dependent preferences*³³ (so that one might also say, by contrast, that the cautious husband case and the ascetic husband case illustrate *the problem of state-dependent utility with state-dependent preferences*). Taken together, these last two variations indicate that Savage's third and fourth postulates are of no help in solving the belief identification problem. Indeed, even when they are satisfied, only *one* unquestionable implication can be drawn from the agent's preferences: if an event is revealed non-null, then it must be true that the agent associates a non-null subjective probability value to it. As the variations above illustrate, no other conclusion is

³²The proof of the Anscombe-Aumann theorem illustrates the fact that a collection of state-dependent utility functions is admissible here if and only if all (non-constant) state-dependent utility functions are related by a positive affine transformation. It is then always possible to *interpret* the transformation coefficients as probability weights. Among others, Karni often makes the following algebraic observation (see e.g. Karni and Mongin, 2000, p. 238) which is suggestive of how arbitrary such a decomposition is. Consider a preference relation \succsim that is represented according to the subjective expected utility model, with π the probability function on a (for the sake of convenience, finite) state space S , and u the utility function. Pick an *arbitrary* function $a : S \rightarrow \mathbb{R}^{+*}$. Define a new utility w based on $w_s(x) = u_s(x) / a(s)$, for all x and all s , and a new probability σ based on $\sigma(s) = \pi(s)a(s) / \sum_{t \in S} \pi(t)a(t)$, for all s . It is readily checked that the product of σ and w represents the same preferences. Focusing on the probability function in all these representations, this illustrates the extent to which preferences underdetermine beliefs.

³³See e.g. Karni, 1996, p. 259. In light of this issue, there is no reason to think that the Ramseyan approaches to the identification of beliefs (see Ramsey, 1931 and e.g. Bradley, 2004 for a modern exposition) are not exposed to the problem of state-dependent utility. More precisely, either these approaches are not able to express this problem (this will be the case if their framework is such that no consequence is available in two different states of nature, see Bradley, 2004, p. 494 for a discussion), or they are also exposed to it. Either way, they do not solve this problem.

robust, in particular none of the form “the decision-maker believes the event E is more likely to occur than the event E' ”. In other words, the accurate uniqueness clause of Savage’s representation theorem is in fact the following. In (1), π^{\succsim} is replaceable by another probability function π'^{\succsim} if and only if, for any event E , if E is revealed non-null, then $\pi'^{\succsim}(E) > 0$. Under this weak constraint, it will always be possible to preserve the representation of \succsim by appropriately adjusting the utility side of it. Clearly, such minimal identification of beliefs could be achieved even without Savage’s third and fourth postulates.

The extensive underdetermination stated above is rarely recognized. Arguably, again, this is best explained by highlighting how mathematical analysis is oriented by pre-existing interpretations. For the motivating semantics of beliefs in terms of willingness to bet to make sense, so that preferences deliver beliefs directly as Savage’s fourth postulate suggests, utility needs to be state-independent. Accordingly, whenever such a state-independent utility is axiomatically available, decision theorists focus on it, at the risk of misstating the uniqueness of the representations they propose. They endorse the *implicit* assumption of state-independent utility, which entails but exceeds *explicit* assumptions of state-independent preference such as Savage’s third and fourth postulates. Here more than elsewhere, it is hard to argue for this assumption on methodological grounds of simplicity. Simplicity considerations are especially convincing when the question is whether to keep some property that is crucial from a theoretical point of view, but questionable from an empirical point of view.³⁴ Notice that this case does not apply here. First, as our last two variations illustrate, some forms of state-dependent utility are empirically indistinguishable from state-independent utility, in the sense that the underlying preferences are the same. Accordingly, the state-independent utility assumption has a non-empirical component, and it

³⁴In decision theory under risk, the respect of *von Neumann and Morgenstern independence* is sometimes defended against the Allais-type behavior on such grounds of simplicity.

is questioned here not as an empirical approximation, but as a metaphysical commitment of sorts. Second, the state-independent utility assumption is all the more resistible, that only the betting approach seems to need it. As mentioned in the introduction with a reference to insurance economics (to which we related our first variation), state-dependent utility is well established elsewhere, essentially, whenever the main problem is not to identify beliefs. In other words, state-dependent utility is not an artifact of this problem and, considering economics as a whole, state-independent utility is not a crucial property from a theoretical point of view.

2.6 Coda: the generality of the problem

Because of the various state-dependent utility issues, the betting approach fails to identify Mr. Smith's beliefs which are given by assumption in π^* . Following Savage's method, there is either no function π^\succsim , or too many functions π^\succsim suitably compatible with a given preference relation \succsim . The following table summarizes the essential aspects of the previous variations.

VARIATION	STATE-DEPENDENT UTILITY ISSUE ILLUSTRATED
1 : $u_{E^*}(x) = \sqrt{x}$	\succsim does not induce a likelihood order on the set of events
2 : $u_{E^*}(x) = -x$	\succsim does not induce a qualitative probability relation on the set of events
3 : $u_{E^*}(x) = c$	\succsim underdetermines the subjective probability of null events
4 : $u_{E^*}(x) = \frac{1}{2}x$	\succsim underdetermines the subjective probability of non-null events

SUMMARY OF THE FOUR VARIATIONS ON THE CASE OF MR. SMITH

The problem of state-dependent utility is not a problem for Savage only. Much of what has been said applies outside his framework. First, the problem has been presented with respect to a Bayesian decision-maker, but it could be presented with respect to non-Bayesian ones as well. Consider for instance the existence issues illustrated by our first two variations. Savage's third

and fourth postulates are present in most representation theorems of models of decision-making under uncertainty designed to accommodate Allais' or Ellsberg's paradox.³⁵ Besides, the issues illustrated by our last two variations also apply in these cases. The problem of state-dependent utility proves orthogonal to familiar decision-theoretic disputes about Bayesianism.

Second, the problem is relevant not only for decision theorists, but also for practitioners who want to identify beliefs. This is because their concrete elicitation techniques eventually rely on results just like Savage's theorem.³⁶

Third, the problem of state-dependent utility should concern not only anyone interested in identifying beliefs, but also anyone interested in identifying the utility function characterizing a decision-maker. The crux of state-dependent utility issues is precisely that these two questions are inseparable, and the four variations on the case of Mr. Smith could have been presented on the utility side rather than on the probability side. As a result, these issues are of interest to insurance professionals for more than one reason. On the one hand, such professional need to identify utility to draw specific implications from general insurance economics results. On the other hand, they might also need to identify beliefs, as in circumstances in which the use of actuarial statistics is problematic. This can be the case, in particular, when so-called *catastrophic* or *extreme* risks are at stake.³⁷ Such risks pertain to rare events on which limited statistical data is available, so that the insurer usually has to rely on expert opinion. On a different score, the little

³⁵For a model accommodating Allais' paradox and relying on Savage's third and fourth postulates, see e.g. Machina and Schmeidler, 1992. For a model accommodating Ellsberg's paradox and relying on these postulates as well, see e.g. Gilboa, 1987 (Savage's third postulate is weakened in this case, but the key implication highlighted in section 2.3 is still in place).

³⁶Thus, the problem of state-dependent utility is yet another topic of discussion for the methodological literature on *scoring rules*. Deviations from risk-neutrality already prove challenging for standard scoring rules (for a presentation of the problem, see e.g. Schlag, Tremewan, and van der Weele, 2015, section 2.4, and see Karni, 2009 for a solution). State-dependent utility, which includes *variations* in risk attitudes as illustrated in section 2.2, is therefore even more challenging for these rules (see Karni, 1999 for a partial investigation).

³⁷Extreme risks, such as exceptional natural cataclysms, epidemics, or terrorist attacks, raise distinctive problems for the economic theory of insurance (see e.g. Gollier, 1997).

statistical data available typically is a poor guide to the beliefs guiding the insuree's behavior regarding the risks which the insurer can offer to cover.

3 Solving the problem of state-dependent utility

3.1 Facing the problem of state-dependent utility: an overview

The problem of state-dependent utility refers not to one of the issues considered in the previous section, but to all of them taken together. We now give a brief overview of the four most significant responses to this problem which have appeared in the literature. Our main purpose is to explain why the rest of our paper focuses on only one of those responses. Incidentally, we want to highlight some features this response shares with those other proposals.

Let us first mention a radical reaction to state-dependent utility issues, which several decision theorists have explored.³⁸ It is based on the following fact. Although state-dependent utility impedes the identification of subjective probability, a weaker form of identification might nonetheless be achieved, and it suffices to predict many aspects of individual behavior. Specifically, under some conditions, preferences over uncertain prospects allow representations which, like (1), are additively separable over the state space, but which, unlike (1), do not propose any separation of utility from probability. The conditions at stake are satisfied by Bayesian decision-makers. They are weaker than those of Savage's theorem or related results. In particular, Savage's third and fourth postulates or related axioms can be dispensed with. For terminological clarity, we say that in these cases, preferences are representable by a collection of *state-indexed value* functions. From the standpoint of interpretation, given a state and a consequence, a

³⁸See in particular Wakker and Zank, 1999 and Hill, 2010. As these references discuss, it is particularly challenging to obtain a general additively separable representation in Savage's framework (recall that, in this framework, the state space must be infinite and notice that, without a probability function, integration is undefined). As the rest of our paper illustrates, this is a much simpler task in the framework of Anscombe and Aumann.

value is the unanalyzed product of the decision-maker’s beliefs regarding the occurrence of this state, and the desirability for her of this consequence in this state. Remarkably, such values suffice to predict many aspects of behavior relevant to economics.³⁹ The response proposed is to make do with the identification of such values. As announced above, this is a radical response to the problem of state-dependent utility, because it amounts to abandoning the goal of identifying beliefs to replace it with that of predicting behavior. As our paper explicitly endorses the former goal, by contrast with the latter, we consider this response as unsatisfactory.⁴⁰ One needs to find a way to further analyze the kind of additively separable representations mentioned in this paragraph, so as to separate out their probabilistic component.

Another typical reaction to our problem is to argue that cases like those presented in the previous section are *misdescribed* and that, appropriately re-described, they would cease to display any form of state-dependent utility. Specifically, the argument is that unlike in our presentation, the consequences which Mr. Smith contemplates are not just sums of money, such as “\$50”, but rather sums of money in a given state, such as “\$50 and the operation has failed” or “\$50 and the operation has succeeded”. His case should be remodeled accordingly, the argument goes, before it is examined in Savage’s framework. Notoriously, this is Savage’s own reaction to the case of Mr. Smith and related puzzles.⁴¹ Formally, it consists in redefining the consequence set X as $X \times S$, and in applying the whole approach to this new consequence set. Trivially, the violations of Savage’s third and fourth postulates sketched in the previous section would thus be dissolved, which seems

³⁹See Nau, 2001 for a detailed defense of this claim.

⁴⁰We would dismiss by the same argument the counter-objections to our last two variations on the case of Mr. Smith according to which, when all of Savage’s postulates are satisfied, the state-independent decomposition in (1) is as good as any alternative state-dependent decomposition because it is sufficient to predict individual behavior accurately. This last point is correct (although it is important that it can be incorrect in some dynamic strategic settings, as detailed in Karni, 2008). Yet, if the goal is to identify beliefs, this observation is a non-starter and the problem of state-dependent utility remains open.

⁴¹See Savage, 1954/1972, p. 25 and Savage and Aumann, 1987, pp. 78-80.

to re-open the way for the identification of a subjective probability. Nevertheless, this approach has several unattractive features. On the one hand, it generates infinitely many *impossible acts* in $(X \times S)^S$, the new set of acts, as for example the constant act giving the consequence “\$50 and the operation has succeeded” in all states, including those in which the operation fails. On the other hand, this approach overgeneralizes by excluding *any* conceivable instance of state-dependent utility. This includes any case that could violate Savage’s third and fourth postulates, while instead of being directly built in the formalism, they are precisely left as “postulates”, i.e., conditions on preferences that might not obtain. This also concerns any uncontroversial case studied elsewhere, e.g. in insurance economics, where the main topic is not the identification of beliefs. What is needed, however, is a solution to the problem of state-dependent utility which lets state-dependent utility be expressed. The response sketched in this paragraph does not meet this requirement. Notice this is not because it introduces the set $X \times S$ as such, but because it takes this set as the consequence set for another application of Savage’s approach.

A far more ambitious response to the problem of state-dependent utility is based on the consideration of moral hazard. Arguably, the betting approach is exposed not only to the problem of state-dependent utility, but also to *the problem of moral hazard*, understood as follows in our context. By her actions, a decision-maker might influence the likelihood of some events occurring, which complicates the identification of her beliefs regarding these events. For a suggestive example, imagine that you are facing an archer and that you want to identify his beliefs regarding his capacity to hit various targets.⁴² When offered a bet such that he receives \$100 if he misses a given target, nothing if he hits it, he will have an incentive to miss it deliberately and to accept this bet more readily than other bets. As a result, it will be

⁴²This is essentially Drèze’s seminal example, see e.g. Drèze, 1987, p. 25.

particularly difficult to infer from the bets he takes the beliefs he holds on the events of interest. The moral hazard approach to the identification of beliefs⁴³ aims at solving *at once* the problem of moral hazard and the problem of state-dependent utility. This demands analytical frameworks which are significantly different from Savage's or related ones. Indeed, in the new approach, the decision-maker is not just a bettor as we initially defined it, i.e., someone trying to benefit from the way the world goes, but a full-fledged agent that has some capacity to change the way the world goes. This cannot be integrated to the traditional frameworks merely by revising the conditions imposed on preferences, such as Savage's third and fourth postulates. This needs to be reflected in a new formalism. We also want to stress that the problem of state-dependent utility can be solved following this approach *only if* the problem of moral hazard can also be solved. The two problems, however, appear conceptually independent. For example, it is one thing that the archer cheats by missing the target deliberately (a moral hazard issue), it is another that he is a bad loser and disregards any stakes he could win, were he to miss the target (a state-dependent utility issue). What we need here is a solution to the problem of state-dependent utility *specifically*, i.e., one that is available even when the decision-maker has no influence on the events regarding which his beliefs are of interest. Mr. Smith, for one, has no bearing on whether Mrs. Smith will survive the operation. A similar remark would apply to an expert on macroeconomic conditions, for instance.

There is one approach in the literature which meets all the requirements above, that is, it focuses specifically on the problem of state-dependent utility, leaves room for state-dependent utility, and aims at identifying beliefs. It is *the hypothetical preference approach*, which leads back to a result by Karni

⁴³See Drèze, 1961 for the pioneering version of this approach, Karni, 2011a,b for a more recent version. A key difference between these models is that, unlike the former, the latter explicitly articulates the actions by which the agent affects the likelihood of events. This leads to significantly different representations and interpretations of subjective probability.

and Schmeidler.⁴⁴ In fact, as it will appear more explicitly in section 3.4, we have an inclusive understanding of the hypothetical preference approach, which goes beyond this particular theorem and its immediate generalizations. However, for the sake of concreteness and without loss of generality regarding the point we want to make, this is the result which we present in some detail. First, we sketch it technically. Second, we examine the concept of preference underlying it. Third, we extract a philosophical message from the way it bypasses state-dependent utility issues.

3.2 What the hypothetical preference approach allows

Mr. Smith is a Bayesian decision-maker. Accordingly, relying on his well-defined conditional preferences, it should be possible to provide an additively separable representation of his preferences. Following the terminology previously introduced, his preferences should be representable by a collection of state-indexed value functions. Assume further that Mr. Smith's state-(in)dependent utility function is identified. Then, his subjective probability function could be unambiguously factored out from such state-indexed value functions. In essence, this is the strategy of the Karni-Schmeidler theorem, the simplest of all the results following the hypothetical preference approach.

This result is established in three steps. The first step is set in a variant of Savage's framework, popularized by the Anscombe-Aumann theorem.⁴⁵ In this setting, besides the first and main source of uncertainty, the natural one which is assumed to be the object of the decision-maker's beliefs, a second

⁴⁴It is stated in print as Theorem 1.4 in Karni, 1985. This result has been significantly generalized in different directions (see Karni, 2003 and Grant and Karni, 2004). These generalizations, however, do not matter for the philosophical discussion to come.

⁴⁵From a mathematical point of view, it is also possible to think of Savage's framework as a variant of Anscombe and Aumann's framework in which the consequences are all degenerate lotteries. In the rest of this paper, for the sake of convenience, S is assumed to be finite (see Fishburn, 1970, p. 179, for a generalization of the Anscombe-Aumann theorem when S is infinite). As mentioned before, this assumption is not compatible with Savage's postulates, but it is compatible with the new set of postulates which we need from now on.

source of uncertainty is introduced. The new source is independent from the natural one, and it follows a probabilistic law that is known to the decision-maker. The combination of the two sources provides much structure to the options she is presented with. Specifically, consequences are now *lotteries* that should be played, once the true state of nature is revealed, for the uncertainty she is facing to be fully resolved. Accordingly, let us keep the same notation as before but interpret now the set X of consequences as the set $\Delta(Z)$ of probability distributions with finite support in some set Z , for instance the real interval $[0, 100]$ still. Being asked to give a preference relation \succsim on $F := X^S$, Mr. Smith is asked to decide between options such as “a lottery $(\frac{1}{2} : \$100, \frac{1}{2} : \$0)$ if the operation succeeds, \$50 otherwise” and “\$50 if the operation succeeds, a lottery $(\frac{1}{2} : \$100, \frac{1}{2} : \$0)$ otherwise”. Under some conditions that express nothing more than his Bayesianism and do not bear on state-dependent utility issues,⁴⁶ it is possible to provide an additively separable representation of the relation \succsim .

Second, Mr. Smith is asked to express preferences on options of a new kind. For his state-(in)dependent utility function to be identified as such, it is necessary that his beliefs are somehow neutralized. This is what motivates the introduction of the following special objects, which are tailored to that effect. Consider *hypothetical outcomes* such as “\$50 and the operation has succeeded”. Notice they are distinct from more familiar *conditional outcomes* such as “you receive \$50 *if* the operation succeeds (and nothing otherwise)”. Mr. Smith is asked to express preferences not only between such hypothetical outcomes, but more generally between the lotteries over these outcomes. He is asked to give a preference relation $\widehat{\succsim}$ on the set $\Delta(S \times Z)$. Notice that this set $\Delta(S \times Z)$ is *not* the set $(S \times Z)^S$, which would be the analogue in

⁴⁶They are the conditions of the von Neumann - Morgenstern theorem. Nevertheless, in the proof of the representation theorem, another supposition is needed. It can be interpreted as a form of indifference regarding the order in which uncertainty is resolved. Under this interpretation, it has been criticized by Drèze in relation with the problem of moral hazard (see e.g. Drèze, 1987, p. 27).

our context of Savage’s contentious set of (mostly) impossible acts. Unlike $(S \times Z)^S$, $\Delta(S \times Z)$ contains no self-contradictory option. Upon deciding which of the lotteries in $\Delta(S \times Z)$ he prefers, Mr. Smith can focus on only one aspect, namely, the utility for him of any consequence x in any state s . The probability that a state obtains is given by assumption in any such option. Hypothetical lotteries are not *uncertain*, but *risky* prospects. As a result, Mr. Smith does not have to draw on his own beliefs regarding the likelihood of the possible states of nature. By construction, they are somehow neutralized. Assume that, faced with hypothetical lotteries, Mr. Smith has Bayesian preferences, just like he does when faced with ordinary lotteries. Then it is possible to represent the relation $\widehat{\succsim}$ according to the expected utility model, and it is easy to make a state-(in)dependent utility function appear explicitly in this representation. Admittedly, this function might prove state-independent. Yet, this would be nothing more than a particular case, moreover, one that would be deduced rather than assumed.

Third, the two relations \succsim and $\widehat{\succsim}$ are not necessarily compatible. Specifically, there are well-defined conditional preferences in both cases and these conditional preferences are comparable to some extent. But they might not agree. Nevertheless, if they do,⁴⁷ then it is possible to distinguish two com-

⁴⁷This is what is required by the key *linkage* postulate of the Karni-Schmeidler theorem, which consists in coordinating the conditional preferences of \succsim and $\widehat{\succsim}$. Some notation is needed to introduce it formally. Denote by \widehat{F} the set $\Delta(S \times Z)$. For a generic element \widehat{f} of \widehat{F} , denote by $\widehat{f}(s, z)$ the probability value associated by \widehat{f} to the hypothetical outcome (s, z) . Denote by $\widehat{\succsim}_s$ the hypothetical preference of the decision-maker conditional on state s obtaining, defining it as follows: $\widehat{f} \widehat{\succsim}_s \widehat{g}$ if $\widehat{f} \widehat{\succsim} \widehat{g}$, with \widehat{f} and \widehat{g} such that for all $t \neq s \in S$, and all $z \in Z$, $\widehat{f}(t, z) = \widehat{g}(t, z)$. (Given that $\widehat{\succsim}$ respects the so-called von Neumann - Morgenstern independence, such conditional preferences are well-defined.) Denote by \widehat{F}^* the set of all elements of \widehat{F} , the marginal probability of which have full support on S . On the other hand, consider F and the relation \succsim which is defined over it. For a generic element f of F , denote by $f_s(z)$ the probability value associated by f to outcome z in state s . Denote by \succsim_s the preference of the decision-maker on $\Delta(Z)$, conditional on state s obtaining, defining it as usual. A bijection H can be defined between \widehat{F}^* and F , as follows: for all \widehat{f}^* , let $H(\widehat{f}^*)$ denote the element $f \in F$ such that for all $z \in Z$, and all $s \in S$, $f_s(z) = \widehat{f}^*(s, z) / \sum_{y \in Z} \widehat{f}^*(s, y)$. The key linkage postulate of the Karni-Schmeidler theorem (labeled “strong consistency” by the authors) requires that for all $\widehat{f}^*, \widehat{g}^* \in \widehat{F}^*$, and for any non-null $s \in S$, we have $\widehat{f}^* \widehat{\succsim}_s \widehat{g}^*$ if and only if $H(\widehat{f}^*) \succsim_s H(\widehat{g}^*)$.

ponents in the state-indexed value functions of the initial representation: one is the state-(in)dependent utility coming from the auxiliary representation, the remaining component corresponds to the subjective probability.

This is a new approach to the identification of beliefs. It differs from Savage’s or Anscombe and Aumann’s approach. Preferences are not assumed to deliver beliefs directly, as Savage’s fourth postulate suggests. They are barely assumed to contain the right information. More generally, the subjective probability measure does not emerge from *cross-state*, but from *within-state* comparisons, namely, the comparison of state-indexed value and state-dependent utility functions. Given this is how the hypothetical preference approach defines subjective probability, it can be applied to circumvent almost all of the state-dependent utility issues listed in the previous section. Following this approach, it is possible to identify the beliefs of cautious, ascetic or downhearted husbands. Admittedly, it is not possible to provide a complete identification of a devastated husband’s beliefs. But a major progress is made on this case too. Using hypothetical preferences, considering any null event, it is always possible to decide whether it qualifies as cognitively null or as conatively null. Indeed, while this is undecidable in most models, Karni and Schmeidler’s framework is rich enough for the following definitions to be articulated. Denote by \succ_E the preference conditional on event E occurring, and similarly with $\widehat{\succ}_E$, the hypothetical preference conditional on E occurring. By definition, E is null if $\succ_E = \emptyset$. Take such an E . If it is also the case that $\widehat{\succ}_E = \emptyset$, define E as conatively null. Otherwise, define it as cognitively null. The following table⁴⁸ summarizes the conceptual refinement proposed.

⁴⁸It is adapted from Karni, Schmeidler, and Vind, 1983, p. 1025. These authors, however, have a terminology less informative than ours on this topic. They oppose “evidently null” events to the “indeterminate case[s]”. They would be more specific in opposing “cognitively” null events to “conatively” null ones instead.

	$\widehat{\succ}_E = \emptyset$	$\widehat{\succ}_E \neq \emptyset$
$\succ_E = \emptyset$	E is conatively null	E is cognitively null
$\succ_E \neq \emptyset$	- this case is excluded by the axioms -	E is non-null

HYPOTHETICAL PREFERENCES AND THE PROBLEM OF NULL EVENTS

Notice that, accordingly, it becomes possible to assume away conatively null events by a specific axiom without, more indiscriminately, requiring that all events are non-null. If conatively null events are allowed and if there are some such events for the decision-maker, then it can be not merely conjectured, but demonstrated that her preferences underdetermine her beliefs. Furthermore, using the uniqueness clause that has been presented before, this underdetermination can be made precise. This is the only kind of underdetermination that can remain, if the hypothetical preference approach is followed. Accordingly, we propose to say that beliefs are then identifiable *up to a conatively null event (which can be categorized as such)*.

3.3 What the hypothetical preference approach requires

Such success comes at a price. First, there is loss of generality with respect to Savage's original framework, because beliefs are identified thanks to *exogenous probabilities*. This refers not only to the elements of $\Delta(Z)$, but also to the elements of $\Delta(S \times Z)$.⁴⁹ Savage committed to working with his bare hands. This constraint is not respected any more in Karni and Schmeidler's approach. Second, more importantly, beliefs are identified thanks to *hypothetical preferences*, the operational meaning of which is doubtful. This

⁴⁹However, those two kinds of exogenous probabilities are unequally essential to the hypothetical preference approach. As illustrated in Karni, 2003 following previous work by Wakker, under some conditions, the preliminary additively separable representation can be obtained without exogenous probabilities of the first kind. Exogenous probabilities of the second kind, by contrast, are instrumental in neutralizing the decision-maker's beliefs and identifying state-(in)dependent utility. It is more difficult to imagine how they could be dispensed with (but see the alternative approach mentioned in footnote 57).

accusation is worth being detailed, as follows.⁵⁰ Each element \hat{f} of $\Delta(S \times Z)$ induces a marginal probability $\pi^{\hat{f}}$ on S . Undoubtedly, the intended interpretation is that when a decision-maker evaluates \hat{f} , she should adopt $\pi^{\hat{f}}$ as a *hypothetical belief* on the state of nature. Yet, if this is correct, what does it mean to prefer, say, \hat{f} over \hat{g} ?

To answer this question, notice that the set $\Delta(S \times Z)$ can be partitioned according to the different marginal probabilities induced on S by its elements. Given that we are considering a Bayesian agent, the following distinctions are relevant. To start with, there is a cell of the partition, denoted by Δ^* , which contains all the hypothetical lotteries inducing the marginal probability π^* corresponding to the agent's actual beliefs. Even though it cannot be identified at first, it is known to exist, by assumption. Any other cell of the partition, denoted generically by Δ° , contains all the hypothetical lotteries inducing some other marginal probability π° . Restricted to Δ^* , preferences in $\hat{\succsim}$ can be interpreted as *ordinary choices* no more and no less than preferences in \succsim . Indeed, Δ^* corresponds to the set X^S as it is seen by the decision-maker given what she believes the state of nature is. Restricted to another cell Δ° ,⁵¹ preferences in $\hat{\succsim}$ can be interpreted as *counterfactual choices*. With \hat{f} and \hat{g} two hypothetical lotteries inducing the same marginal probability π° , $\hat{f} \hat{\succsim} \hat{g}$ could mean that, if she had this other set of beliefs rather than the one she actually holds, the decision-maker would choose the prospect described in \hat{f} , rather than the one described in \hat{g} .

⁵⁰See Karni and Mongin, 2000, section 4.3, for another methodological discussion of hypothetical preferences. By contrast with ours, Karni and Mongin's discussion is more concerned with assessing whether ordinary preferences really are immune to the criticisms usually levied against hypothetical preferences. We focus here on the preliminary step of clarifying those criticisms more completely than elsewhere in the current literature.

⁵¹Notice that in some cases, π° may correspond to a Bayesian updating of π^* after the reception of some information by the decision-maker. The restriction of hypothetical preferences to a generic cell Δ° has been investigated in a theorem related to the Karni-Schmeidler theorem (see Karni, Schmeidler, and Vind, 1983). It has been proved that it would not be sufficient for beliefs to be identified uniquely (see Karni and Mongin, 2000, section 3.4). To that end, it is indispensable to let hypothetical preferences be defined across different cells of the partition, i.e., necessary to introduce an even stronger form of hypotheticality than the one of counterfactual choices.

Nevertheless, in general, preferences in $\hat{\succsim}$ are defined over pairs of elements of $\Delta(S \times Z)$ belonging to different cells of the partition of interest. When they are defined over lotteries that induce mutually incompatible marginal probabilities, preferences in $\hat{\succsim}$ differ most radically from ordinary choices. With two hypothetical lotteries \hat{f} and \hat{g} inducing conflicting marginal probabilities, $\hat{f} \hat{\succsim} \hat{g}$ would mean that the agent would rather choose the prospect described in \hat{f} , having the beliefs given in $\pi^{\hat{f}}$, than the one described in \hat{g} , having the beliefs given in $\pi^{\hat{g}}$.⁵² But, for a Bayesian decision-maker, no such choice exists. She acts given some beliefs on the state of nature, never *across conflicting beliefs*. Notice that the conflict is unavoidable in Karni and Schmeidler’s approach, because it underlies the basic preferences on which the representation is built. Indeed, assume for instance that Mr. Smith prefers the hypothetical outcome “\$100 and the operation has succeeded” to this other one, “\$0 and the operation has failed”. Such a preference cannot be equated with a choice, as both Mrs. Smith’s death and Mrs. Smith’s survival are presented as certain at the same time. Neither such basic comparisons, nor more sophisticated ones can be interpreted as choices, be it counterfactually. They amount to *impossible choices*.⁵³

Admittedly, the hypothetical preference approach should be considered as a whole. In this approach, unlike in any approach which consists simply in asking the agent about the beliefs she holds, beliefs are inferred from preference data, *part of which* is indisputably interpretable as choice data. Specifically, this is true of \succsim and the restriction of $\hat{\succsim}$ to Δ^* . To this extent,

⁵²Thus, hypothetical preferences are reminiscent of *extended preferences* which have been considered in social choice theory (see in particular Mongin, 2001, sections 4 and 6). Such preferences typically read as follows: “from the point of view of individual i , it is better to be individual j in social state x , than to be individual k in social state y ”. Like hypothetical preferences, in general, extended preferences cannot be related to ordinary or counterfactual choices.

⁵³Notice the following contrast with Savage’s own proposal regarding state-dependent utility issues. Savage’s approach introduces impossible options and, to that extent, impossible choices. Karni and Schmeidler’s approach introduces impossible choices, but no impossible option (recall the difference between $\Delta(S \times Z)$ and $(S \times Z)^S$, which we highlighted above).

the beliefs identified are the ones upon which the agent is ready to act in her own decision-making. Nevertheless, the beliefs identified are only *partially* grounded in choices. For the identification issues raised by state-dependent utility to be solved following the hypothetical preference approach, as we detailed above, the crucial requirement is that preference data have to be enlarged beyond choice data. We now claim that this is not an artifact of this approach, but a necessary feature of any complete solution to the problem of state-dependent utility, which deserves philosophical attention.

3.4 Hypothetical preferences in philosophical perspective

Assume you are facing a decision-maker who has no influence on the events regarding which her beliefs are of interest to you. Then, we argue, enlarging preference data beyond choice data is *necessary* to identify her beliefs, that is, to solve the problem of state-dependent utility.

First, recall the last two variations on the case of Mr. Smith. They illustrate the problem of state-dependent utility *without* state-dependent preferences. To solve this problem, choice data do not suffice. This point is in fact consensual. Even critics of this problem agree, as they argue for the identification of beliefs given in (1) by stressing that decompositions alternative to (1) cannot be distinguished from (1) based on observable choices. This is also indirectly confirmed by Savage's own approach. In effect, Savage proposes a solution based on the assumption of state-independent utility. This assumption can only be implicit in his framework, in the sense that it cannot be fully articulated using his own primitive concepts. In Karni and Schmeidler's framework, by contrast, it can be made explicit. For Bayesian decision-makers at least, it corresponds to the statement that, for any consequence x , and for all states s and t , we have $(x, s) \sim (x, t)$. Yet, as we previously observed, such preferences are hypothetical ones which cannot be interpreted as choices (neither as ordinary ones, nor as counterfactual

ones). Thus, instead of being an assumption of the *decision theorist*, state-independent utility can truly represent the preferences of the *decision-maker*, but this is only if preferences are allowed to extend beyond choices.

Second, recall the first two variations on the case of Mr. Smith. They illustrate the problem of state-dependent utility *with* state-dependent preferences. Considered in isolation from the previous one and by contrast with it, this problem may seem more familiar and tractable. Yet we emphasize that, as we proved in the previous section, it has two logically independent sides. The issues illustrated by each side do not cancel out, but multiply. Specifically, a robust solution to this problem should be able to define a subjective probability for a decision-maker whose conditional utility functions would be, say, $u(x) = x$ on E_1 , $u(x) = -x$ on E_2 , $u(x) = \sqrt{x}$ on E_3 and $u(x) = -\sqrt{x}$ on E_4 , with $\{E_i\}_{i=1}^4$ a partition of the state space. The fact that such a preference structure would be implausible in most economic contexts is irrelevant for the conceptual point at stake here. We are unaware of *any* generalization of Savage's theorem or related results that could construct a subjective probability in this case based on the kind of preference information available in Savage's framework, i.e., without letting preferences extend beyond ordinary choices. Some approaches aim at solving the problem of state-dependent utility while keeping Savage's informational basis. They try to do so by relaxing his state-independence postulates. But their weaker postulates too would be violated in this case, and no subjective probability

could be constructed based on these postulates.⁵⁴ We take cases like the one above to illustrate the fact that with state-dependent utility, there might never be enough structure common to the decision-maker’s conditional preferences for any belief to emerge based on cross-state comparisons. Yet such comparisons, put forward e.g. in Savage’s fourth postulate and related conditions, are unavoidable if the beliefs underlying choices are to be inferred from choice data alone. They can be dispensed with in the construction of subjective probability only if some access is gained to state-(in)dependent utility as such, like in Karni and Schmeidler’s approach. This demands to neutralize beliefs, hence to move away from ordinary choices. To sum up, even putting aside the problem of state-dependent utility *without* state-dependent preferences, it seems unlikely that the full problem of state-dependent utility *with* state-dependent preferences can be solved without letting preferences extend beyond choices. This is a new line of argument for the hypothetical preference approach, which deserves to be highlighted.

The restriction of preference data to choice data has a name in the economic literature: it defines the *revealed preference framework*. Strictly speaking, revealed preference results⁵⁵ examine the conditions under which a choice function can be represented as induced by an underlying preference

⁵⁴Consider for example the result in Hill, 2009, which is set in Savage’s framework. It is the only such result we are aware of which proposes an identification of subjective probability applicable when *neither* Savage’s third postulate, *nor* his fourth is satisfied. It relies on generalizations of these two postulates, which would be violated in the case above. Admittedly, in this case, the state space might be partitioned in such a way that, conditional on any event within each cell of the partition (by contrast with events across such cells), Savage’s third postulate holds, and likewise for his fourth. This last case is investigated in Karni and Schmeidler, 1993. However, the partition relevant for Savage’s third postulate might differ from the one relevant to his fourth, e.g., the third postulate could be respected on either $E_1 \cup E_3$ or $E_2 \cup E_4$, while the fourth postulate would be respected on either $E_1 \cup E_2$ or $E_3 \cup E_4$. Karni and Schmeidler’s 1993 result does not cover this case nor indicates how to do so. We conjecture that it is intractable in Savage’s setup.

⁵⁵See especially Samuelson, 1950 for the pioneering original competitive consumer version of those results, and e.g. Sen, 1971 for the later more abstract set-theoretical version. Sen, 1973 illustrates the now familiar philosophical discussions of the revealed preference *semantics* according to which, in the results above, preference is *defined* in terms of choices (see Hausman, 2000 for a more recent example of such discussions). Our discussion of the revealed preference *framework* is distinct. In particular, the methodological issues considered below are relevant even if “choice” and “preference” are treated as distinct concepts.

relation, i.e., they characterize the choice patterns associated with various kinds of preferences. By extension, in a revealed preference framework, any statement about preference translates in one about choice, and preferences are observable in the sense choices are. This is not the case, for instance, when preferences are the object of introspective reports only. Notoriously, the revealed preference framework prevails in economic theory, which is one aspect of its aiming to be an empirical theory. More importantly for our purposes, it is usually taken for granted that this framework fits the needs of economics, that is, *it is understood that the bulk of economic theory can be developed wholly within the revealed preference framework*. Admittedly, this framework might be restrictive with respect to our informal understanding of preference, following which the domain of preference exceeds that of choice. Nonetheless, as it is usually held, dissatisfaction with the revealed preference framework grows only from outside economic theory.

This is where we wish to locate the problem of state-dependent utility, which is a rare counter-example to the analysis above. Indeed, the take-home message of this problem is that *beliefs cannot be identified within a revealed preference framework*. Thus one thing has to go, i.e., one should either weaken the goal of identifying beliefs, or relax the revealed preference constraint under which this goal is aimed at. The qualifications given at the very beginning of this section notwithstanding, economic theory cannot abandon the goal of identifying beliefs. In particular, too many important results rely on assumptions about the beliefs of a decision-maker, or explore the assumption that several decision-makers share the same prior beliefs, for

instance.⁵⁶ Therefore, to identify beliefs, the revealed preference constraint has to be relaxed to some extent. This is what is done in Karni and Schmeidler’s approach, by complementing revealed preference data with one carefully selected kind of non-revealed preference data. Karni and Schmeidler’s hypothetical preferences are better described as genuine preferences over hypothetical lotteries. Yet, in general, they amount to impossible choices. Introducing such non-revealed preferences is a feature which is common to all the approaches aiming at providing a comprehensive solution to the problem of state-dependent utility.⁵⁷ There might be other topics that illustrate the fact that the revealed preference framework does not fully fit the needs of economic theory.⁵⁸ But the problem of state-dependent utility is the most convincing case we are aware of. In the technical literature on which our paper is based, this problem is not presented in this wider perspective about the internal limits of the revealed preference framework.

In light of the above, we argue that philosophers would benefit from paying closer attention to the problem of state-dependent utility. To our

⁵⁶For example, some results in portfolio theory rely on assumptions about the investor’s beliefs regarding the return of assets (see Karni and Schmeidler, 1993, p. 272 for a discussion of a classical theorem by Arrow in light of the problem of state-dependent utility with state-dependent preferences). On the other hand, many results rely on a shared prior assumption, especially in the interactive epistemology literature which we mentioned in the introduction. Of special interest in our context is the use of such assumption to circumscribe the correct use of Pareto conditions in the literature about the aggregation of preferences under uncertainty (see Gilboa, Samet, and Schmeidler, 2004, and a discussion in Karni, 2007, section 3.1 in light of the problem of state-dependent utility without state-dependent preferences). If beliefs are unidentified, none of these results is applicable.

⁵⁷For instance, another important stream in the literature is the *conditional expected utility* approach (see Luce and Krantz, 1971, Fishburn, 1973, and more recently Karni, 2007). In this approach, preferences are defined over acts conditioned on *different* events, as when Mr. Smith prefers a certain act, knowing that the operation succeeds, to another act, knowing that it fails. This amounts to acting across conflicting beliefs, and large parts of the discussion of section 3.3 would apply there as well. This is why we claimed that focusing on Karni and Schmeidler’s result entailed no loss of generality for our topic.

⁵⁸For example, it is sometimes argued that, in order to carry out welfare evaluations, but also to have a more unified theory of decision-making, economics needs a cardinal utility applicable to decisions under certainty (see e.g. the discussion in Wakker, 1994, section 2). In general, however, preferences over certain options cannot deliver a cardinal utility, unless one introduces a notion of *preference differences* in this context. Nevertheless, such preference differences are generally held to be incompatible with the revealed preference framework (see Fishburn, 1970, section 6.1 for a representative statement of this view).

knowledge, they rarely consider it in great detail. This may be related to the fact that they are often sympathetic to the re-description strategy sketched at the beginning of this section⁵⁹ (its other defects being set apart, this strategy tends to hide the full extent of the problem). Nevertheless, they often think about the appropriate method for investigating beliefs. In particular, they have written extensively about the shortcomings of the “pragmatic” approach to this issue. Admittedly, their typical targets are pragmatic arguments such as the one trying to establish that the decision-maker’s beliefs *should* take the form of a unique probability measure over a given state space. This is the so-called Dutch Book argument, which many philosophers claim to be a wrong kind of argument for a right normative conclusion. But we take it that their interest for the appropriate method for investigating beliefs extends beyond this case to any approach linking such investigation with that of desires and choices.⁶⁰ Accordingly, the problem of state-dependent utility should be integrated into their reflection. This might be done in several ways. Some might try to turn this problem into a general internal objection to the pragmatic approaches. Specifically, they might try to argue that this problem illustrates the self-defeatingness of the project aiming at inferring beliefs from choices, instead of following another method more adapted to epistemic issues. Other philosophers might scrutinize more carefully the conclusion presented in this section, which involves more subtly not two, but three concepts, namely, choice, belief, and preference. They could argue that the problem illustrates the need of separating conative and cognitive matters altogether, and thus of parting with the standard betting approach more radically than it has been considered here. We leave it to them to judge which of these and other ways best fits their own concerns. It is enough for us

⁵⁹See e.g. the discussion about outcome individuation in Joyce, 1999, section 2.2.

⁶⁰See e.g. Joyce, 1998 for a critique of the Dutch Book argument along the line mentioned. Regarding the wider understanding of the methodological topic under discussion, see e.g. Joyce, 1999, p. 89, where “pragmatism” is defined as the claim according to which “we can learn everything we need to know about epistemology by doing decision theory”.

to have extracted from the problem of state-dependent utility a conceptual message which is of clear interest to them.

4 Conclusion

The problem of state-dependent utility is usually underestimated. First, the problem has many different aspects. As we illustrated with Savage's theorem, four types of state-dependent utility issues stand in the way of the betting approach to the identification of beliefs. Any specific model is challenged to solve each of those four issues. Second, as we illustrated with Karni and Schmeidler's theorem, the solution to this problem has a high methodological cost. In order to identify beliefs under state-dependent utility, with or without state-dependent preferences, it is necessary to let preferences extend beyond choices. This is in itself a significant conceptual conclusion. Given the importance of belief identification, this indicates that the revealed preference framework does not fully fit the needs of economic theory. This also suggests that philosophers, who think about the appropriate method for investigating beliefs, should pay closer attention to the problem of state-dependent utility.

5 Bibliography

- ALLAIS, M. (1953): “Le comportement de l’homme rationnel devant le risque : critique des postulats et axiomes de l’école américaine”, *Econometrica*, 21(4), 503–546.
- ANSCOMBE, F., AND R. AUMANN (1963): “A Definition of Subjective Probability”, *The Annals of Mathematical Statistics*, 34(1), 199–205.
- ARROW, K. (1953): “Le rôle des valeurs boursières pour la répartition la meilleure des risques”, *Colloques Internationaux du Centre National de la Recherche Scientifique, Économétrie*, 40, 41–58.
- (1965): *Aspects of the Theory of Risk-Bearing*. Helsinki : Yrjö Jahnssoonin Säätiö.
- (1973): “Optimal Insurance and Generalized Deductibles”, *Office of Economic Opportunity Report*, R-1108-OEO.
- AUMANN, R. (1976): “Agreeing to Disagree”, *The Annals of Statistics*, 4(6), 1236–1239.
- BONANNO, G., AND K. NEHRING (1997): “Agreeing to Disagree: A Survey”, *U.C.-Davis Department of Economics Working Paper No. 97-18*.
- BRADLEY, R. (2004): “Ramsey’s Representation Theorem”, *Dialectica*, 58(4), 483–497.
- CLEMEN, R., AND R. WINKLER (2007): “Aggregating Probability Distributions”, in *Advances in Decision Analysis*, ed. by W. Edwards, R. Miles, and D. von Winterfeldt, pp. 154–176. Cambridge : Cambridge University Press.
- DE FINETTI, B. (1937): “La prévision : ses lois logiques, ses sources subjectives”, *Annales de l’Institut Henri Poincaré*, 7(1), 1–68.

- DRÈZE, J. (1961): “Les fondements logiques de la probabilité subjective et de l’utilité”, *La décision, Colloques Internationaux du Centre National de la Recherche Scientifique*, pp. 73–87.
- (1987): *Essays on Economic Decisions under Uncertainty*. Cambridge : Cambridge University Press.
- ELLSBERG, D. (1961): “Risk, Ambiguity, and the Savage Axioms”, *The Quarterly Journal of Economics*, 75(4), 643–669.
- FINKELSTEIN, A., E. LUTTMER, AND M. NOTOWIDIGDO (2013): “What Good is Wealth Without Health? The Effect of Health on the Marginal Utility of Consumption”, *Journal of the European Economic Association*, 11(1), 221–258.
- FISHBURN, P. (1970): *Utility Theory for Decision Making*. New York : Wiley.
- (1973): “A Mixture-Set Axiomatization of Conditional Subjective Expected Utility”, *Econometrica*, 41(1), 1–25.
- GENEST, C., AND J. ZIDEK (1986): “Combining Probability Distributions: A Critique and an Annotated Bibliography”, *Statistical Science*, 1(1), 114–135.
- GILBOA, I. (1987): “Expected Utility with Purely Subjective Non-Additive Probabilities”, *Journal of Mathematical Economics*, 16(1), 65–88.
- (2009): *Theory of Decision under Uncertainty*. New York : Cambridge University Press.
- GILBOA, I., D. SAMET, AND D. SCHMEIDLER (2004): “Utilitarian Aggregation of Beliefs and Tastes”, *Journal of Political Economy*, 112(4), 932–938.
- GOLLIER, C. (1997): “About the Insurability of Catastrophic Risks”, *Geneva Papers on Risk and Insurance*, 22(83), 177–186.

- GOOD, I. (1950): *Probability and the Weighing of Evidence*. London : C. Griffin and Company Limited.
- GRANT, S., AND E. KARNI (2004): “A Theory of Quantifiable Beliefs”, *Journal of Mathematical Economics*, 40(5), 515–546.
- HÁJEK, A. (2008): “Dutch Book Arguments”, in *The Handbook of Rational and Social Choice*, ed. by P. Anand, P. Pattanaik, and C. Puppe, pp. 173–196. Oxford : Oxford University Press Oxford.
- HAUSMAN, D. (2000): “Revealed Preference, Belief, and Game Theory”, *Economics and Philosophy*, 16(1), 99–115.
- HILL, B. (2009): “When Is There State Independence?”, *Journal of Economic Theory*, 144(3), 1119–1134.
- (2010): “An Additively Separable Representation in the Savage Framework”, *Journal of Economic Theory*, 145(5), 2044–2054.
- HORN, A., AND A. TARSKI (1948): “Measures in Boolean Algebras”, *Transactions of the American Mathematical Society*, 64(3), 467–497.
- JOYCE, J. (1998): “A Nonpragmatic Vindication of Probabilism”, *Philosophy of Science*, 65(4), 575–603.
- (1999): *The Foundations of Causal Decision Theory*. New York : Cambridge University Press.
- KARNI, E. (1985): *Decision Making under Uncertainty: The Case of State-Dependent Preferences*. Cambridge : Harvard University Press.
- (1996): “Probabilities and Beliefs”, *Journal of Risk and Uncertainty*, 13(3), 249–262.
- (1999): “Elicitation of Subjective Probabilities when Preferences are State-Dependent”, *International Economic Review*, 40(2), 479–486.

- (2003): “On the Representation of Beliefs by Probabilities”, *Journal of Risk and Uncertainty*, 26(1), 17–38.
- (2007): “Foundations of Bayesian Theory”, *Journal of Economic Theory*, 132(1), 167–188.
- (2008): “On Optimal Insurance in the Presence of Moral Hazard”, *The Geneva Risk and Insurance Review*, 33(1), 1–18.
- (2009): “A Mechanism for Eliciting Probabilities”, *Econometrica*, 77(2), 603–606.
- (2011a): “Subjective Probabilities on a State Space”, *American Economic Journal: Microeconomics*, 3(4), 172–185.
- (2011b): “A Theory of Bayesian Decision Making with Action-Dependent Subjective Probabilities”, *Economic Theory*, 48(1), 125–146.
- KARNI, E., AND P. MONGIN (2000): “On the Determination of Subjective Probability by Choices”, *Management Science*, 46(2), 233–248.
- KARNI, E., AND D. SCHMEIDLER (1993): “On the Uniqueness of Subjective Probabilities”, *Economic Theory*, 3(2), 267–277.
- KARNI, E., D. SCHMEIDLER, AND K. VIND (1983): “On State Dependent Preferences and Subjective Probabilities”, *Econometrica*, 51(4), 1021–1031.
- KARNI, E., AND M.-L. VIERØ (2013): ““Reverse Bayesianism”: A Choice-Based Theory of Growing Awareness”, *The American Economic Review*, 103(7), 2790–2810.
- KOOPMAN, B. (1940): “The Axioms and Algebra of Intuitive Probability”, *The Annals of Mathematics*, 41(2), 269–292.

- KRAFT, C., J. PRATT, AND A. SEIDENBERG (1959): “Intuitive Probability on Finite Sets”, *The Annals of Mathematical Statistics*, 30(2), 408–419.
- LUCE, R., AND D. KRANTZ (1971): “Conditional Expected Utility”, *Econometrica*, 39(2), 253–271.
- MACHINA, M., AND D. SCHMEIDLER (1992): “A More Robust Definition of Subjective Probability”, *Econometrica*, 60(4), 745–780.
- MONGIN, P. (2001): “The Impartial Observer Theorem of Social Ethics”, *Economics and Philosophy*, 17(2), 147–179.
- NAU, R. (1995): “Coherent Decision Analysis with Inseparable Probabilities and Utilities”, *Journal of Risk and Uncertainty*, 10(1), 71–91.
- (2001): “de Finetti Was Right: Probability Does Not Exist”, *Theory and Decision*, 51(2-4), 89–124.
- RAMSEY, F. (1931): “Truth and Probability”, *The Foundations of Mathematics and Other Logical Essays*, pp. 156–198.
- SAMUELSON, P. (1950): “The Problem of Integrability in Utility Theory”, *Economica*, 17(68), 355–385.
- SAVAGE, L. (1954/1972): *The Foundations of Statistics*. New York : Dover.
- SAVAGE, L., AND R. AUMANN (1987): “Letters Between Leonard Savage and Robert Aumann (1971)”, in *Essays on Economic Decisions under Uncertainty*, ed. by J. Drèze, pp. 76–81. Cambridge : Cambridge University Press.
- SCHERVISH, M., T. SEIDENFELD, AND J. KADANE (1990): “State-Dependent Utilities”, *Journal of the American Statistical Association*, 85(411), 840–847.

- SCHLAG, K., J. TREMEWAN, AND J. VAN DER WEELE (2015): “A Penny for Your Thoughts: A Survey of Methods For Eliciting Beliefs”, *Experimental Economics*, 18(3), 457–490.
- SEN, A. (1971): “Choice Functions and Revealed Preference”, *The Review of Economic Studies*, 38(3), 307–317.
- (1973): “Behaviour and the Concept of Preference”, *Economica*, 40(159), 241–259.
- WAKKER, P. (1987): “Subjective Probabilities for State Dependent Continuous Utility”, *Mathematical Social Sciences*, 14(3), 289–298.
- (1993): “Clarification of Some Mathematical Misunderstandings about Savage’s Foundations of Statistics, 1954”, *Mathematical Social Sciences*, 25(2), 199–202.
- (1994): “Separating Marginal Utility and Probabilistic Risk Aversion”, *Theory and Decision*, 36(1), 1–44.
- WAKKER, P., AND H. ZANK (1999): “State Dependent Expected Utility for Savage’s State Space”, *Mathematics of Operations Research*, 24(1), 8–34.

6 Appendix

In this appendix, it is proved that preferences representable as in (2) with $u_{\overline{E^*}}(x) = -x$ and $E^* = (0, \frac{2}{3}]$ respect Savage's fourth postulate, namely:

$$\forall E, E' \in 2^S, \forall x, y, x', y' \in X \text{ s.t. } x \succ y, x' \succ y' : xEy \succ xE'y \Leftrightarrow x'E'y' \succ x'E'y'.$$

First, it is readily checked that for those preferences, $x \succ y \Leftrightarrow x > y$. Next, suppose that $x \succ y$ and $xEy \succ xE'y$. Using (2), this is true if and only if:

$$\begin{aligned} & \pi^*(E).u_E(x) + \pi^*(\overline{E}).u_{\overline{E}}(y) > \pi^*(E').u_{E'}(x) + \pi^*(\overline{E}').u_{\overline{E}'}(y) \\ \Leftrightarrow & \pi^*(E \cap E^*)x - \pi^*(E \cap \overline{E^*})x + \pi^*(\overline{E} \cap E^*)y - \pi^*(\overline{E} \cap \overline{E^*})y > \\ & \pi^*(E' \cap E^*)x - \pi^*(E' \cap \overline{E^*})x + \pi^*(\overline{E}' \cap E^*)y - \pi^*(\overline{E}' \cap \overline{E^*})y \\ \Leftrightarrow & x [\pi^*(E \cap E^*) - \pi^*(E' \cap E^*) + \pi^*(E' \cap \overline{E^*}) - \pi^*(E \cap \overline{E^*})] > \\ & y [\pi^*(\overline{E}' \cap E^*) - \pi^*(\overline{E} \cap E^*) + \pi^*(\overline{E} \cap \overline{E^*}) - \pi^*(\overline{E}' \cap \overline{E^*})]. \end{aligned}$$

Noticing that $\pi^*(E' \cap \overline{E^*}) + \pi^*(\overline{E}' \cap \overline{E^*}) = \pi^*(\overline{E^*}) = \pi^*(E \cap \overline{E^*}) + \pi^*(\overline{E} \cap \overline{E^*})$ and that $\pi^*(E' \cap E^*) + \pi^*(\overline{E}' \cap E^*) = \pi^*(E^*) = \pi^*(E \cap E^*) + \pi^*(\overline{E} \cap E^*)$, using the fact that $x \succ y \Leftrightarrow x > y$ and $x' \succ y' \Leftrightarrow x' > y'$, the supposition is true if and only if:

$$\begin{aligned} & x [\pi^*(E \cap E^*) - \pi^*(E' \cap E^*) + \pi^*(E' \cap \overline{E^*}) - \pi^*(E \cap \overline{E^*})] > \\ & y [\pi^*(E \cap E^*) - \pi^*(E' \cap E^*) + \pi^*(E' \cap \overline{E^*}) - \pi^*(E \cap \overline{E^*})] \\ \Leftrightarrow & [\pi^*(E \cap E^*) - \pi^*(E' \cap E^*) + \pi^*(E' \cap \overline{E^*}) - \pi^*(E \cap \overline{E^*})] > 0 \\ \Leftrightarrow & x' [\pi^*(E \cap E^*) - \pi^*(E' \cap E^*) + \pi^*(E' \cap \overline{E^*}) - \pi^*(E \cap \overline{E^*})] > \\ & y' [\pi^*(E \cap E^*) - \pi^*(E' \cap E^*) + \pi^*(E' \cap \overline{E^*}) - \pi^*(E \cap \overline{E^*})] \\ \Leftrightarrow & x' [\pi^*(E \cap E^*) - \pi^*(E' \cap E^*) + \pi^*(E' \cap \overline{E^*}) - \pi^*(E \cap \overline{E^*})] > \\ & y' [\pi^*(\overline{E}' \cap E^*) - \pi^*(\overline{E} \cap E^*) + \pi^*(\overline{E} \cap \overline{E^*}) - \pi^*(\overline{E}' \cap \overline{E^*})] \\ & \Leftrightarrow x'E'y' \succ x'E'y'. \end{aligned}$$