The Problem of State-Dependent Utility: A Reappraisal

Jean Baccelli[*]

Munich Center for Mathematical Philosophy

June 2, 2019

Abstract

State-dependent utility is a problem for the behavioral branch of decision theory under uncertainty. It questions the very possibility that beliefs be revealed by choice data. According to the current literature, all models of beliefs are equally exposed to the problem. Moreover, the problem is solvable only when the decision-maker can influence the resolution of uncertainty. This paper gives grounds to reject these two views. The various models of beliefs can be shown to be unequally exposed to the problem of state-dependent utility. The problem can be argued to be solved even when the decision-maker has no influence over the resolution of uncertainty. The implications of such reappraisal for a philosophical appreciation of the revealed preference methodology are discussed.

1 Introduction

A large stream of contemporary decision theory, originating in the works of Ramsey, de Finetti, and especially Savage, adopts a behavioral approach and is based on a methodological promise. The promise is that the beliefs of a decision-maker regarding the realization of uncertain events can be satisfactorily identified by observing her choice behavior. In the words of de Finetti, "the degree of probability attributed by an individual to a given event is revealed by the conditions under which he would be disposed to bet on that event" (de Finetti, 1937, p. 6, as in Kyburg and Smokler, 1980, p. 61; similar passages can be found in Ramsey, 1931, p. 172 or Savage, 1972, p. 27-28). By providing powerful assumptions under which a decision-maker's choices can be represented by the expectation of a utility function with respect to a probability measure, Savage's axiomatization of subjective expected utility—"the crowning glory of choice theory" (Kreps, 1988, p. 120)—seems to deliver on that promise. Specifically, the uniqueness properties of Savage's representation seem to vindicate the view according to which his result provides the entire edifice of Bayesian epistemology with secure behavioral foundations.

However, it is now well understood in the literature that the above conclusion sweeps at least one unpleasant fact under the rug. There is a fundamental difficulty with the uniqueness properties of the subjective expected utility representation. For brevity, it can be called the problem of state-dependent utility (Seidenfeld et al., 1990; Schervish et al., 1990; 2013; Karni, 1996). The problem is that Savage's axioms prove to exclude only some, not all forms of state-dependent utility. This leaves the subjective probability measure essentially unidentified in the representation. Therefore, the representation does not deliver the promised behavioral foundations. Methodologically speaking, the problem seems to run deep. Indeed, the current state of the state-dependent utility literature in the behavioral tradition of decision theory (as reflected in overviews such as Karni, 2008; Baccelli, 2017) is best described by the conjunction of the following two key views. First, expected utility and non-expected utility are equally exposed to the problem of state-dependent utility. This means in particular that the problem arises whatever the form of the beliefs to be revealed by choice data—be it a classical probability measure, a set thereof, or a non-additive probability measure, to mention but a few options. Call this "the general problem view". Second, the beliefs of a decision-maker cannot be behaviorally identified, unless the decision-maker has some capacity to influence the resolution of uncertainty. This is what decision theorists call a case of "moral hazard". Traditionally, they are keen on assuming moral hazard away, rather than bringing it in, as here suggested. Call this "the radical solution view". The overall methodological situation, then, looks almost aporetic. According to some, the philosophical lesson is that we should simply drop the project of providing the

various models of beliefs with behavioral foundations. There would be no viable alternative to accepting these models as primitives (Nau, 2001).

The main contribution of this paper will be to give grounds to reject the two views above. As I will show under some assumptions to be discussed in due course, first, non-expected utility proves structurally less exposed to the problem of state-dependent utility than expected utility (the former being naturally understood, in this statement, as excluding the latter). The problem is more model-sensitive, and expected-utility-specific, than currently assumed. Second, the beliefs of a decision-maker can be behaviorally identified even though the decision-maker has no influence over the resolution of uncertainty. Moral hazard may be sufficient, but the very assumptions used to show this also entail that moral hazard is not necessary to solve the problem of state-dependent utility. From all this, one should not infer too quickly that I am bringing exonerating news for the behavioral methodologies, however. As I will go on to explain, the take-home message of this paper is certainly not that the problem of state-dependent utility is methodologically less embarrassing than currently understood in the literature. It is only that it is more complex. The conclusions thus previewed contribute primarily to decision theory. They also contribute to the philosophy of economics, as economics is based on decision theory, and the philosophy of psychology, because decision theory investigates the measurability of beliefs.

The paper is organized as follows. Section 2 recaps the problem of state-dependent utility. Section 3 spells out the general problem view and the radical solution view in more detail. Section 4 sheds more light on the latter view, admittedly the most enigmatic of the two. Section 5 explains why both the general problem view and the radical solution view must be abandoned. Section 6 reflects on the methodological implications. A brief conclusion ensues.

2   The problem of state-dependent utility: an exposition

The analysis will be conducted in a finite-state variant of the Savage framework (Savage, 1954; 1972). The decision-maker is uncertain about the true state of nature, s^*, which she knows to be in $S = \{s_1, \ldots, s_n\}$, the set of all possible states of nature. Any subset of S is called an event. X is the set of all possible payoffs.[1] $F = X^S$ is the set of all possible choice options, called acts. Let $\succcurlyeq$ denote a binary preference relation over F, with symmetric and asymmetric parts $\sim$ and $\succ$, respectively. The relation $\succcurlyeq$ is interpreted as a

---

[1]The richness conditions required for the construction of a representation will be imposed on X. By building on the work of Köbberling and Wakker, referred to next, I assume that X satisfies a "solvability" condition (Köbberling and Wakker, 2003, p. 398, Def. 3).

set of observed choice data.[2]

Axiomatizations of the classical subjective expected utility model are available in this special setting (e.g., Köbberling and Wakker, 2003, especially Sec. 4.1, and the references therein). These results provide axiomatic conditions on $\succeq$ under which the following claims hold. First, there exists a non-constant state-independent utility function $u : X \to \mathbb{R}$ and a probability measure $p : 2^S \to \mathbb{R}$ such that $\succeq$ can be represented by the expectation of u with respect to p, i.e., the function $v : F \to \mathbb{R}$ in (1):

$$v(f) = \sum_{i=1}^{n} p(s_i) \cdot u\big(f(s_i)\big). \tag{1}$$

Second, the existence claim above is complemented by a uniqueness claim, traditionally formulated as follows. In (1), u is unique up to a strictly positive affine transformation and p is absolutely unique. A particular implication seems to be, then, that the subjective probability measure is fully identified by the decision-maker's observed behavior.

The methodological significance of this very last claim cannot be overestimated. On it depends the received view according to which the charaterizations of subjective expected utility to be found in axiomatic decision theory provide behavioral foundations for subjective probability measures wherever they appear in science—e.g., in statistics or in interactive epistemology, where Bayesian priors are often accepted as primitives. In the terminology of economics, the axiomatizations of subjective expected utility are considered successful applications of the "revealed preference" methodology (for an overview, see Hands, 2013; also Chambers and Echenique, 2016). Conceived as precisely as our topic requires, this is the methodology demanding that any unobservable psychological concept, like the concept of belief, be put in correspondence with—which need not be interpreted as a reduction to—observable choice patterns.

However, the received view faces a potentially devastating problem, viz., the problem of state-dependent utility (Seidenfeld et al., 1990; Schervish et al., 1990, 2013; Karni, 1996). The simplest way to raise the problem is as follows. Take any $\succeq$ representable as in (1). For simplicity, further assume that $\succeq$ is such that no state is null, i.e., p in (1) has full support over S.[3] Take any other full support q. Notice that, with the definition $u_{s_i}(x) = \frac{p(s_i)}{q(s_i)} u(x)$ for all $i = 1, \ldots, n$, the following equality holds:

---

[2]This can be considered a faithful variant of the Savage framework inasmuch as the setting is purely subjective, i.e., no probabilistic notion is listed among the primitives. Contrast with, e.g., the Anscombe-Aumann framework (Anscombe and Aumann, 1963).

[3]The problem extends to when there are null events (e.g., Baccelli, 2017, Sec. 2.4). But the solutions examined in this paper do not lead to significant progress on this front. To evaluate these solutions at their best, I assume that no state is null.

$$\sum_{i=1}^{n} p(s_i) \cdot u\big(f(s_i)\big) = \sum_{i=1}^{n} q(s_i) \cdot \frac{p(s_i)}{q(s_i)} u\big(f(s_i)\big) = \sum_{i=1}^{n} q(s_i) \cdot u_{s_i}\big(f(s_i)\big). \quad (2)$$

Denote by u_s the collection of u_{s_i} thus constructed. Since attributing the same expected utility values to all options suffices to represent the same order, the equality in (2) establishes that u and p represent $\succeq$ if and only if u_s and q represent $\succeq$. It is worth highlighting that, by construction, the representation of $\succeq$ by the expectation of u_s with respect to q thus satisfies all the axioms that are necessary to its representation by the expectation of u with respect to p. This applies, in particular, to the standard state-independent utility axioms—i.e., Savage's axioms P3 and P4, (the core of) which can be found in the axiomatizations of many non-expected utility models as well. They are satisfied in both cases. The observed choice data cannot adjudicate between the two alternative representations, or indeed infinitely many others.

Consider the following concrete illustration of the implications on the identification of subjective probability. With $E \in 2^S$, let xEy stand for the act, the payoff of which is x if the true state of nature is in E, y otherwise. Imagine that for a given decision-maker, there is a distinguished event E^* such that for all $x, y \in X$, $xE^*y \sim yE^*x$. Further assume that the preferences of this decision-maker satisfy all the the traditional axioms of subjective expected utility. Then, in the representation of her preferences, the remarkable indifference above will be represented as follows:

$$\frac{1}{2}u(x) + \frac{1}{2}u(y) = \frac{1}{2}u(y) + \frac{1}{2}u(x). \quad (3)$$

But notice that, following the construction in (2), (3) holds if and only if—among infinitely many similar equalities—the following holds:

$$\frac{1}{3}u_{E^*}(x) + \frac{2}{3}u_{\overline{E^*}}(y) = \frac{1}{3}u_{E^*}(y) + \frac{2}{3}u_{\overline{E^*}}(x), \quad (4)$$

with $u_{E^*}(\cdot) = \frac{3}{2}u(\cdot)$ and $u_{\overline{E^*}}(\cdot) = \frac{3}{4}u(\cdot)$ or, equivalently, $u_{E^*}(\cdot) = u(\cdot)$ and $u_{\overline{E^*}}(\cdot) = \frac{1}{2}u(\cdot)$.[4] Thus, the example establishes that contrary to what is traditionally held, under the constraint of the traditional axioms of subjective expected utility, an event like E^* need not receive subjective probability

---

[4]Explained as in Schervish et al., 1990, the example could go as follows. Take $X = X_1 \times X_2$, with for $i = 1, 2$, X_i a monetary interval expressed in currency c_i. Assume that for each f and any of $i = 1, 2$, there is an indifferent $g := f_i$ that pays in currency c_i alone (i.e., the consequences of f_i are in $X_i \times 0$). Imagine now—this is the key—that E^* and $\overline{E^*}$ correspond to different exchange rates between c_1 and c_2. Specifically here, imagine that $c_1 = c_2$ on E^*, but $c_1 = \frac{1}{2}c_2$ on $\overline{E^*}$. The equality in (3) would then be related to a payment entirely in c_1, that in (4), to a payment entirely in c_2. Thus, the probability values associated to E^* and $\overline{E^*}$ would vary with the conventional choice of a numéraire.

value $\frac{1}{2}$. It can receive any strictly positive probability value, be it smaller or greater than $\frac{1}{2}$ (witness the alternative probability values for E^* and $\overline{E^*}$, respectively). Symmetrically and more generally, whatever the prior p featured in the initial representation (1), one can, following (2), find a behaviorally equivalent expected utility representation featuring a prior q that is uniform over S. Identification fails.

The gist of the issue is well understood (e.g., Schervish et al., 1990). It is that in (1), the utility function and the probability measure are identified not separately from one another, but as a pair. Crucially, under the existence conditions of (1), the correct pair is not behaviorally identified.[5] In particular, there are no behavioral grounds for considering the state-independent pair in (1) as the correct one. Put differently, the correctness of this representation is predicated on the assumption that all forms of state-dependent utility are excluded, and this assumption is not a behavioral one. Therefore, under the existence conditions of (1), the subjective probability measure is not fully identified by the observed behavior of the decision-maker. On the contrary, within the class of all full support probability measures, it is fully unidentified.

Incidentally, one can also present the main issue with reference to Ramsey's notion of an "ethically neutral" event. For some following the spirit—if not the letter—of Ramsey's approach (e.g., Davidson and Suppes, 1956, p. 265; Bradley, 2004, p. 488), this is a behavioral notion to be cashed out as follows. E^* is defined as ethically neutral (sometimes added: "of probability one-half") if, as in the preceding example, it is such that for all $x, y \in X$, $xE^*y \sim yE^*x$. The intended function of the assumption that such an event exists is to provide the decision theorist with an event of subjective probability value $\frac{1}{2}$, based on which to construct representation (1). In the actual writings of Ramsey (Ramsey, 1931, p. 177) and the work of its closest followers, an ethically neutral event is defined, rather, as an event E^* such that the decision-maker is indifferent between it or its complement obtaining. Evidently, the point of assuming that such an event exists is to try to ensure that, in the targeted representation, $u_{E^*}(\cdot) = u_{\overline{E^*}}(\cdot)$. Either way, unless one equates solving a problem with assuming that it is solved, a behavioral Ramseyan approach to subjective expected utility does not solve the problem of state-dependent utility. The force of the assumption that an ethically neutral event exists comes largely from the force of the non-behavioral assumption that all forms of state-dependent utility are excluded. In this specific respect,

---

[5]Thus, the main question is the following. When a state-independent utility representation exists, does there exist an alternative state-dependent utility representation? This should be carefully distinguished from the following question (investigated in, e.g., Karni, 1993; Wakker and Zank, 1999; Hill, 2009). When no state-independent utility representation exists, does there exist a state-dependent utility representation? Furthermore, one can show that, under expected utility, the identification issues raised in the present section would recur even in the latter case.

there is no significant difference between Ramsey's and Savage's approaches.

3 The current state of the state-dependent utility literature

The main implications of the problem of state-dependent utility are methodological in nature. Karni, for instance, summarizes the problem as follows: "the uniqueness of the probability in [the axiomatizations of subjective expected utility] is predicated on the convention that the utility function be state-independent (...). This convention is not implied by the axioms and, consequently, lacks choice-behavioral implications" (Karni, 2008, p. 227). Thus, although their other merits are left untouched, these axiomatizations do not provide Bayesian epistemology with the promised behavioral foundations.

However, to fully appreciate the methodological significance of the problem, one needs to answer the following two questions. First, how general is the problem? Indeed, it has been introduced with respect to one particular model of belief, namely, the Bayesian prior. Would the same obstacle arise if one varied the targeted model of belief? Does the problem also plague the revelation of, say, sets of Bayesian priors, like in maxmin expected utility or related models (Gilboa and Schmeidler, 1989; Levi, 1974), or non-additive probability measures, like in Choquet expected utility or related models (Schmeidler, 1989; Buchak, 2013)?[6] Second, what would it take to solve the problem, that is to say, achieve an acceptably unique behavioral identification of a decision-maker's beliefs? Assume for the sake of argument that, to that end, one is not ready to relax the constraints of the revealed preference methodology. The restrictiveness of this assumption should not be underestimated. It excludes, for instance, bringing in verbal reports as admissible additional data in the identification exercise (especially Karni and Schmeidler, 2016). It also excludes from consideration the introspective tradition of decision theory (Jeffrey, 1983; Joyce, 1999), where state-dependent utility has been fruitfully discussed, yet not from the specific perspective of observing choice data, which is the only perspective relevant to the present discussion. Then, if one is not ready to relax the constraints of the revealed preference methodology—which is, notoriously, the official doctrine of economics and parts of psychology—, how radical are the methodological adjustments necessary for revealing beliefs from choice data?

What best describes the current state of the state-dependent utility lit-

---

[6]I defer to Section 6 the discussion of whether it is problematic to interpret such functions as measuring the decision-maker's beliefs (by contrast with a compound of these and other attitudes). I observe that, in any case, this interpretation is unproblematic when one starts from these functions as primitives, instead of trying to recover them from choice data.

erature in the behavioral tradition of decision theory is, I contend, the following two answers to the questions above. First, the problem applies to expected and non-expected utility alike. It is a problem with revealing beliefs in general, not just with revealing Bayesian priors in particular. Call this "the general problem view". The most explicit endorsement is in Nau, 2001 (p. 101; see also Karni, 2011, p. 126): "all of the above remarks about the indeterminacy of true probabilities in Savage's model apply to every other theory of choice under uncertainty that uses '[payoffs]' as a primitive, including (...) Choquet expected utility, maxmin expected utility, probabilistic sophistication, biseparable preferences, and quantifiable beliefs". Second, the only situations in which choice data uniquely identifies the generating beliefs are so-called "moral hazard" (a curious, but established terminology which I will here accept without discussion) situations. Those are cases where, contrary to common assumption, the decision-maker has the capacity to influence the resolution of uncertainty.[7] Call this "the radical solution view". The most explicit endorsement is in Drèze and Rustichini, 2004 (p. 883; see also Karni, 2008, p. 228): "[the] separate identification of subjective probabilities and state-dependent utilities is possible only to the extent that the realisation of the corresponding states can be influenced by the decision maker".[8]

Before embarking on explicating and eventually arguing against these two established views, it is worth pausing to realize just how uncomfortable the overall methodological situation is, then, supposed to be. If the problem of state-dependent utility arises for all models of belief, and if moral hazard is the only behavioral way out, then one seems to be stuck with two equally unattractive options for identifying the beliefs of a decision-maker. The first option is to partially or totally abandon the revealed preference methodology, so as to achieve identification with non-choice data. The other option is to profoundly transform—as would be inevitable under moral hazard—one's conception of what beliefs are, so as to make beliefs amenable to that methodology. Nau seems to advocate taking the first branch of the dilemma: "true probability is a concept that makes sense only when it is regarded as a primitive (...), and revealed by unsubstantiated verbal reports" (Nau, 2001 p. 101). Drèze, on the other hand, seems to advocate taking the second branch, without even trying to hide that this has the flavour of a paradox when writing: "the subjective probability of an event can be defined precisely only when that probability is chosen by the decision-maker" (Drèze, 1961, p. 86, as in Drèze, 1987, p. 101).

---

[7]Notice that, in this decision-theoretic definition of moral hazard, the observability of the means by which the resolution of uncertainty is influenced plays no role. This is unlike in the initial game-theoretic definition of the notion (e.g., Laffont and Martimort, 2002).

[8]As it is at first so enigmatic, one might wonder whether the claim is not meant to hold only within Drèze and Rustichini's model. First, this narrow interpretation is not supported by reading the whole piece from which the quotation is taken, or the rest of the relevant literature. Second, even within this model, the claim proves not to hold (see fn. 13).

4   On state-dependent utility and moral hazard

The general problem view is quite intuitive. The axiomatizations of the non-expected utility models come with uniqueness claims that are structurally similar to the one discussed in the case of expected utility. They typically state the existence of an absolutely unique belief parameter and a state-independent utility function that is unique up to a strictly positive affine transformation.[9] Accordingly, it is natural to conjecture that the same uniqueness issues will recur under non-expected utility. By contrast, the radical solution view is not intuitive. In this section, I provide some insight into this view. I take an apparent detour. I contribute a simple explanation of why, under some assumptions, moral hazard is sufficient to solve the problem of state-dependent utility. This will shed some light on the radical solution view, i.e., the claim that moral hazard is also necessary to solve the problem. Most importantly, understanding why the former claim is true will prove instrumental in understanding (in the next section) why, under the very same assumptions, the latter is false. Still by way of anticipation, a by-product will be that the general problem view must be rejected as well.

The simplest cases of moral hazard lead to patterns of behavior that can be analyzed with the tools of expected utility theory, only, with irreducibly act-dependent probability measures. That is to say, unlike in the classical expected utility model examined in the previous section, the probabilities of the states now depend on the act considered. This reflects the fact that the decision-maker can influence the resolution of uncertainty. This influence is exerted in response to the incentives given to her in each act, i.e., each state-distribution of payoffs. Accordingly, moral hazard is often referred to, simply, as "act-state dependence" (e.g., Joyce, 1999; Bradley, 2017), and the two terminologies are typically considered equivalent. Formally, henceforth assume that there exists a state-independent utility function $u : X \to \mathbb{R}$ and, for each act, some optimal act-dependent probability measure $p^f : 2^S \to \mathbb{R}$, such that the observed behavior $\succcurlyeq$ of the decision-maker can be represented by the associated act-dependent expectation of u, i.e., the function $w : F \to \mathbb{R}$ in (5):

$$w(f) = \sum_{i=1}^{n} p^f(s_i) \cdot u\big(f(s_i)\big). \tag{5}$$

Two comments are in order. First, the classical expected utility model corresponds to the particular case of (5) where there exists a representation such

---

[9]Prospect theory would be an apparent exception (e.g., Köbberling and Wakker, 2003, Sec. 4.3). Indeed, this non-expected utility model is rather exceptional in that it features a utility function that is not just an interval scale, but a ratio scale. However, notice that (2) applies to a ratio scale no less than to an interval scale. On a different score, in the present paper, I will not consider any of the non-expected utility models that feature not one, but several utility functions (e.g., Levi, 1986; Galaabaatar and Karni, 2013).

that for all $f, g \in F$, $p_f = p_g$. Henceforth, whenever referring to (5) without further qualification, I will typically be referring to the cases where there does not exist such an act-independent representation, i.e., I will be excluding the limit case of the classical expected utility model. Second, notice the following fact about representation (5). Although the optimal probability measure depends on the act, the utility function does not. This can be interpreted as a form of moral hazard without differential effort. This means that the decision-maker can influence the resolution of uncertainty at no utility cost—more generally, possibly at a cost, but then at a cost that would be constant across all the ways she has to exert her influence. This is, in general, a major restriction. But it can be granted for the purposes of this section, which is to establish a methodologically interesting possibility result.[10]

To investigate the uniqueness properties of representation (5), I will make three further assumptions. None of these assumptions is behaviorally innocuous, but they are to be granted on the above mentioned grounds. First, I will assume that there exists a best and a worst act in F, and that the expected utility values attributed to them by w have been normalized to 1 and 0, respectively. This is merely to ensure that for preserving the representation of $\succeq$, preserving the expected utility values given by w is not only sufficient, but also necessary. Second, I will assume that the collection P of all optimal act-dependent probability measures is finite, i.e., one can write it as $P = \{p^1, \ldots, p^m\}$, for some m. In fact, it is enough to assume that P is characterized by one of its finite subsets. This is to make uniqueness arguments more tractable. Third, I will assume that all elements of the collection P have full support over S. For brevity, I will also say that P itself has full support. This suffices to keep null states out of the picture. Once again, I emphasize that these are behaviorally restrictive assumptions. But I leave it to further work to determine if the possibility result they lead to can be obtained under weaker assumptions.

The key question for our purposes is how the analogue of argument (2) would fare in the context of representation (5). To answer this question, consider a full support $Q = \{q^1, \ldots, q^m\}$ introduced as an alternative to the initial $P = \{p^1, \ldots, p^m\}$. Like in (2), with p^f and q^f the relevant elements of P and Q, the following equality holds:

$$w(f) = \sum_{i=1}^{n} p^f(s_i) \cdot u\big(f(s_i)\big) = \sum_{i=1}^{n} q^f(s_i) \cdot \frac{p^f(s_i)}{q^f(s_i)} u\big(f(s_i)\big) \qquad (6)$$

But notice that, unlike in (2), this would induce a vector of state-dependent utility weights that, in general, would vary across acts. That is to say, with

<hr />

[10]As this assumption suggests and the rest of the section will confirm, only some forms of moral hazard can be of any help in solving the problem of state-dependent utility. Moral hazard does not, absent special assumptions, suffice to solve the problem (more on this non-sufficiency in Baccelli, forthcoming, Sec. 3).

f and g any act, s_i any state, and Q any alternative full support collection, it cannot be excluded, and will generically hold, that $\frac{p^f(s_i)}{q^f(s_i)} \neq \frac{p^g(s_i)}{q^g(s_i)}$. This simple observation is the key to all that follows.

Assume that one is concerned with solving the problem of state-dependent utility as it has been initially raised under expected utility, i.e., with an act-independent form of state-dependence. This amounts to assuming away specifically act-dependent forms of state-dependent utility. Evidently, this is a logically weaker assumption than the exclusion of all forms of state-dependent utility, as needed by Savage, Ramsey, and others. Besides, specifically act-dependent forms of state-dependent utility do not seem to admit of any clear interpretation anyhow. Nevertheless, I emphasize that this needs to be listed here, as it should be elsewhere, as an additional assumption.[11] As the sequel illustrates, this intermediary assumption is well worth expliciting, instead of being left implicit as in the current literature.

Indeed, under this assumption, the simple observation following (6) immediately indicates a form of partial identification that has no counterpart in the baseline case of expected utility. Given a preference relation $\succcurlyeq$ representable as in (5) based on collection P, say that Q induces an alternative state-dependent representation if, through (6), Q induces the same state-dependent utility weights *for all acts*. The above observation indicates that, for any $\succcurlyeq$ representable as in (5) based on collection P, there exists some alternative collection Q that does not induce an alternative state-dependent representation. Here is a non-trivial example. Assume that $S = \{s_1, s_2, s_3\}$ and $P = \{p^1, p^2, p^3\}$, with probability vectors $p^1 = (\frac{1}{6}, \frac{1}{3}, \frac{1}{2})$, $p^2 = (\frac{1}{8}, \frac{1}{2}, \frac{3}{8})$, and $p^3 = (\frac{1}{12}, \frac{2}{3}, \frac{1}{4})$. Then, the alternative collection $Q = \{q^1, q^2, q^3\}$, with probability vectors $q^1 = (\frac{1}{4}, \frac{1}{4}, \frac{1}{2})$, $q^2 = (\frac{1}{2}, \frac{1}{6}, \frac{1}{3})$, $q^3 = (\frac{1}{5}, \frac{2}{5}, \frac{2}{5})$, does not induce an alternative state-dependent representation. This is because through (6), Q generates state-dependent utility functions that vary across acts—e.g., $\frac{p^1(s_1)}{q^1(s_1)} = \frac{2}{3} \neq \frac{1}{4} = \frac{p^2(s_1)}{q^2(s_1)}$. Accordingly, under the assumption that specifically act-dependent forms of state-dependent utility are excluded, the alternative collection $Q = \{q^1, q^2, q^3\}$ can be deemed incompatible with the observed choice data. Notice that would P and Q be singletons, like in the traditional expected utility model, no such partial identification could be claimed.

More remarkably, one can prove the following, stronger claim, which is the promised methodologically interesting possibility result. For some $\succcurlyeq$ representable as in (5) based on collection P (once again: under the assumption that specifically act-dependent forms of state-dependent utility are ex-

---

[11]One can show that, without this assumption, the problem is hopeless in any of the standard decision-theoretic frameworks. The reason is simple to explain. If both the probability measure and the utility function depend on the act, then the problem presented in Section 2 will recur *for each act*. Accordingly, without this assumption, the various solutions that have been offered—at different methodological costs—to the problem (such as the one in Drèze, 1987, but also the one in Karni and Schmeidler, 2016) would fail.

cluded), there does not exist any alternative state-dependent representation. When such is the case, the state-independence of the utility function in (5) is entailed by the choice data, rather than merely compatible with it. Accordingly, P proves absolutely unique, rather than merely unique relative to some admissible transformation of the utility function. This claim is the content of Proposition 1, the main insight of which goes back to Drèze, 1961. Proposition 1 is proved in the Appendix as a corollary of a stronger—but conceptually less transparent—characterization result.[12]

Proposition 1 *Assume that $\succcurlyeq$ is representable as in (5) based on a finite full support collection P. If P contains $\big|S\big|$ linearly independent probability vectors, then there exists no alternative state-dependent representation.*

For instance, if P is taken as in the previous example, then the sufficient condition in the proposition is not satisfied. In fact, one can go on and check that the collection $Q = \{q^1, q^2, q^3\}$ with probability vectors $q^1 = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, $q^2 = (\frac{1}{4}, \frac{1}{2}, \frac{1}{4})$, $q^3 = (\frac{1}{6}, \frac{2}{3}, \frac{1}{6})$, would induce an alternative state-dependent representation. But if one takes instead $P = \{p^1, p^2, p^3\}$ with probability vectors $p^1 = (\frac{1}{6}, \frac{1}{3}, \frac{1}{2})$, $p^2 = (\frac{5}{8}, \frac{1}{8}, \frac{1}{4})$, and $p^3 = (\frac{3}{10}, \frac{1}{10}, \frac{3}{5})$, for instance, then the sufficient condition is satisfied. Under the above assumption, then, no alternative state-dependent representation exists. In the context of the present section, the sufficient condition in Proposition 1 can be interpreted as a richness condition on the agency of the decision-maker. It indicates a maximal number of degrees of freedom in how she can exert her influence over the state space. Proposition 1 thus offers an explanation of why, under some assumptions, moral hazard suffices to solve the problem of state-dependent utility.

5   The problem of state-dependent utility: a reappraisal

The previous section provided a simple argument establishing that, under some assumptions, moral hazard is sufficient to solve the problem of state-dependent utility. The same argument can now be used to show that, under the same assumptions, moral hazard is not necessary to solve the problem, and non-expected utility is structurally less exposed to the problem than expected utility (the former being naturally understood, in this statement, as excluding the latter). In a nutshell, the reason for this is as follows. Representation (5) is compatible with, but does not entail, an interpretation

---

[12]Within the limits of the assumptions made earlier in the section, the result improves on Drèze's (compare with Drèze, 1987, Thm. 8.2) in the following respects. First, it explicits an assumption—the exclusion of specifically act-dependent forms of state-dependent utility—that is necessary for the result to go through. Second, the proof is shorter, simpler, and self-contained. Third, it covers several decision rules under moral hazard, while Drèze's result focuses on only one of them (see Drèze, 1987, Thm. 8.1, as well as fn. 13).

in terms of moral hazard. It is also a mould in which most non-expected utility models can be cast. The standard account of these models, on the other hand, does not require that the decision-maker have any influence over the resolution of uncertainty, which completes the argument.

In more detail, the key fact— that is of interest beyond the state-dependent utility literature—is as follows. Most non-expected utility models can be seen as departing from expected utility *exactly* in that they operate with act-dependent, rather than act-independent, probabilities. Consider, for example, the familiar maxmin expected utility model.[13] This is a non-expected utility model generalizing representation (1) in that the decision-maker is endowed not with just one prior, but with a (typically: closed and convex) set of priors. Denote this set by Π. The fact that Π does not reduce to a singleton is not meant to indicate that the decision-maker has any richer agency than customary in decision theory. Indeed, perhaps the most widespread interpretation is that this indicates, rather, her being more uncertain about the true state of nature, than the single-prior framework of expected utility can express. The model is called "maxmin expected utility" because the decision rule built in the representation is that the decision-maker maximizes, over Π, the minimum expectation of some state-independent utility function u. Accordingly, the model is a particular case of representation (5), with the definition $p^f = \arg\min_{p \in \Pi} \Big(\sum_{i=1}^{n} p(s_i) \cdot u\big(f(s_i)\big) \Big)$. Therefore, under the assumptions mentioned in the previous section, all that has been said regarding (5) in general can be applied to the maxmin expected utility model in particular. To wit, if Π is finitely generated by a set P satisfying the conditions in Proposition 1, under these assumptions, only the state-independent utility representation will be compatible with the observed choices of the decision-maker. Consequently, only the set Π featured in that representation—i.e., the act-independent set from which the act-dependent expected utility minimizers are taken—will correctly quantify her beliefs.

In the previous argument, nothing hinges on the specific properties of the maxmin expected utility model. An equally good (albeit more advanced) illustration would be the Choquet expected utility model, that can be explained to be a "rank-dependent" generalization of the baseline representation (1). Indeed, through a different mechanics than the one followed in maxmin expected utility, Choquet expected utility also leads to representa-

---

[13] This is a useful example to consider not just because maxmin expected utility is so familiar. It is also the case that Drèze's pioneering model of moral hazard is behaviorally indistinguishable from the dual of it, i.e., the maxmax expected utility model (see Drèze, 1987, Thm. 8.1). As this dual model can be interpreted without any reference to moral hazard, the claim of Drèze that has been mentioned on p. 7 must be rejected, even within his own model.

tion (5) (e.g., Köbberling and Wakker, 2003, Sec. 4.2; also Buchak, 2013).[14] In fact, most non-expected utility models are, likewise, special cases of representation (5) (see Chambers and Echenique, 2016, Sec. 8.4; also Cerreia-Vioglio et al., 2011, Cor. 3). By construction, merely by their being act-dependent probability representations, such non-expected utility representations can be argued—once again: under the assumption that specifically act-dependent forms of state-dependent utility are excluded—to achieve a form of partial identification that has no counterpart in the case of expected utility. Besides, whenever their belief parameter can be characterized by a finite collection P of act-dependent probability measures satisfying the richness condition in Proposition 1, this proposition can be invoked to argue that the decision-maker's beliefs are fully identified. As the standard interpretation of these models does not require, and typically excludes, that the decision-maker have any influence over the resolution of uncertainty, I conclude that both the general problem view and the radical solution view should be abandoned.

6 Implications for the revealed preference methodology

What are the methodological implications of realizing that the general problem view and the radical solution view should be abandoned? Arguably, the most important lesson is simply the following. The problem of state-dependent utility is more model-sensitive, and expected-utility-specific, than currently understood in the literature. Outside expected utility, under some assumptions, the problem can be solved. Within expected utility, under the same assumptions, it cannot. It would be too restrictive to say that state-dependent utility is a problem for revealing Bayesian priors only. But it would not be informative enough to say that it is a problem for revealing beliefs of any kind, in general.

Is this all good news for the revealed preference methodology, however? Undeniably, there is some good news, and it was much awaited for in this context. Some belief models can, in some cases, be given acceptable behavioral foundations. This need not come at the expense of such a significant change of the name of the game, as to allow for moral hazard. Thus, on the own terms of those sketching it, one can escape the dilemma presented at the end of Section 3. Naturally, a philosophical appreciation of the revealed preference methodology does not reduce to a discussion of the behavioral identification of beliefs in decision theory under uncertainty (e.g., Sen, 1973;

---

[14]Let $\nu : 2^S \to \mathbb{R}$ be a "capacity", as defined in Choquet expected utility theory. Given any f, let σ be a permutation on $\{1, \ldots, n\}$ such that $f(s_{\sigma_1}) \succcurlyeq \cdots \succcurlyeq f(s_{\sigma_n})$, with the notation $\sigma_j = \sigma^{-1}(j)$. Then, Choquet expected utility is the special case of (5) where p^f is defined by $p^f(s_i) = \left(\nu \left(\bigcup_{j=1}^{\sigma(i)} s_{\sigma_j} \right) - \nu \left(\bigcup_{j=1}^{\sigma(i)-1} s_{\sigma_j} \right) \right)$, with the convention $\bigcup_{j=1}^{0} s_{\sigma_j} = \varnothing$.

Hausman, 2000). But it must be said clearly and stressed that reassuring, non-paradoxical possibility results can be claimed for this methodology in this branch of decision theory.

On the other hand, the bigger picture hardly delivers the kind of good news the revealed preference methodologist could have hoped for. First, like the analysis which it rebuts, the foregoing analysis relies on the assumption that utility cannot depend on the states in a specifically act-dependent way. This is logically weaker than, but methodologically comparable to, the assumption that utility cannot depend on the states in any way whatsoever. Therefore, the offered solution cannot be claimed to constitute more than a partial solution to the problem.

Second, even granting the above assumption, within the classical expected utility model, the problem of state-dependent utility remains unsolvable, i.e., the decision-maker's beliefs are not identifiable. Given the special normative status of this model, this is already, in itself, an unpleasant fact to live with.

Third, even putting aside the special status of expected utility, one should not misappreciate the contrast with non-expected utility, where, it has been explained, the problem proves to be less acute in all cases and to vanish in some cases. To start with, it is now well-appreciated in the literature that there are substantial conceptual asymmetries between expected and non-expected utility. To wit, while it is almost unproblematic to interpret the priors of expected utility as measuring the decision-makers' beliefs, there are problems with similarly interpreting, say, the sets of priors of maxmin expected utility or the capacities of Choquet expected utility (see, e.g., Hill, 2019, Sec. 3.3). Specifically, when they come from choice data, these more general forms of measure appear to amalgamate both beliefs, that are cognitive attitudes, and some conative attitudes, such as uncertainty aversion and the like. Thus, if one is reluctant to reduce the belief identification problem under scrutiny to a sheer question of uniqueness, there is a sense in which the beliefs of a decision-maker become transparent only when they are also at least partially obscured by the presence of other mental attitudes.

Next, even accepting that the problem be boiled down to a sheer question of uniqueness, the very contrast between expected and non-expected utility indicates that the answer to the question "can beliefs be uniquely revealed from choice data?" will depend on the form, and even the content, of the beliefs examined. The form—If the beliefs of the decision-maker appear as some single Bayesian prior over the state space in conjunction with the expected utility rule, then the answer is "Not at all"; otherwise, the answer is "Perhaps". The content—For a given non-Bayesian model of beliefs, if the beliefs of the decision-maker happen to satisfy the richness condition given in the result of which Proposition 1 is a corollary, then the answer is "Yes"; otherwise, the answer is "Not fully". From a methodological point of view, the simple fact that the question above does not admit of a more

straightforward answer is an admission of failure.

Therefore, the take-away of the present paper is certainly not that the problem of state-dependent utility is methodologically less serious than currently presented in the literature (especially Karni, 1996; Baccelli, 2017). It is only that it is more complex. The problem of state-dependent utility still stands as a formidable challenge for the revealed preference methodology. Admittedly, the project of providing the various models of beliefs with behavioral foundations is not—pace Nau, 2001—entirely hopeless. But it remains fraught with considerable difficulties, that cannot be philosophically ignored.

7 Conclusion

One important stream of literature in the behavioral branch of decision theory under uncertainty is tasked with delivering revealed preference foundations for the various models of beliefs. However, its capacity to fulfil that mission is challenged by the problem of state-dependent utility. The uniqueness properties of the proposed representations seem simply too weak to fit the bill. According to the current literature, one can be more assertive about the extent of the problem. First, the problem of state-dependent utility arises for expected and non-expected utility alike, i.e., whatever the targeted model of beliefs. Second, the problem is solvable only under moral hazard, i.e., the assumption that the decision-maker can influence the resolution of uncertainty. The overall methodological situation is, then, almost aporetic.

I have given grounds to reject the two claims above. More specifically, under assumptions that do not set my paper apart from the rest of the relevant literature, I have reached the following two conclusions. First, non-expected utility is less exposed to the problem of state-dependent utility than expected utility (the former being naturally understood, in this statement, as excluding the latter). The problem is more model-sensitive, and expected-utility-specific, than currently assumed. Second, the beliefs of a decision-maker can be behaviorally identified even though the decision-maker has no influence over the resolution of uncertainty. The very assumptions under which moral hazard proves sufficient to solve the identification problem can be used to show that it is not necessary to this effect. Arguably, the restrictiveness of these assumptions nothwithstanding, this constitutes a significant reappraisal of the problem. But it is at least as important to realize that this reappraisal brings only limited good news for the revealed preference methodology. Among other considerations, the answer to the question "can beliefs be uniquely revealed from choice data?" turns out to depend on the form, and even the content, of the decision-maker's beliefs. This is, in itself, an admission of methodological failure.

8   Appendix

Proof of Proposition 1. $P = \{p^1, \ldots, p^m\}$ is a collection of probability measures that have full support over the state space $S = \{s_1, \ldots, s_n\}$. For non-triviality, assume that $m \geq n$. Of interest are the conditions under which there exists another collection $Q = \{q^1, \ldots, q^m\}$ of full support probability measures such that, for all $j = 1, \ldots, n$, and any $i, i' \in \{1, \ldots, m\}$, we have:

$$\frac{p^i(s_j)}{q^i(s_j)} = \frac{p^{i'}(s_j)}{q^{i'}(s_j)}. \tag{7}$$

Equivalently, taking p^1 and q^1 as units, the question regards the existence of another collection $Q = \{q^1, \ldots, q^m\}$ of full support probability measures such that, for any $i = 1, \ldots, n$ and $j = 1, \ldots, n$, we have:

$$q^i(s_j) = \frac{p^i(s_j)}{p^1(s_j)} \cdot q^1(s_j). \tag{8}$$

Thus, the problem consists in finding one full support probability measure q^1 such that the q^i induced by (8) are probability measures. By assumption, $\frac{p^i(s_j)}{p^1(s_j)} > 0$ for all i and j. Thus, taking one full support probability measure q^1, the only question is whether the q_i induced by (8) sum up to 1 over S.

To investigate this, construct the m by n matrix $\boldsymbol{A} = \left[a^i_j \right]^{i=1,\ldots,m}_{j=1,\ldots,n}$, with $a^i_j = \frac{p^i(s_j)}{p^1(s_j)}$. Denote by $\boldsymbol{x}$ the (n by 1) probability vector induced by p^1. By assumption, we have $\boldsymbol{Ax} = \boldsymbol{1}$ (with $\boldsymbol{1}$ a m by 1 vector). The question is whether there exists a strictly positive probability vector $\boldsymbol{y}$, with $\boldsymbol{y} \neq \boldsymbol{x}$, such that $\boldsymbol{Ay} = \boldsymbol{1}$. Were such a $\boldsymbol{y}$ to exist, we would have $\boldsymbol{A}(\boldsymbol{x} - \boldsymbol{y}) = \boldsymbol{A}(\boldsymbol{y} - \boldsymbol{x}) = \boldsymbol{0}$, with non-zero vectors $(\boldsymbol{x} - \boldsymbol{y})$, $(\boldsymbol{y} - \boldsymbol{x})$. Therefore, the existence of such a $\boldsymbol{y}$ is characterized by the following condition: There exists a vector $\boldsymbol{z}$ in the null space of $\boldsymbol{A}$ such that for all $j = 1, \ldots, n$, either $x_j - z_j > 0$ or $x_j + z_j > 0$. By the Rank-Nullity Theorem, for such a $\boldsymbol{z}$ to exist, it is necessary that $\mathrm{rank}(\boldsymbol{A}) < n$. Thus, for no such $\boldsymbol{z}$ to exist, if suffices that $\mathrm{rank}(\boldsymbol{A}) = n$. In other words, the following is a corollary of the above characterization result. For no suitable alternative collection Q to exist, it suffices that the collection P contains $\left| S \right|$ linearly independent probability vectors. $\square$

References

ANSCOMBE, F., AND R. AUMANN (1963): "A Definition of Subjective Probability", *The Annals of Mathematical Statistics*, 34(1), 199–205.

BACCELLI, J. (2017): "Do Bets Reveal Beliefs?", *Synthese*, 194(9), 3393–3419.

———— (forthcoming): "Moral Hazard, the Savage Framework, and State-Dependent Utility", *Erkenntnis*.

BRADLEY, R. (2004): "Ramsey's Representation Theorem", *Dialectica*, 58(4), 483–497.

———— (2017): *Decision Theory with a Human Face*. Cambridge: Cambridge University Press.

BUCHAK, L. (2013): *Risk and Rationality*. Oxford: Oxford University Press.

CERREIA-VIOGLIO, S., P. GHIRARDATO, F. MACCHERONI, M. MARINACCI, AND M. SINISCALCHI (2011): "Rational Preferences under Ambiguity", *Economic Theory*, 48(2-3), 341–375.

CHAMBERS, C., AND F. ECHENIQUE (2016): *Revealed Preference Theory*. New York: Cambridge University Press.

DAVIDSON, D., AND P. SUPPES (1956): "A Finitistic Axiomatization of Subjective Probability and Utility", *Econometrica*, 24(3), 264–275.

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): "Logical Foundations of Cardinal Utility and Subjective Probability", in *Essays on Economic Decisions under Uncertainty*, ed. by J. Drèze, pp. 90–104. Cambridge: Cambridge University Press.

DRÈZE, J., AND A. RUSTICHINI (2004): "State-Dependent Utility Theory", in *Handbook of Utility Theory, Volume II: Extensions*, ed. by S. Barbera, P. Hammond, and C. Seidl, pp. 839–892. Boston: Kluwer Academic Press.

GALAABAATAR, T., AND E. KARNI (2013): "Subjective Expected Utility with Incomplete Preferences", *Econometrica*, 81(1), 255–284.

GILBOA, I., AND D. SCHMEIDLER (1989): "Maxmin Expected Utility with Non-Unique Prior", *Journal of Mathematical Economics*, 18(2), 141–153.

HANDS, W. (2013): "Foundations of Contemporary Revealed Preference Theory", *Erkenntnis*, 78(5), 1081–1108.

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.

——— (2019): "Confidence in Beliefs and Rational Decision-Making", *Economics & Philosophy*, 35(2), 223–258.

JEFFREY, R. (1983): *The Logic of Decision*. Chicago: Chicago University Press (Second Edition).

JOYCE, J. (1999): *The Foundations of Causal Decision Theory*. New York: Cambridge University Press.

KARNI, E. (1993): "Subjective Expected Utility Theory with State-Dependent Preferences", *Journal of Economic Theory*, 60(2), 428–438.

——— (1996): "Probabilities and Beliefs", *Journal of Risk and Uncertainty*, 13(3), 249–262.

——— (2008): "State-Dependent Utility", in *The Handbook of Rational and Social Choice*, ed. by P. Anand, P. Pattanaik, and C. Puppe, pp. 223–238. Oxford: Oxford University Press.

——— (2011): "A Theory of Bayesian Decision Making with Action-Dependent Subjective Probabilities", *Economic Theory*, 48(1), 125–146.

KARNI, E., AND D. SCHMEIDLER (2016): "An Expected Utility Theory for State-Dependent Preferences", *Theory and Decision*, 81(4), 467–478.

KÖBBERLING, V., AND P. WAKKER (2003): "Preference Foundations for Nonexpected Utility: A Generalized and Simplified Technique", *Mathematics of Operations Research*, 28(3), 395–423.

KREPS, D. (1988): *Notes on the Theory of Choice*. Boulder: Westview Press.

KYBURG, H., AND H. SMOKLER (eds.) (1980): *Studies in Subjective Probability*. New York: Robert E. Krieger Publishing Company.

LAFFONT, J.-J., AND D. MARTIMORT (2002): *The Theory of Incentives: The Principal-Agent Model*. Princeton: Princeton University Press.

LEVI, I. (1974): "On Indeterminate Probabilities", *The Journal of Philosophy*, 71(13), 391–418.

——— (1986): *Hard Choices: Decision-Making under Unresolved Conflict*. Cambridge: Cambridge University Press.

NAU, R. (2001): "de Finetti Was Right: Probability Does Not Exist", *Theory and Decision*, 51(2-4), 89–124.

RAMSEY, F. (1931): "Truth and Probability", in RAMSEY, F., *The Foundations of Mathematics and Other Logical Essays*, pp. 156–198.

SAVAGE, L. (1954): *The Foundations of Statistics*. New York: Wiley (First Edition).

——— (1972): *The Foundations of Statistics*. New York: Dover (Second Edition).

SCHERVISH, M., T. SEIDENFELD, AND J. KADANE (1990): "State-Dependent Utilities", *Journal of the American Statistical Association*, 85(411), 840–847.

——— (2013): "The Effect of Exchange Rates on Statistical Decisions", *Philosophy of Science*, 80(4), 504–532.

SCHMEIDLER, D. (1989): "Subjective Probability and Expected Utility without Additivity", *Econometrica*, 57(3), 571–587.

SEIDENFELD, T., M. SCHERVISH, AND J. KADANE (1990): "When Fair Betting Odds Are Not Degrees of Belief", *Philosophy of Science*, 1990(1), 517–524.

SEN, A. (1973): "Behaviour and the Concept of Preference", *Economica*, 40(159), 241–259.

WAKKER, P., AND H. ZANK (1999): "State Dependent Expected Utility for Savage's State Space", *Mathematics of Operations Research*, 24(1), 8–34.