Theory Dec. (2009) 67:405–432
DOI 10.1007/s11238-008-9127-x

Living without state-independence of utilities

Brian Hill

Published online: 23 December 2008

Abstract This article is concerned with the representation of preferences which do not satisfy the ordinary axioms for state-independent utilities. After suggesting reasons for not being satisfied with solutions involving state-dependent utilities, an alternative representation shall be proposed involving state-independent utilities and a *situation-dependent factor*. The latter captures the interdependencies between states and consequences. Two sets of axioms are proposed, each permitting the derivation of subjective probabilities, state-independent utilities, and a situation-dependent factor, and each operating in a different framework. The first framework involves the concept of a *decision situation*—consisting of a set of states, a set of consequences and a preference relation on acts; the probabilities, utilities and situation-dependent factor are elicited by referring to other, appropriate decision situations. The second framework, which is technically related, operates in a fixed decision situation; particular "subsituations" are employed in the derivation of the representation. Possible interpretations of the situation-dependent factor and the notion of situation are discussed.

Keywords Elicitation · Subjective probability · Subjective expected utility · State-dependent utility · Small worlds

JEL Classification D81 · C60

B. Hill (✉)
GREGHEC,
HEC Paris, 1 Rue de la Libération,
78351 Jouy-en-Josas, France
e-mail: hill@hec.fr

1 Introduction

It has been known for many years that the axioms implying the state-independence of utilities which appear in many classical theories of decision under uncertainty, such as those of Savage (1954) and Anscombe and Aumann (1963), do not hold in many important situations, such as cases of life insurance or health insurance (Arrow 1974; Cook and Graham 1977). The conclusion normally drawn is that a representation involving *state-dependent* utilities is required; thus the work on this subject (for example, Karni et al. 1983; Karni and Mongin 2000; Karni and Schmeidler 1993; Karni 2008; Drèze 1987; Drèze and Rustichini 2004). However, this is not the only option; in this article, an alternative representation shall be proposed, and a representation theorem for it proved.

To introduce the problems posed by violations of the traditional state-independence axioms, and the motivation for representation theorems for state-dependent utility, consider an example closely related to that proposed by Aumann in his correspondence with Savage (reproduced in Drèze 1987, pp. 76–81; this version is taken from Karni 1996). A woman is to undergo a potentially fatal operation, with 50–50 chances of survival. Her husband, who knows the odds, is offered two bets: one pays \$100 if the operation is successful and nothing if not; the other pays \$100 if the operation is not successful and nothing if it is. A common representation of the bets is given in Table 1. They are considered to be acts: functions from a set of states—taken to be success and failure of the operation—to a set of consequences (elements to which the husband allocates utility values)—taken to be winning \$100 and winning \$0. Although the chances of winning the money are equal in both cases, the husband would apparently choose the former bet, because the \$100 would be more precious to him if his wife were with him to enjoy it.

Such behaviour may pose two distinct challenges for state-independent utility theories. First of all, it may cast doubt on the accuracy of such theories as representations of *behaviour*. Secondly, it may cast doubt on the claim that such theories *elicit* the agent's beliefs and desires, in the form of probabilities and utilities. The second challenge is more widespread and more serious than the first.

Consider, firstly, the extent of the challenges. There are many cases where agents who make the choices described in the example have preferences which do not satisfy some axioms of the standard theory—most notably, cases where the axioms ensuring state-independence of utility (P3 and P4 in Savage (1954) and Monotonicity in Anscombe and Aumann (1963); see Sect. 2.1) are violated. Such cases pose problems for the behavioural and the elicitation claims of the theory: on the one hand, these are cases where the theory is shown to be behaviourally inaccurate; on the other hand, since the axioms need to be satisfied for the theorems to apply, and for the probability and utility to be elicited using them, these are cases where the theory cannot be used to elicit the agent's attitudes.

Table 1 2-consequence formulation of decision problem

		Success	Failure
A	Bet on success	\$100	\$0
B	Bet on failure	\$0	\$100

However, although Aumann and Savage assumed in their discussion that the state-independence axioms are violated, the pattern of choices is not in fact incompatible with satisfaction of these axioms. Karni (1996) gives a development of Aumann's example where all the standard Savage axioms are satisfied, so that probabilities and utilities can be elicited using Savage's theorem. The probability function elicited assigns more weight to the success of the operation than to its failure; given that the utilities are state-independent, this is the only way to account for the husband's strict preference over the bets. Such examples do not pose a particular challenge for Savage's theory on the behavioural front: to the extent that all the axioms are satisfied, the theory adequately describes the agent's behaviour. However, as Karni points out, they seem to pose a challenge to the theory's pretension to elicit the agent's beliefs (and desires). The probability function elicited allocates different probabilities to the success and failure of the operation, and this contradicts the intuition that the husband believes success to be as likely as failure—an intuition which is sustained not only by the assumption that he is fully cognizant of the odds of survival, but also by some of his other actions (for example, he advises a friend, who is offered the same pair of bets, that the odds of winning are the same). Karni concludes that the probabilities furnished by Savage's theorem do not properly represent the husband's beliefs. This is a case where Savage's theory applies and provides a probability and utility function, but it is doubtful whether these accurately represent the agent's attitudes.

The behavioural problem is also less serious than the elicitation problem. Strictly speaking, one could reply to the former problem by resorting to a representation which does not rely on the state-independence axioms. In the case of expected utility, one obtains an additively separable representation, that is, a representation in terms of a real-valued function on pairs of states and consequences—which we shall call the *evaluation function*[1]—that is of the following form:

$$f \preceq g \text{ iff} \sum_{s\in S} U(s, f(s)) \leq \sum_{s\in S} U(s, g(s)) \tag{1}$$

At a push, one could decompose the evaluation function U into a probability and state-dependent utility ($U(s, x) = p(s) \cdot u(s, x)$). However, such a decomposition is entirely arbitrary: any (reasonable) probability function could be used, for example. Representations such as this, perhaps with arbitrary decompositions, are sufficient for many economic applications (though not all: see Karni and Schmeidler 1993, §4, for an example). However, they do not provide a viable solution to the elicitation problem: here, one needs a (preferably principled) way to decompose the evaluation function U into a *unique* probability function and a (suitably) *unique* utility function.

As regards this problem, which is doubtless the main challenge posed by the failure of state-independence axioms, two strategies have been proposed. The first is promoted by theorists working on state-dependent utilities. The intuition is that, in the example above, the utility of the consequences *depends* on the states: \$100 is more desirable if the wife survives than if she does not. Thus the representation should feature state-dependent utilities. Many state-dependent utility theorems are supposed

[1] Thanks to Mark Machina and Edi Karni for suggesting this term.

Table 2 4-consequence formulation of decision problem

		Success	Failure
A	Bet on success	$100 with wife alive	$0 without wife
B	Bet on failure	$0 with wife alive	$100 without wife
C	Bet on success (but wife murdered)	$100 without wife	$0 without wife
D	Bet on failure (but wife resuscitated)	$0 with wife alive	$100 with wife alive
E	Bet on failure (but wife murdered or resuscitated)	$0 without wife	$100 with wife alive

to provide an answer to the elicitation question: they yield a *unique* probability and a (suitably) *unique* state-dependent utility function, the uniqueness being necessary for the claim that these functions represent the agent's beliefs and utilities (Karni et al. 1983; Karni and Schmeidler 1993; Karni 1993a,b, 2007, 2008).

The second strategy is the one proposed by Savage in his reply to Aumann. He suggests representing the decision by the same set of states, but using as consequences the following four elements: $100 with survival, $100 with demise, $0 with survival, $0 with demise (Table 2). He then evokes the "make-believe" situation where the husband considers that any act (function from states to consequences) taking values in this set of consequences is an available option—and in particular, the acts which, in the case where the operations fails, would yield $100 and the wife returned to him in good health (acts D and E). In this situation, the axioms of state-independent utility apply and the classical representation theorems can be employed. These yield, according to a proponent of this strategy, the agent's beliefs and utilities (it is easy to show, in particular, that the problem raised by Karni's aforementioned development of the example no longer applies).

Each of these strategies has its price. Under the state-dependent utility approach, the agent's utilities for the consequences depend on the states involved; it thus becomes impossible to use these utilities in decision situations where the same set of consequences are on offer, but the states of the world are different. Consider the case where the agent is offered a bet on a horse race, with the consequences being appropriate combinations of $100 and $0: wouldn't one expect him to have the same utilities in this situation as in the one described in the story above? Under the state-dependent analysis, this cannot be the case, since the states on which the utilities are dependent (success and failure of the operation) are not involved. This naturally casts doubt on whether these theorems are yielding his "real" utilities, for one would expect him to have the same utilities for the same consequences considered in different situations (when betting on his wife's operation, when betting on a horse race, when making investment decisions, when taking out life insurance and so on), and this is not a property of the utilities that the state-dependent utility theorems provide.[2]

To understand the weakness of Savage's strategy, note that it operates in two stages. In the first stage, one rewrites the decision problem with an enlarged set of consequences; this corresponds to moving from Table 1 to Table 2. In the second stage, one

[2] See Sect. 3 for further discussion of the relationship between attitudes in different situations.

asks the decision maker to make his choice *as if* a certain number of the acts—such as acts involving murder or resuscitation (acts C–E in Table 2)—were not fantastic or ridiculous. The first stage of the strategy involves a simple change in the formulation of the decision problem and no behavioural change on the part of the agent. By contrast, the second stage involves an explicit request that the agent alters his behaviour. This can be seen by the fact that the agent, answering in the "as if" mode, is indifferent between the acts A and E in Table 2, whereas in reality he would never consider picking act E.

In the terminology which will be introduced in the next section, we will say that the second stage involves two different decision situations: both the decision situations have the same states and consequences (and thus acts), but the decision maker's preferences differ between the two. One of the decision situations is the one the agent is actually in (those are the preferences he actually has); the challenge is to understand his behaviour and attitudes in this situation. The other decision situation is "make-believe"—the preferences he has in that situation are hypothetical, and conflict with his real preferences. That is the situation where one can elicit his beliefs and utilities using standard state-independent utility results.

Although these are behaviourally distinct decision situations, for the beliefs and utilities in the former situation to be elicited in the way that Savage suggests, it needs to be assumed that the agent has the same attitudes in the two situations. What permits this assumption? An adequate answer to this question should make clear the relationship between the two situations. The simplest and most intuitive response is that the former situation is as the latter, with an added restriction on the set of acts available (namely, that the fantastic or ridiculous acts are not available). However, the status of this restriction is problematic. First of all, it is posed exogenously to the decision theory used. Furthermore, it is known that one cannot incorporate such restrictions endogenously into representation theorems such as Savage's, since these depend on the availability of a full set of acts (that is, all functions from states to consequences). Indeed, such restrictions generally lead to representations in terms of state-dependent utilities, which is exactly what Savage's strategy was meant to avoid (Hammond 1998, §6 considers the case, pertinent for this example, where the consequences attainable from different states are distinct). To sum up, the strategy proposed by Savage relies on a decision situation where the agent shows preferences which differ from those he actually has; the weakness in the strategy is that, although the strategy relies on assumptions about the relationship between this situation and the situation the decision maker is actually in (his real preferences), this relationship is poorly understood. Indeed, any attempt to account for the relationship in the theory leads us back to state-dependent utilities.

Two morals can be drawn from these considerations. The discussion of Savage's strategy makes it clear that there is an (inevitable) interdependence between states and consequences in situations such as the one in the example. However, given what has been said about the solution offered by state-dependent utility theorems, it is not clear that this interdependence should be built into the utilities. This seems to suggest that, if there is interdependence of states and consequences, then it must be specific to situations where these states and consequences are involved. A natural suggestion is to represent the interdependence by a *situation-dependent factor* expressing the relationship between states and consequences. The representation would then be as follows. For S the set of states involved in the situation, C the set of consequences and $\preceq$ a preference relation on acts (functions from S to C) there is a probability function p on S, a (state-independent) utility function u on C and a function γ on $S \times C$—the situation-dependent factor—such that, for any acts f, g,

$$f \preceq g \text{ iff} \sum_S p(s) \cdot \gamma(s, f(s)) \cdot u(f(s)) \leq \sum_S p(s) \cdot \gamma(s, g(s)) \cdot u(g(s)) \qquad (2)$$

The main aim of this article is to propose a representation theorem for this sort of representation: that is, to propose a set of conditions under which there is a *unique* probability and a (suitably) *unique* utility and situation-dependent factor such that (2) holds. This can be considered as another reply to the behavioural problem posed by examples such as Aumann's, insofar as it represents correctly the preferences of the agent, using not only the probabilities and utilities, but also the situation-dependent factor. It constitutes a reply to the elicitation problem, insofar as it yields (suitably) unique probabilities and utilities, which feature in the representation of preferences and can be thought of as capturing the agent's beliefs and desires. Indeed, the representation theorem relies on an elementary but attractive intuition on how to elicit probabilities and utilities: for elicitation, use situations which are as simple as possible. To elicit the probabilities and utilities in a situation where there is interdependence between states and consequences, the technique will be to look in situations where either the sets of states or sets of consequences are different, but where there is no interdependence between the states and the consequences, so that the traditional (state-independence) results can be applied.

In Sect. 2, a formal notion of situation shall be defined, postulates and axioms shall be proposed and a representation theorem shall be stated. In Sect. 3, the concepts and techniques introduced, as well as possible interpretations, applications and relations to existing approaches, shall be considered at length. It turns out that the use of other decisions situations is, to a large extent, an artifice of the way the decision problem is modelled; in Sect. 4, a representation theorem for (2) which operates in a single situation shall be stated. Proofs are to be found in the Appendix.

2 Situational version

2.1 Preliminaries and axioms

A decision situation is characterised by a set of states and a set of consequences, with a preference on the acts.

Definition 1 A *decision situation* σ consists of a set S_σ (of states in the situation), a set C_σ (of consequences in the situation) and a binary relation $\preceq_\sigma$ on the set of functions $\mathcal{A}_\sigma$ from S_σ into C_σ (preferences in the situation over the acts in the situation).

A set of situations $\mathcal{S}$ is assumed to be given; this is the set of decision problems which the agent might have faced at the current moment, with his preferences over the acts he would have faced. Two sorts of conditions will need to be proposed on $\mathcal{S}$ to obtain the representation. On the one hand, there will be general, more or less standard conditions guaranteeing that the agent's choices are consistent and that he is an expected utility maximiser (rather than a non-expected utility maximiser, for example). On the other hand, there will be particular conditions regarding the existence of specific types of situations, which are required for the elicitation technique. The two postulates stating the former conditions will be presented first; then we will turn to the axioms corresponding to the latter conditions.

Basic postulates. According to the definition above, $\mathcal{S}$ may contain two situations with the same states and consequences but different preferences. This represents an agent whose preferences are inconsistent: they may differ on the same set of acts. The theorem below does not apply to such agents. It employs a minimal consistency constraint on the agent's preferences which implies that he is not of this type: namely, that he can only have one order of preference on any set of acts envisaged.

Postulate 1 For all $\sigma_1, \sigma_2 \in \mathcal{S}$, if $S_{\sigma_1} = S_{\sigma_2}$ and $C_{\sigma_1} = C_{\sigma_2}$, then $\preceq_{\sigma_1} = \preceq_{\sigma_2}$.

This constraint allows the situations to be thought of extensionally; that is, any two distinct situations differ either in their sets of states or in their sets of consequences (or both).

Remark 1 There is another consistency condition which may be imposed, and which relates to the possibility of "identifying" acts between situations. An act is a function from states to consequences: there is thus no (evident) way of saying that two acts are the same if they belong to situations with different sets of states (they have different domains). By contrast, it does seem that one could identify acts which belong to situations with the same set of states and which give the same consequences (they have the same domain and the same image), although the sets of consequences available in the situations to which they belong differ (they have different ranges). Consider, for example, $\sigma = (S_\sigma, C_\sigma, \preceq_\sigma)$ and $\sigma' = (S_{\sigma'}, C_{\sigma'}, \preceq_{\sigma'})$, where $S_\sigma = S_{\sigma'}$ and $C_\sigma \subset C_{\sigma'}$: it seems intuitive to identify an act in σ with the act in σ' that gives the same consequences for each state, and thus to demand that the preferences over the acts in σ are the same as the preferences over the corresponding acts in σ'. This is expressed by the following axiom.

Axiom A1 For any pair of situations $\sigma_1 = (S_{\sigma_1}, C_{\sigma_1}, \preceq_{\sigma_1})$ and $\sigma_2 = (S_{\sigma_2}, C_{\sigma_2}, \preceq_{\sigma_2})$ with $S_{\sigma_1} = S_{\sigma_2} = S$ and $C_{\sigma_1} \cap C_{\sigma_2} \neq \emptyset$, $\preceq_{\sigma_1} |_{\mathcal{A}_{\sigma_1} \cap \mathcal{A}_{\sigma_2}} = \preceq_{\sigma_2} |_{\mathcal{A}_{\sigma_1} \cap \mathcal{A}_{\sigma_2}}$.

(Standard mathematical notation shall be employed whereby, for an order $\leq$ on a set Y and $Y' \subseteq Y$, $\leq |_{Y'}$ is the restriction of $\leq$ to Y'.)

This axiom shall *not* be necessary for the results in this section, and shall not be assumed to hold unless explicitly stated. However, we shall have reason to discuss it in Sects. 3 and 4.

The decision-theoretic framework used here is that proposed in Anscombe and Aumann (1963), which shall now be briefly summarised.

For a given situation σ, S_σ is assumed to be finite, and C_σ is assumed to be the set of lotteries over a finite set X_σ—the set of outcomes. Acts—functions from states to lotteries—are thought of as functions from $S_\sigma \times X_\sigma \rightarrow \Re$; that is, the set of acts is $\mathcal{A}_\sigma = \{f : S_\sigma \times X_\sigma \rightarrow \Re | \sum_{x \in X} f(s,x) = 1\}$.[3] Under these assumptions, $\mathcal{A}_\sigma$ is a mixture set with the mixture relation defined pointwise: for f, h in $\mathcal{A}_\sigma$ and $a \in \Re$, $0 < a < 1$, the mixture $af + (1-a)h$ is defined by $(af + (1-a)h)(s,x) = af(s,x) + (1-a)h(s,x)$ (Fishburn 1970, Chap. 13). So the von Neumann-Morgenstern axioms can be stated for the preference orders $\preceq_\sigma$.

Axiom A2 (*Weak order*) (a) For all f, g in $\mathcal{A}_\sigma$, $f \preceq_\sigma g$ or $g \preceq_\sigma f$. (b) For all f, g and h in $\mathcal{A}_\sigma$, if $f \preceq_\sigma g$ and $g \preceq_\sigma h$, then $f \preceq_\sigma h$.

Axiom A3 (*Independence*) For all f, g and h in $\mathcal{A}_\sigma$, and for all $a \in \Re$, $0 < a < 1$, if $f \preceq_\sigma g$ then $af + (1-a)h \preceq_\sigma ag + (1-a)h$.

Axiom A4 (*Continuity*) For all f, g and h in $\mathcal{A}_\sigma$, if $f \preceq_\sigma g$ and $g \preceq_\sigma h$, then there exist $a, b \in (0, 1)$ such that $af + (1-a)h \preceq_\sigma g$ and $g \preceq_\sigma bf + (1-b)h$.

Anscombe and Aumann (1963) add the following axiom:

Axiom A5 (*Monotonicity*) For any elements c_1 and c_2 of C_σ, any f in $\mathcal{A}_\sigma$ and any s in S_σ, if $\mathbf{c_1} \preceq_\sigma \mathbf{c_2}$ then $f_s^{c_1} \preceq_\sigma f_s^{c_2}$

where $\mathbf{c}$ is the constant act taking value c (for all $s \in S_\sigma$, $f(s,x) = c(x)$) and f_s^c is identical to f, except on s, where it takes value c (that is, $f_s^c(s,x) = c(x)$ and, for $s' \neq s$, $f_s^c(s',x) = f(s',x)$).

Fishburn (1970, p. 146) has shown that Weak Order, Independence and Continuity imply that there is an additively separable representation of $\preceq_\sigma$ (a representation of form (1); see Sect. 1), such that the evaluation function is unique up to simple positive affine transformation: for any pair of functions U and U' representing $\preceq_\sigma$, there is a positive real number a and real numbers b_s for each $s \in S_\sigma$ such that $U'(s,x) = aU(s,x) + b_s$ for all $s \in S_\sigma$ and $x \in X_\sigma$. Anscombe and Aumann (1963) show furthermore that adding the Monotonicity axiom yields a decomposition of this function into a unique probability function over the set of states and a state-independent utility function over the set of outcomes which is unique up to positive affine transformation. Indeed, the Monotonicity axiom is so closely connected to state independence that it is often simply called the axiom of state independence (in Hammond (1998), for example).

This study will follow much of the literature on state dependence of utilities and suppose that the expected utility axioms other than those yielding state independence hold in every situation. This is expressed by the following postulate.

Postulate 2 For any situation $\sigma \in \mathcal{S}$, $\preceq_\sigma$ satisfies Weak Order, Independence and Continuity.

[3] In thinking of acts in this way, Anscombe and Aumann's "Reversal of Order" axiom is assumed to hold (Anscombe and Aumann 1963). As is standard in much work with this framework, this axiom shall be assumed to hold in all situations, throughout this article. The consequences of rejecting it are explored by Drèze (1987), for example.

In the light of what has been said above, this postulate implies that the agent always maximises an additive evaluation function, whether or not this function can be decomposed (in a non-arbitrary manner) into a probability function and a utility function. Behaviourally, an agent satisfying this postulate always acts as if he is an expected utility maximiser (rather than a non-expected utility maximiser), but, from an elicitation point of view, satisfying the postulate is not enough to allow identification of his probabilities and utilities.

The situations where there exists a decomposition of the evaluation function into a probability and state-independent utility function are of special importance.

Definition 2 σ is called a *simple* situation if and only if $\preceq_\sigma$ also satisfies Monotonicity.

The simple situations are those which permit the application of the Anscombe and Aumann representation theorem. Let us state this explicitly.

Theorem 1 (Anscombe and Aumann 1963) *For a simple situation σ, there is a unique probability distribution p_σ on S_σ and a utility function u_σ on X_σ which is unique up to positive affine transformation, such that, for $f, g \in \mathcal{A}_\sigma$*

$$f \preceq_\sigma g \text{ iff } \sum_{S_\sigma, X_\sigma} p_\sigma(s)u_\sigma(x)f(s,x) \leq \sum_{S_\sigma, X_\sigma} p_\sigma(s)u_\sigma(x)g(s,x) \tag{3}$$

An unsympathetic reading of the literature would have it that *all* situations are simple, so that theorems such as this one can always be applied (Savage 1954, §5.5, for example, seems to assume that his axioms apply in all "small world" decision situations). This cannot be assumed here, for the challenge is to elicit probabilities and utilities in non-simple situations (i.e. situations where the state-independence axioms do not apply). A milder assumption will be made: that there are *enough* simple situations to permit elicitation of (unique) probabilities and utilities. The axioms in the representation theorem will cash out this assumption in formal terms. First, it is necessary to introduce some preliminary definitions.

Definition 3 For a set of consequences C, let $\Sigma_C = \{\sigma \in \mathcal{S} | \ C = C_\sigma, \sigma \text{ simple}\}$. This is the set of simple situations with C as set of consequences.

Define $\|\Sigma_C\| = \{u | \ \exists\sigma \in \Sigma_C, p \text{ on } S_\sigma, \text{ s.t.} p, u \text{ represent } \preceq_\sigma\}$, the set of utilities involved in the representation of the preferences (according to (3)) in the simple situations. By Theorem 1 and Definition 2, $\|\Sigma_C\|$ is non-empty if Σ_C is. The elements of this set will be considered up to positive affine transformation, to avoid unnecessary repetitions.

Similarly, for a set of states S, define $\Xi_S = \{\sigma | \ S = S_\sigma, \sigma \text{ simple}\}$, and $\|\Xi_S\| = \{p | \ \exists\sigma \in \Xi_S, u \text{ on } C_\sigma, \text{ s.t.} p, u \text{ represent } \preceq_\sigma\}$.

The idea is to use the representations of the preferences in simple situations to ascertain the utility in non-simple situations which have the same set of consequences; similarly for probabilities. For this several existence and uniqueness constraints are required on the set of simple situations; these will be the main axioms in the representation theorem. The axioms are presented here; they will be discussed at greater length in Sect. 3.

Main axioms. Let us first deal with the situations sharing the same consequences. Consider the following two axioms.

Axiom A6 (*Consequence Richness*) For all C, Σ_C is non-empty.

Axiom A7 (*Conative Consistency*) For all C, for any $\sigma_1, \sigma_2 \in \Sigma_C$, $\preceq_{\sigma_1}^{\text{const}} = \preceq_{\sigma_2}^{\text{const}}$, where $\preceq^{\text{const}}$ is the restriction of $\preceq$ to the constant acts.

Consequence Richness states that there are enough situations such that, for each set of consequences, an "independent" set of states can be found—independent in the sense that Anscombe and Aumann's state-independence axiom, and thus their expected utility result, applies. As such, it can be considered to be a structural or technical axiom: its behavioural consequences are negligible, and, due to the existential quantifier, it is difficult to refute. From a technical point of view, this axiom guarantees that $\|\Sigma_C\|$ has at least one element. See Sect. 3 for an extended discussion of the plausibility of the axiom.

Conative Consistency is equivalent to the existence of at most one element in $\|\Sigma_C\|$ (Lemma 1 in the Appendix). By contrast to the previous axiom, it does have some behavioural content: it demands that the agent's preferences over constant acts—acts which give the same result on all states—is independent of what the states are (and thus, given Definition 1 and Postulate 1, of what the situation is). An agent who prefers \$50 to a 50–50 objective lottery yielding \$100 when in the context of a horse race (states: winners of the race) but prefers the objective lottery to the sure amount in the context of investment on the markets (states: stock price tomorrow) violates this axiom. Certainly, such preferences do not seem particularly consistent: he prefers receiving the sure amount, no matter who wins the horse race, to the playing objective lottery, no matter who wins the horse race, but he prefers the lottery, no matter what the stock price is tomorrow, to the getting the sure amount, no matter what the stock price is tomorrow.

One proceeds in a similar way for states and probabilities; once again, there are two axioms.

Axiom A8 (*State Richness*) For all S, Ξ_S is non-empty.

Axiom A9 (*Doxastic Consistency*) For all S, for any $\sigma_1, \sigma_2 \in \Xi_S$, for any $a, b \in X_{\sigma_1}$ and $c, d \in X_{\sigma_2}$ with $\mathbf{a} \prec_{\sigma_1} \mathbf{b}$ and $\mathbf{c} \prec_{\sigma_2} \mathbf{d}$, let $\tau : L(\{a, b\})^S \to L(\{c, d\})^S$ be the bijection between the set of acts in σ_1 with values in the set of lotteries on $\{a, b\}$ and the set of acts in σ_2 taking values in the set of lotteries on $\{c, d\}$, defined by $\tau(f)(s, c) = f(s, a)$ and $\tau(f)(s, d) = f(s, b)$ for all $s \in S$.[4] Then, for all $f, g \in L(\{a, b\})^S$, $f \preceq_{\sigma_1} g$ iff $\tau(f) \preceq_{\sigma_2} \tau(g)$.

State Richness is the equivalent for states of Consequence Richness (A6): it ensures that $\|\Xi_S\|$ has at least one element. As for A6, it is a largely structural assumption.

[4] $L(X)$ is the set of lotteries on the set X; $L(X)^S$ is the set of functions from S into $L(X)$—that is, the set of acts taking values in these lotteries. Recall that the acts are considered as functions from pairs of states and outcomes to the real numbers.

Doxastic Consistency is the equivalent of Conative Consistency (A7): it is true if and only if $\|\Xi_S\|$ has at most one element (Lemma 2). As for Conative Consistency, it does have behavioural content: more or less, it demands that, to compare probabilities of events (sets of states) using acts which differ in their consequences depending on whether the event holds, it does not matter if one uses consequences in C_{σ_1} or C_{σ_2}—the answer will be the same. Consider two situations: one where the agent is betting on a horse race (states: results of the race; consequences: monetary prizes) and another where he is advising his friend how to bet on the race (states: results of the race; consequences: monetary prizes for the friend). Suppose furthermore that the agent prefers a sure \$100 to a sure \$0 in the first situation, and that he prefers that his friend wins \$100 for sure to his friend winning nothing for sure in the second situation. Then, according to Doxastic Consistency, he should prefer a bet of \$100 on horse α over a bet on horse β if and only if he prefers advising his friend to take the bet on α to advising his friend to take the bet on β. Indeed, an agent who did not have such preferences would seem rather inconsistent.

Given the preceding remarks, there is a resemblance between Doxastic Consistency (A9) and Savage's P4, which states that the preferences over bets on events are independent of the consequences of the bets. But P4 holds *within* simple situations, whereas Doxastic Consistency holds *between* different simple situations. So, if Doxastic Consistency holds between situations having different sets of consequences, one might expect P4 to hold in the situation having the union of these sets as consequences. This is essentially what is expressed by the following axiom, which implies Doxastic Consistency, in the presence of A1 (Proposition 1 in the Appendix).

Axiom A10 For any S, if $\sigma_1, \sigma_2 \in \Xi_S$, then there is a situation $\sigma_{12} \in \Xi_S$ with $C_{\sigma_{12}} = C_{\sigma_1} \sqcup C_{\sigma_2}$, where $C_1 \sqcup C_2$ is the set of lotteries on $X_1 \cup X_2$.

This axiom is not required for the result stated below; it shall, however, prove relevant in Sect. 4.

Null events. One final axiom is required to deal with the possibility of null events. In the state-independent representation, null events are generally those which are allocated probability 0; it follows that the order is indifferent between any pair of acts which differ only on such events. For the probabilities elicited in simple situations to be valid in non-simple situations, the preference order in these non-simple situations must show the same sort of indifference. Thus the following axiom is posed as a consistency constraint. Recall the classic definition of a null event (Savage 1954; Karni et al. 1983): an event A in a situation σ is *null* iff, for any pair of acts $f, g \in \mathcal{A}_\sigma$ such that $f(s) = g(s)$ for $s \notin A$, $f \sim_\sigma g$. (Null states are those whose singletons are null events.) The axiom is as follows.

Axiom A11 (*Null Consistency*) For any situation σ and any event $A \subseteq S_\sigma$, if A is null in every $\sigma' \in \Xi_{S_\sigma}$, then A is null in σ.

2.2 Theorem

The postulates and axioms give the following representation theorem.

Theorem 2 *Assume Postulates 1 and 2. Moreover, let Consequence Richness, Conative Consistency, State Richness, Doxastic Consistency and Null Consistency hold. Then, for any situation $\sigma \in \mathcal{S}$, there exists a probability distribution p on S_σ, a utility function u on X_σ, and a function $\gamma : S_\sigma \times X_\sigma \rightarrow \Re$ such that,*

- *for all $f, g \in \mathcal{A}_\sigma$*

$$f \preceq_\sigma g \text{ iff } \sum_{\substack{s\in S_\sigma \\ x\in X_\sigma}} p(s)\cdot\gamma(s,x)\cdot u(x)\cdot f(s,x) \leq \sum_{\substack{s\in S_\sigma \\ x\in X_\sigma}} p(s)\cdot\gamma(s,x)\cdot u(x)\cdot g(s,x) \tag{4}$$

- *for each $\sigma_1 \in \Sigma_{C_\sigma}$, there exists a probability p_1 such that, for all $f_1, g_1 \in \mathcal{A}_{\sigma_1}$*

$$f_1 \preceq_{\sigma_1} g_1 \text{ iff } \sum_{\substack{s\in S_{\sigma_1} \\ x\in X_{\sigma_1}}} p_1(s)\cdot u(x)\cdot f_1(s,x) \leq \sum_{\substack{s\in S_{\sigma_1} \\ x\in X_{\sigma_1}}} p_1(s)\cdot u(x)\cdot g_1(s,x) \tag{5}$$

- *for each $\sigma_2 \in \Xi_{S_\sigma}$, there exists a state-independent utility u_2 such that, for all $f_2, g_2 \in \mathcal{A}_{\sigma_2}$*

$$f_2 \preceq_{\sigma_2} g_2 \text{ iff } \sum_{\substack{s\in S_{\sigma_2} \\ x\in X_{\sigma_2}}} p(s)\cdot u_2(x)\cdot f_2(s,x) \leq \sum_{\substack{s\in S_{\sigma_2} \\ x\in X_{\sigma_2}}} p(s)\cdot u_2(x)\cdot g_2(s,x) \tag{6}$$

Furthermore, if p', u', γ' is another representation satisfying (4–6), then there exist positive real numbers a and c, and real numbers b and d_s for each $s \in S_\sigma$, such that $p'(s) = p(s)$, $u'(x) = a\cdot u(x) + b$ and $\gamma'(s,x) = c\cdot\gamma(s,x) - \frac{b\cdot c}{u'(x)}\cdot\gamma(s,x) + \frac{d_s}{u'(x)}$, for all $s \in S_\sigma$, $x \in X_\sigma$, and $d_s = 0$ if s is null.

Remark 2 The probability p and utility u have the same uniqueness properties as in typical state-independent utility theorems (for example, Theorem 1); to this extent, they can be said to be elicited in Theorem 2. γ has new degrees of freedom which arise largely from the fact that, whereas state-independent utilities are unique up to positive affine transformation, state-dependent utilities are unique up to simple positive affine transformation (see for example, Karni et al. 1983, who call these transformations "cardinal unit comparable transformations"). Indeed, the representation (4) naturally yields a state-dependent utility function as the product of γ and u (see Sect. 3.1), and it is easy to see that this function is unique up to simple positive affine transformation, just as for the state-dependent utilities elicited by Karni et al. (1983). That is, for another triple p', γ', u' satisfying the conditions of Theorem 2, there is a positive real number a' and real numbers b'_s for each state $s \in S_\sigma$ such that $\gamma'(s,x)\cdot u'(x) = a'\cdot\gamma(s,x)\cdot u(x) + b'_s$.

Remark 3 The theorem does not assume A1, and so allows differences between the preferences on acts taking the same set of states to the same consequences, but which are elements of different situations (notably, situations having the same sets of states and different but overlapping sets of consequences). In such cases, the representations of the preferences in the different situations may differ in γ and u.

On the other hand, if A1 holds, the preferences agree (on common acts) between such situations; it follows that u takes the same values in situations with at least two outcomes in common and γ takes the same values in all situations which share the same set of states and which have at least two outcomes in common (up to the transformations given in Theorem 2).

3 Discussion

3.1 The situation-dependent factor

The factor γ is a function of states and consequences. Since different situations necessarily have different states and consequences, they will necessarily involve different γ's. Thus γ is properly thought of as a situation-dependent factor: particular to the decision situation under consideration, rather than applicable in different decision situations, as, say, utilities and probabilities are. Put succinctly, γ captures the contextual factors of the decision situation.

It should be noted that a single decision problem can sometimes be rewritten in several ways. The discussion in Sect. 1 provides an example: the problem the husband faces can be formulated either as a choice between acts yielding one of two consequences (Table 1) or as a choice among acts yielding one of four consequences (Table 2). These formulations of the decision problem are two different situations, in the sense of Definition 1, with different sets of consequences; therefore, the representations will involve different factors γ.

Consider, for example, under the 4-consequence formulation (Table 2), 50–50 (subjective) probabilities, a utility of 150 for \$100 with the wife alive, of 50 for \$100 with the wife dead, of 75 for \$0 with the wife alive and of 25 for \$0 with the wife dead, with γ taking value 1 when the state and consequence are compatible, and 0 when they are not. This is a typical representation of the husband's preferences, which explains the preference for the bet on success (bet A in Table 2) over failure (bet B) by a larger utility of \$100 with the wife alive, and the preference of the bet on success (bet A) over a bet on resuscitation (bet E) by a low value of γ which represents the fact that the latter bet is unbelievable. However, it is equally possible to represent the husband's preferences when the 2-consequence formulation is used (Table 1), just this will require a different utility function (because the consequences are different) and a different situation-dependent factor γ. For example, 50–50 probabilities, a utility of 100 for \$100 and 50 for \$0, and a γ with $\gamma(\text{success}, \$100) = 1.5$, $\gamma(\text{failure}, \$100) = 0.5$, $\gamma(\text{success}, \$0) = 1.5$, $\gamma(\text{failure}, \$0) = 0.5$ represents "the same" preferences as described above in the 4-consequence formulation.[5] Notably, they account for the

[5] Naturally, any claim that the preferences are the "same" has to be qualified by the fact that, because of the different sets of consequences, the objects of preference—the acts—differ between the two formulations. For further discussion of the relationships between the situation-dependent factors and utilities in such different but related situations, see Sect. 3.2.

preference for the bet on success, whilst retaining the 50–50 beliefs about the result of the operation. The use of the 2-consequence formulation or the 4-consequence formulation is largely a decision for the theorist; as this example indicates, the proposed representation (2) applies whatever choice is made.

As for all theorems treating decision under uncertainty, Theorem 2 elicits the probabilities and utilities; moreover, it also elicits the situation-dependent factor γ. The latter is thus to be thought of as "subjective" rather than "objectively" given. It represents aspects of the way that the agent thinks of the decision problem and the contextual factors involved, rather than the contextual factors "really" involved. For example, if the decision maker did not realise that resuscitation was impossible, the γ described above (for the 4-consequence formulation of the example) would not feature in an accurate description of his preferences (for he would be indifferent between bet A and bet E in Table 2), even if it is an objectively accurate representation of the relationship between states and consequences. The topic of objective situation-dependent factors is beyond the scope of this article.

A situation-dependent factor of this sort, and a representation of the form (2) proves useful in the analysis of the phenomenon of adaptive preferences. In particular, it allows a distinction between the agent's "absolute" utility u—which is independent of the situation in which he finds himself—and his "situation-relative" utility or "utility in practice". The latter is the state-dependent utility function obtained as the product of the factor γ and the absolute utility u. See Hill (2009) for an extended discussion of this distinction and the importance for the problem of adaptive preferences.[6]

The role of γ as characterising the interdependence between states and consequences, or the difference between the "absolute" and "situation-relative" utilities may prove useful in other areas of economic analysis. It is known that the agent's attitude to risk may differ depending on the state realised: his attitude to risk is different in the case of illness as opposed to health, or in the case of life as opposed to death (Karni 1983; Drèze 1987, Chap. 8). Under state-dependent utility representations of his preferences, this difference is built into the utility function: the function has different properties (form, curvature and so on) on each state. Under the proposed representation, by contrast, it is separated from his general attitude to risk. The "absolute" utility represents the agent's general attitude to risk (in the standard way), whereas the situation-dependent factor γ captures the aspects which are specific to the interdependence between states and consequences involved in the particular decision situation.

Consider the example of health insurance, to which state-dependent utilities are often applied: the states are states of health and the consequences are monetary outcomes. γ represents the differences in utility and risk attitude with respect to change in state of health, whereas the utility function measures his utility for money (and the risk attitude associated) and applies even beyond decisions regarding health insurance (to investment decisions, career decisions, and so on). For an agent who is more risk averse in case of ill health than good health, this fact would be represented by the fact that γ is more concave when the state is ill health than when it is good health.

[6] Another possible interpretation is to think of u as the "intrinsic" utility and the product with the situation-dependent factor as the "experienced" utility. The author wishes to thank a referee for pointing out this interpretation, and the relation to the distinction drawn by Kahneman et al. (1997).

Future research would investigate whether tools for the analysis of risk attitude in cases where there is interdependence between states and consequences can be developed employing the situation-dependent factor proposed here. One would expect to obtain techniques for doing comparative statics on the situation-dependent factor: one could thus compare the different attitudes to health insurance of individuals with the same beliefs about their future health and the same general utility for money.[7]

It should be clear from the discussion so far that, beyond the observation that γ is a contextual or situation-dependent factor, there are several ways to give it a more concrete interpretation. A full discussion of the range of interpretations is beyond the scope of this article; we shall just present one further interpretation which is perhaps applicable in some if not all cases, and which has already been suggested in the discussion of adaptive preferences (Hill 2009).

Consider once again the 4-consequence formulation of Aumann's example. As noted above, the situation-dependent factor which corresponds to the described preferences typically takes lower values on pairs of states and consequences which are inconsistent (failed operation and \$100 with wife alive, for example) than on pairs which are consistent. A natural interpretation of this is in terms of reliability. The act purporting to take the state where the operation fails to the consequence where he wins \$100 and his wife is healthy is not possible. If offered such an act (or such a bet), the agent would not trust the bookey. The act is *unreliable* (or, more precisely, the part of the act purporting to send this state to this consequence in unreliable). The expected utility should take this into account, and it is precisely the situation-dependent factor γ that does this. An interpretation of the situation-dependent factor which seems plausible in at least some cases is as a measure of the reliability of acts purporting to take particular states to particular consequences. It should be noted finally that under this interpretation it is most natural to expect γ to have maximum and minimum values (utterly reliable and utterly unreliable). It is not too difficult to show, by scaling the utilities so that they are positive, that Theorem 2 can always yield a representation with $\gamma \in [0, 1]$. This form of the theorem is not given here, for such a range of γ may not be natural for other interpretations. Indeed, it is worth re-emphasising that this interpretation is one of many: to the extent that γ represents the effect of contextual or situation-dependent aspects on the agent's preferences, there may be as many interpretations as there are sorts of contextual or situation-dependent factors.

3.2 The elicitation technique

The elicitation technique is based on a simple idea: instead of trying to elicit attitudes in cases which are complicated (for example, where there is interdependence between states and consequences), use cases which are straightforward (where there is no such interdependence). Such a technique invokes situations other than the one in question; it thus involves not a single situation, but a set of situations, $\mathcal{S}$.

[7] It may be perhaps be possible to use differences in the situation-dependent factor to account for the observation that agents' risk aversion depends on the circumstance or domain (MacCrimmon and Wehrung 1990; Weber et al. 2002); if so, the notion might also prove useful in studies of such differences in risk aversion.

The situations in S are the decision problems the agent could have been facing at the particular moment in question (given by the sets of states and consequences) and the preferences he would have had were he faced with these problems (given by the preference relation). One of the situations is actual: that is the situation which he is really facing at a given moment. The others are hypothetical: he could have faced them, and would have had the preferences specified if he had, but he is not actually facing them (see below for further discussion of the notion of hypothetical choice involved here). Some of these situations are simple, insofar as they do not involve interdependence between states and consequences: they are used to elicit the agent's probabilities and utilities. The method of elicitation relies on a basic assumption: that the agent's attitudes (probabilities and utilities) have a certain degree of stability or constancy. For it to make sense to elicit the agent's probabilities and utilities in the actual situation by looking in other situations, it must be assumed that the same probabilities and utilities are involved in different situations.

This assumption is not only natural, but a necessity for the task of elicitation or measurement to make any sense. First of all, one would expect rational agents to have probabilities and utilities which are stable enough not to be excessively dependent on the precise decision problem with which they are faced; and many real agents' attitudes do have this property. Secondly, if one does not suppose such constancy, one cannot use the measurements of an agent's probabilities and utilities in any situation other than that in which it was measured. For example, probabilities and utilities measured in an experiment could not be used in models of agents' behaviour in the market. There is a risk that this problem does indeed affect some experimental work (Harrison et al. 2007); nevertheless, for there to be any interest in measuring attitudes, an assumption such as constancy seems to be necessary.

Furthermore, many elicitation techniques do in fact seem to make implicit use of an assumption like the constancy assumption. Often different sorts of questions are used to elicit different components of the representation (probabilities, utilities, decision weights and so on); these different questions correspond, in the terminology of this section, to different situations (see Sect. 4 for a construal of these situations as subsituations of a single fixed situation). One must thus assume that the probability function elicited in one situation is the same as the one involved in the other situation. Naturally, were an experimental setup to be proposed based on Theorem 2, it would use exactly this method. First the agent's probabilities would be elicited using standard techniques but in situations which do not necessarily involve the consequences in the decision problem of interest. Then the same would be done for utilities. Finally, by eliciting his evaluation function in the situation of interest, the situation-dependent factor would be derived.

Beyond being essential to the elicitation technique employed in the theorem, the constancy assumption does have some behavioural content. Indeed, the Conative and Doxastic Consistency axioms (A7 and A9)—which constitute, along with the requirement of basic consistency of preferences (Postulate 1) and the requirement that the agent's preferences have an additive representation (Postulate 2), the main behavioural content of the theorem—follow from the constancy assumption. If the probabilities and utilities are the same in all situations, then in those situations where they can be elicited, one obtains the same answer. The axioms imply that this is the case: that the utility and probability elicited in any simple situations are always the same, for a given set of consequences and states.

The elicitation technique rests upon another assumption, regarding the richness of the set of situations $\mathcal{S}$. To elicit the agent's probabilities in a given non-simple situation, the technique looks in other situations, which have the same states, but different and independent consequences, so that elicitation of probabilities is simple. Hence it needs to be assumed that such situations exist. Similarly, to elicit utilities, it needs to be assumed that simple situations with the same consequences but perhaps different states exist. This is basically the content of the Consequence and State Richness axioms (A6 and A8). As noted in Sect. 2.1, these can be thought of as structural or technical; nevertheless, they are not unreasonable. To take the example discussed in Sect. 1 (in the 2-consequence version), consider the decision situation where the states of the world are results of a horse race and the consequences are \$100 or \$0. In this situation the consequences are most likely independent of the states (assuming the husband does not know any of the horse-owners or jockeys), so that the state-independence axioms apply and the utility of the money can be elicited using standard techniques. This situation has the same consequences but different states from the situation where the husband bets on the outcome of the operation, so that the former situation can be used to elicit the utilities in the latter one. Hence Consequence Richness (A6) is satisfied in this case. Similar examples can be given for states and State Richness (A8). All that is required are situations with the same states (success or failure of the operation) but consequences which are independent, so that the state-independent axioms and results apply. Examples include situations were the consequences are gains and losses for a friend, and the husband has a choice about the advice he gives (see also Karni 1996, p. 256), or situations where the consequences of the husband's bet will be "ethical" issues, such as differing numbers of deaths in Myanmar.[8] In cases such as these, the utility of the consequence does not depend on the outcome of the operation (the state of the world), and so the state-independence axioms apply; they show that, for the example in Sect. 1, State Richness is satisfied.

The same points hold for the 4-consequence formulation of the decision problem. The utility function elicited in the 4-consequence formulation will be different from that elicited in the 2-consequence formulation because the sets of consequences on which these functions are defined differ. Equally, the situations used to elicit them will differ: in one case, bets on a horse race yielding monetary prizes will be involved, in the other case, bets yielding combinations of monetary prizes with a deceased or healthy wife will be used. Naturally, one might sometimes expect there to be relations between the utility functions in these different situations, just as one might expect there to be relations between the situation-dependent factors, which also differ between the situations because of their different consequences (Sect. 3.1). Candidate relations would ideally come in the form of equations which connect the utilities, probabilities and state-dependent factor in one situation to those in the other situation, and which correspond to behavioural conditions on the agent's preferences—notably relations between his preferences in the two situations. A simple example of such a relation

[8] Thanks to Itzhak Gilboa for suggesting this example.

has already been mentioned. Axiom A1 presents a condition on the preferences of the agent in situations with different but overlapping sets of consequences; as noted in Remark 3, it ensures that the utilities on common consequences, and the values of the situation-dependent factors on common state–consequence pairs are the same in the different situations. A direction for future development would consist in finding and assessing such conditions on preferences in different situations, and the corresponding consequences for the relation between utilities, probabilities and state-dependent factors in these situations.

Such a project requires a clear idea of the interpretation of the situations, and of the differences between them. Consider once again Aumann's example and the elicitation of utilities by reference to situations where the agent is betting on a horse race. In the situation with 4 consequences, the outcome of the wife's operation is explicit in the consequences, whereas in the situation with 2 consequences, it is not—if it plays a role in the evaluation of the consequences, this role is at most implicit. The explicitness or implicitness is an important aspect of the relationship between the two situations, and may be interpreted in at least two ways.

Firstly, it may be understood as a difference in the way the problem is posed to the agent: he has his wife's health in mind at all times, but has to make choices where the health of his wife is not explicitly at issue. In such cases, one might expect a principled relationship between his utilities in the 2-consequence situation and those in the 4-consequence situation. For example, he might be expected to consider each consequence in the 2-consequence situation (say, \$100) as a bet over pairs of consequences in the 4-consequence situation (\$100 with the wife alive and \$100 with the wife dead), so that his utility for the 2-consequence outcome is the expected utility of \$100 over his subjective probabilities about the outcome of the operation (Savage 1954, §5.5).

Alternatively, the implicitness and explicitness could be interpreted as properties of the way the agent himself sees the problem. The 4-consequence formulation corresponds to the case where he is aware of the effect of the operation's success or failure on his utility for money; the 2-consequence formulation represents the case where he is not aware of this effect. In this case, it is possible that his utilities, probabilities and situation-dependent factor are related in the sort of principled way just described: although he is unaware, he evaluates the options in the same way as if he were aware. On the other hand, it more likely under this interpretation that the relationships between the utilities, probabilities and state-dependent factors in the 2- and 4-consequence situations are more complicated.[9]

A further point should be made about this last interpretation. There are many things about which the agent is not necessarily certain and which could affect his utility for money: the health of himself and his family, his plans and ambitions, his current and future state of wealth, the relation to the state of the economy and so on. It cannot be asked of him that he be aware of all these issues when he chooses acts yielding monetary consequences; at most, they play an implicit role. So one cannot reject the 2-consequence formulation of the husband's decision, the situation required for

[9] Awareness has recently been recognised as important in decision and game theory (Dekel et al. 2001; Heifetz et al. 2006). So far, no accepted theory exists of the relationship between the agent's attitudes when he is aware and those when he is not. See Hill (2007) for an approach.

elicitation of the utility in this formulation, and the utility elicited, on the basis that a relevant factor—the outcome of the operation—was "left out"; for there are more relevant factors than could ever feasibly be explicitly "brought in". The elicitation of utility described yields a utility function in which these factors are, so to speak, implicitly present.

Let us remark finally that this is not the first time that situations other than the one of immediate interest have been used in representation theorems. Examples include Karni (2007, 2008) and Karni et al. (1983); Karni (1985). In the former cases, reference is made to the situation the agent finds himself in after having made a Bayesian update; in the latter cases, the situations involved are hypothetical. To this extent, the latter cases resemble the proposal here, where, as stated above, the natural interpretation of the other situations involved is as hypothetical. Indeed, some of the arguments presented by Karni and Mongin (2000) in favour of the use of hypothetical data in elicitation carry over to the current result; see Hill (2009) for other arguments which are particular to sort of data involved here.

However, beyond this surface resemblance, there is an important difference between the approaches just cited and the one taken here. There, the other situations which are used involve the *same decision problem* (states and consequences) but *different* (hypothetical or updated) *beliefs*, whereas in this article, we have (to the extent that the states and consequences are shared) the *same beliefs and utilities* but *different decision problems*. This difference is crucial, for it means that the articles mentioned above involve elicitation of attitudes with reference to preferences which are incompatible with the preferences the agent actually has, just as for Savage's strategy in the example discussed in Sect. 1. Here, by contrast, there is no incompatibility between the preferences in the situations involved in the elicitation: as the previous discussion of the constancy assumption makes clear, this point is at the core of the present approach. The sense in which there is something hypothetical at issue is thus milder here: whereas previous approaches have considered what would happen with a different decision maker (with hypothetical beliefs) faced with the same decision problem, we are considering what would happen if the same decision maker were faced with a different decision problem. As suggested above, the second sense of hypothetical is easier to overcome in practice: it suffices to get the decision maker to choose in a decision problem in which it is easier to elicit his attitudes.

A further consequence of this difference is that, since there is no incompatibility among his preferences, they can all in principle be considered as "fragments" of a (consistent) set of general preferences. From the point of view of the general preferences, the decision situations we have defined are just fragments of a large decision problem. Since this decision problem is actual, the distinction between the hypothetical and actual decision situations disappears—they are just subsituations of one, large, actual situation. One thus obtains a version of Theorem 2 which makes no reference at all to hypothetical situations: a relief for those who are not convinced by the preceding arguments in favour of the mild notion of hypothetical situation employed here. In the final section, we detail this setup and prove a simple version of this result.

4 Non-situational version

4.1 Preliminaries and axioms

The use of situations other than the one in which the decision maker is choosing may be shunned. Fortunately, they are not strictly necessary for the result given above: instead of thinking of the situations used in the result as different situations, one could consider the sets of states, consequences and acts of each situation to be appropriate *subsets* of the elements (events, consequences, acts) of some large, fixed situation in which the decision maker is making his choice. Situations would thus be considered to be like the "microcosms" of Savage (1954, §5.5), so the relation of their "small world" states and consequences to any fixed "grand world" states and consequences can be understood. Whereas the bets on the success and failure of the operation and the bets on the horse race were considered as acts belonging to different situations in the previous sections, here they are considered as being part of one grand world situation, whose states specify both the outcome of the operation and the results of the race. In this way, it is possible to formulate an equivalent to Theorem 2 in terms of one "grand world" situation. In this section, a simple theorem of this sort is stated. It gives a flavour of how the technique introduced in the previous sections continues to apply in a single, grand world situation.

Let S designate the set of states of the grand world situation, X the set of outcomes, C the set of lotteries on X, and $\preceq$ the preference order on the set of acts $\mathcal{A}$ (functions from S to C). S and X are assumed to be finite.

To obtain a representation of the preferences and elicit unique probabilities and utilities, the idea presented and discussed in Sects. 2 and 3 will be employed: elicit the probabilities and utilities in related, simple, situations. By contrast to the previous sections, the situations used here will be "small world" situations with respect to the fixed grand world situation, or if you prefer *subsituations* of the grand world situation. That is, they will be situations of the form $(S_\sigma, C_\sigma, \preceq_\sigma)$ where S_σ is a set of events which partition S, C_σ is the set of lotteries on X_σ for some $X_\sigma \subseteq X$ and $\preceq_\sigma$ is the restriction of $\preceq$ to acts which are constant on the elements of S_σ and which take values in C_σ. A small world state is a grand world event; the same notation will be used when it is treated as the former as when it is thought of as the latter. An act f in the (small world) situation $(S_\sigma, C_\sigma, \preceq_\sigma)$ is considered as a (grand world) act $\hat{f}$, where, for each $s \in S$, $\hat{f}(s, x) = f(s_\sigma, x)$, for s_σ the element of S_σ such that $s \in s_\sigma$.[10]

Axioms Beyond the axioms introduced in Sect. 2, the following axioms will be required.

Axiom A12 There exists a set of sets of events on S, $\{S_i\}_{i \in I}$, such that,

(i) for each $s \in S$, there is a set $\{s_i^s\}$, with one element s_i^s from each S_i, such that $\bigcap_{i \in I} s_i^s = \{s\}$;

[10] This interpretation differs from Savage's interpretation of microcosms (Savage 1954, §5.5); thus the difference in the interpretation of small consequences, which are not grand world acts, as in Savage.

(ii) for each S_i, there exists $X' \subseteq X$ such that $(S_i, C', \preceq')$ is simple (where, as above, C' is the set of lotteries on X');
(iii) for each $s \in S$, the set $\{s_i^s\}$ is minimal: there is no proper subset $A \subset \{s_i^s\}$ such that $\bigcap_{s' \in A} s' = \{s\}$.

Axiom A13 For each set of events S' which partition S, if there exists $X' \subseteq X$ such that $(S', C', \preceq')$ is simple, then there is a unique such X' which is maximal under inclusion—that is, such that there is no $X'' \supset X'$ with $(S_i, C'', \preceq'')$ simple.

For each S_i from A12, the set of outcomes described by A13 is denoted by X_i, the corresponding set of consequences by C_i, and the preference order by $\preceq_i$.

Axiom A14 For any simple small world situation $(S', C', \preceq')$, if $A \subseteq S'$ is a null event in $(S', C', \preceq')$, then $\bigcup_{s_j \in A} s_j$ is a null event in $(S, C, \preceq)$.

Discussion Just as for the result in Sect. 2, two sorts of conditions are required. On the one hand, there are standard conditions guaranteeing that the agent is a consistent expected utility maximiser. As in the previous sections, the axiom for state-independence (Monotonicity) is not assumed to hold in the grand world situation, but the von Neumann-Morgenstern axioms (Weak Order, Independence and Continuity) are (see Theorem 3). This assumption is logically stronger than the equivalent in the situational setup (Postulate 2): the former implies the latter, because it implies that the axioms hold in individual situations; but the former is not implied by the latter because it involves inter-situational comparisons. Note furthermore, that, since the preference orders in the small world situations are restrictions of a grand world preference order, they agree on the acts they have in common, so Postulate 1 and the stronger A1 hold automatically.

The other conditions involved in Theorem 2 guarantee that the set of situations is suitable for elicitation of probabilities and utilities. In the theorem presented below, there are no assumptions or axioms regarding consequences and utilities. This is because the preference relations in the small world situations are derived from the preference relation in the grand world situation, so that they agree on constant acts, and the restriction to constant acts yields, by the von Neumann-Morgenstern theorem, a utility function which applies in all (small world and grand world) situations. So the use of small world situations is not required to elicit utilities; by contrast, they are required for probabilities, and here supplementary conditions have to be imposed.

Axiom A12 is the equivalent of State Richness (A8) in the grand world situation framework, and, as for State Richness, it is a largely structural or technical axiom. Recall that State Richness guarantees the existence of simple situations with the same set of states as the situation of interest; the probabilities elicited in the former situations is used to deduce the probability in the latter. However, in the grand world version, it is not elicitations in situations with the same set of states which are used to obtain the probabilities in the grand world situation, but rather elicitations in subsituations; thus one needs the existence of a set of simple small world situations which is rich enough that, by eliciting the probabilities in the small world situations, one can deduce the grand world probabilities. This is the role of A12. Clause (ii) implies that probabilities can be elicited on each of the sets S_i, which are the sets of states of simple small world situations. Clause (i) implies that these sets of small world states "cover" the set of grand world states. Formally, this clause implies that the Boolean algebra of events of the grand world situation is the smallest Boolean algebra containing the Boolean algebras generated by the $\{S_i\}_{i \in I}$. In other words, every grand world event can be obtained by taking appropriate intersections and unions of events in the small world situations. If this were not the case, it would never be possible to ascertain the probabilities of some of grand world events using just the probabilities elicited in these small world situations.

Clause (iii) translates a simplifying assumption which is made in this section: namely, that one can work with small world situations whose states are stochastically independent. This assumption is at work in the proof of the theorem, although, as is generally the case for central assumptions in representation theorems, it cannot be stated explicitly: after all, the most natural way to express stochastic independence is in terms of the probability function, and such a function is exactly what is to be elicited. Nevertheless, this assumption has some formal consequences, and clause (iii) of A12 is one of them. If the states are stochastically independent, then they are logically independent (for each state $s_j \in S_i$, $s'_j \in S_{i'}$, the intersection $s_j \cap s'_j$ is non-empty), and clause (iii) guarantees this. Technically, it implies that the set of states of the grand world situation is the Cartesian product of the sets S_i.

The supplementary assumption that there exists not only a set of simple small world situations whose states cover the grand world state space but that there exists such a set of situations whose states are stochastically independent is made here largely for convenience. First of all, versions of the result stated below continue to hold *mutatis mutandis* when this assumption is weakened appropriately. Furthermore, given that the aim here is merely to indicate that the results in the previous sections can be reformulated in the context of a single situation, rather than enter into complicated technical details, the simplicity permitted by the assumption justifies its use. Finally, the assumption is quite plausible in many contexts. Naturally, one can always invent a grand world situation (set of states, set of consequences and preference on acts) where the assumption of the existence of simple small world subsituations with stochastically independent states will not hold (and in particular, where A12 is violated). However, if, as is more often the case, one knows some of the acts, consequences and events involved in a decision problem and one wishes to elicit attitudes, then it is usually possible to find a model of the problem as a single grand world situation which has the required property. Note, for example, that most of the simple situations considered in the previous sections have stochastically independent states: the state of health of the wife is stochastically independent of the result of the horse race, for example. The grand world situation generated by these simple situations has a set of states which is covered by stochastically independent sets of events.[11] Indeed, what was done in previous sections can be thought of as operating in this, single, grand world situation. For these cases, which concern us here insofar as they show that the interpretation in terms of hypothetical situations is unnecessary, the assumption of stochastic independence is not implausible.

[11] For details on how to generate a grand world state space from small world ones, see Hill (2008).

Axiom A13 does the job of Doxastic Consistency (A9); in fact, given the framework, it is equivalent to the stronger A10 (Proposition 2). The comments made in Sect. 2 regarding Doxastic Consistency largely hold for A13. In particular, it has behavioural content, demanding that the agent have a certain consistency: the probabilities which can be deduced from his preferences in simple situations do not depend on the consequences and thus the simple situation used.

Finally, A14 is the equivalent in the current framework of Null Consistency (A11). It just says that, if A as an event in the simple situation $(S', C', \preceq')$ is null, then it is null as an event in $(S, C, \preceq)$.

4.2 Theorem

One can thus obtain an (essentially) unique representation of the preference relation of the form (2) which agrees with the probabilities and utilities elicited in the simple small world situations.

Theorem 3 *Let S be a set of states, X a set of outcomes, C the set of lotteries over X and $\preceq$ an order on $\mathcal{A}$. Suppose that A2–A4 hold. Suppose furthermore that A12 holds with set $\{S_i\}_{i\in I}$ and that A13–A14 hold. Then, there exists a probability distribution p on S, a utility function u on X, and a function $\gamma : S \times X \to \Re$ such that,*

- *for all $f, g \in \mathcal{A}$*

$$f \preceq g \text{ iff } \sum_{s\in S, x\in X} p(s)\cdot\gamma(s,x)\cdot u(x)\cdot f(s,x) \leq \sum_{s\in S, x\in X} p(s)\cdot\gamma(s,x)\cdot u(x)\cdot g(s,x) \tag{7}$$

- *There exist positive α and real β such that, for each $x \in X$*

$$\sum_{s\in S} p(s)\cdot\gamma(s,x)\cdot u(x) = \alpha u(x) + \beta \tag{8}$$

- *For each $i \in I$, there exists positive α_i and real β_i such that, for each $s_j \in S_i$, for each $x \in X_i$,*

$$\sum_{s\in s_j} p(s)\cdot\gamma(s,x)\cdot u(x) = \alpha_i p(s_j)\cdot u(x) + \beta_i p(s_j) \tag{9}$$

Furthermore, if p', u', γ' is another representation satisfying (7–9), then there exist positive real numbers a and c, and real numbers b and d_s for each $s \in S$, such that $p'(s) = p(s)$, $u'(x) = a\cdot u(x) + b$ and $\gamma'(s,x) = c\cdot\gamma(s,x) - \frac{b\cdot c}{u'(x)}\cdot\gamma(s,x) + \frac{d_s}{u'(x)}$, for all $s \in S$, $x \in X$, and $d_s = 0$ if s is null.

Remark 4 Equation (8) corresponds to (5) of Theorem 2 insofar as it states that the preferences on constant acts is represented by u (and that this representation is unique up to positive affine transformation).

Equation (9) corresponds to (6) of Theorem 2 insofar as it reflects the fact that the preference order in situations $(S_i, C_i, \preceq_i)$ is represented by p and u (and that this representation is unique up to positive affine transformation of u).

Appendix

Lemma 1 *A7 holds if and only if, for all C, $\|\Sigma_C\|$ has at most one element (up to positive affine transformation).*

Proof For any C, if there is at most one element of $\|\Sigma_C\|$, then all simple situations accord the same utilities to consequences (up to positive affine transformation), and thus to constant acts, so A7 holds.

On the other hand, if A7 holds, then using the von Neumann-Morgenstern theorem on the preferences on constant acts, one obtains a utility function on constant acts, and thus consequences, which applies in all situations. However, in each situation, the utility provided by Theorem 1 must also represent the preference over constant acts, by (3); by the uniqueness properties, such utilities must agree up to positive affine transformation. □

Lemma 2 *A9 holds if and only if, for all S, $\|\Xi_S\|$ has at most one element.*

Proof Take any S, and suppose that there is at most one element of $\|\Xi_S\|$. Consider situations $\sigma_1, \sigma_2 \in \Xi_S$ and elements $a, b \in X_{\sigma_1}$ and $c, d \in X_{\sigma_2}$ satisfying the conditions of A9. The utilities u_1 and u_2 representing the orders in these situations can be scaled so that $u_1(a) = u_2(c)$ and $u_1(b) = u_2(d)$. But then, for any $f \in L(\{a, b\})^S$, $\sum_{S,X_{\sigma_1}} p(s)u_1(x)f(s,x) = \sum_{S,X_{\sigma_2}} p(s)u_2(x)\tau(f)(s,x)$: so the preference orders behave as in A9.

Suppose now that A9 holds. For any $\sigma_1, \sigma_2 \in \Xi_S$ and any $a, b \in X_{\sigma_1}$ and $c, d \in X_{\sigma_2}$ satisfying the conditions of A9, take the representations p_1, u_1 and p_2, u_2 with $u_1(a) = 0$, $u_1(b) = 1$ and $u_2(c) = 0$, $u_2(d) = 1$. Consider the following acts in σ_1: for each $A \subseteq S$ f_A, with $f_A(s, b) = 1$ for $s \in A$, $f_A(s, a) = 1$ for $s \notin A$ (and $f(s, x) = 0$ elsewhere); for each element y of $L(\{a, b\})$, the constant act g_y taking the value y. Note that $p_1(A) = \sum_{S,X_{\sigma_1}} p_1(s)u_1(x)f_A(s,x) = \inf_{y \in L(\{a,b\}), f_A \preceq_{\sigma_1} g_y} \sum_{X_{\sigma_1}} y(x)u_1(x)$. However, the image of each of these acts under τ is of the same type (respectively, f'_A, $g'_{y'}$), and u_2 takes the same value as u_1 on the images of the elements involved, so, for all $A \subseteq S$, $p_1(A) = p_2(A)$. □

Proposition 1 *In the presence of A1, A10 implies A9.*

Proof Suppose there is more than one element in $\|\Xi_S\|$ for some S. Then there are simple situations $\sigma_1, \sigma_2 \in \Xi_S$, with representations involving p_1 and p_2, where $p_1 \neq p_2$. By A10, the situation σ_{12} with set of states S and set of consequences $C_{\sigma_1} \sqcup C_{\sigma_2}$ is simple, and so has representation with probability p_{12}. By A1, $\preceq_{\sigma_1}$ coincides with

$\preceq_{\sigma_{12}}$ on the acts common to both situations; so by the uniqueness properties of Theorem 1, they must be represented using the same probability on S. The same goes for $\preceq_{\sigma_2}$: so $p_1 = p_2$, contradicting the supposition. By Lemma 2, A9 must hold. □

Proof of Theorem 2 *Existence.* If σ is simple, the existence of the representation (4), with $\gamma(s, x) = 1$ for all s, x, is given immediately by Theorem 1. By A7, A9 and Lemmas 1 and 2, this representation satisfies the other two clauses.

Suppose σ is not simple. By A6–A9 and Lemmas 1 and 2, there is a unique probability function p and a utility function u, unique up to positive affine transformation, satisfying (5) and (6).

By Postulate 2, a version of the von Neumann-Morgenstern theorem applies in σ (Fishburn 1970, p. 176), giving a representation by an evaluation function $U : S_\sigma \times X_\sigma \rightarrow \Re$, unique up to simple positive affine transformation; that is, for all $f, g \in \mathcal{A}_\sigma$,

$$f \preceq_\sigma g \text{ iff} \sum_{S_\sigma, X_\sigma} U(s, x) \cdot f(s, x) \leq \sum_{S_\sigma, X_\sigma} U(s, x) \cdot g(s, x) \tag{10}$$

Pick U such that $U(s, x) = 0$ on the null states (if $\preceq_\sigma$ has any).

Define γ the function on $S_\sigma \times X_\sigma$ such that

$$\gamma(s, x) = \begin{cases} \frac{U(s,x)}{p(s) \cdot u(x)} & \text{if } p(s) \neq 0 \\ 1 & \text{if } p(s) = 0 \end{cases}$$

By construction and by A11, (4)–(6) hold.

Uniqueness. The uniqueness properties of p and u follow from the application of Theorem 1 to the situations in Ξ_S and Σ_C. Moreover, as noted above, U is given up to simple positive affine transformation. So, for another representation p', u', γ', there are a, b, a' and b'_s for each $s \in S_\sigma$, with a and a' positive, such that, for all $s \in S_\sigma, x \in X_\sigma$, $p'(s) = p(s)$, $u'(x) = a \cdot u(x) + b$ and $p(s) \cdot \gamma'(s, x) \cdot u'(x) = a' \cdot p(s) \cdot \gamma(s, x) \cdot u(x) + b'_s$. Substituting for u and solving for γ', one obtains the required result, with $c = \frac{a'}{a}$ and $d_s = \frac{b_s}{p(s)}$ when $p(s) \neq 0$ and $d_s = 0$ when $p(s) = 0$. □

Proposition 2 *A13 holds if and only if A10 does.*

Proof Suppose A10 but not A13. Since X is finite, for each S' partitioning S, there is no infinite ascending chain $X^0 \subset X^1 \subset \ldots$ with $X^j \subseteq X$ for all j and $(S', C^j, \preceq^j)$ simple for all j. So, if A13 is not satisfied, there must be a S' partitioning S and X_1 and X_2 with $X_1 \neq X_2$ such that $(S', C_1, \preceq_1)$ and $(S', C_2, \preceq_2)$ are simple, and X_1 and X_2 are maximal. But, by A10, $X_{12} = X_1 \cup X_2$ also yields a simple situation $(S', C_{12}, \preceq_{12})$, contradicting the maximality of X_1 and X_2.

Suppose A13, and consider any pair of simple small world situations $(S', C'_1, \preceq'_1)$ and $(S', C'_1, \preceq'_1)$, for sets of outcomes X'_1 and X'_2. Since $X'_1, X'_2 \subseteq X'$, the maximal set mentioned in A13, $X'_3 = X'_1 \cup X'_2 \subseteq X'$, and since $\preceq'_3$ coincides with $\preceq'$ on common acts, $(S', C'_3, \preceq'_3)$ is simple; so A10 holds. □

Proof of Theorem 3 The proof proceeds in much the same way as that of Theorem 2.

Existence. By A2–A4, a version of the von Neumann-Morgenstern theorem can be applied (Fishburn 1970, p. 176), yielding an evaluation function $U : S \times X \rightarrow \Re$ which is unique up to similar positive affine transformation. If $\preceq$ has null states, then pick U such that $U(s, x) = 0$ on the null states. Let u be the restriction of U to constant acts. That is, $u(x) = \sum_{s \in S} U(s, x)$. Note that u is unique up to positive affine transformation.

By A12–A13 and Theorem 1, for each S_i, there is a unique $(S_i, X_i, \preceq_i)$ with X_i maximal, and a unique p_i and u_i, unique up to positive affine transformation, representing $\preceq_i$. That is, for all $f, g \in \mathcal{A}_i$,

$$\begin{aligned} f \;\preceq_i g \text{ iff } & \sum_{s_j \in S_i, x_j \in X_i} p_i(s_j) \cdot u_i(x_j) \cdot f(s_j, x_j) \\ \leq & \sum_{s_j \in S_i, x_j \in X_i} p_i(s_j) \cdot u_i(x_j) \cdot g(s_j, x_j) \end{aligned} \tag{11}$$

Since both u_i and u represent the constant acts taking values in C_i, they are positive affine transformations of each other. The orders $\preceq_i$ are thus represented (in the sense of equation (11)) by p_i and u.

By A12, for each $s \in S$, there exists a unique set $\{s_i^s\}_{i \in I}$, $s_i \in S_i$, with $s = \bigcap_{i \in I} s_i^s$. Define p on S as follows:

$$p(s) = \prod_{i \in I} p_i(s_i^s) \tag{12}$$

Finally, define γ to be the function on $S \times X$ such that

$$\gamma(s, x) = \begin{cases} \frac{U(s,x)}{p(s) \cdot u(x)} & \text{if } p(s) \neq 0 \\ 1 & \text{if } p(s) = 0 \end{cases}$$

By construction and by A14, p, u and γ satisfy (7). Therefore $\sum_{s \in S} p(s) \cdot \gamma(s, x) \cdot u(x)$ represents the restriction of $\preceq$ to constant acts. By the uniqueness properties of u, as a representation of this order, there is a positive real number α and a real number β such that $\sum_S p(s) \cdot \gamma(s, x) \cdot u(x) = \alpha \cdot u(x) + \beta$ for all $x \in X$. Thus (8) is satisfied.

In a similar way, for each $i \in I$, both p, γ and u, and p_i and u represent $\preceq_i$: the first via (7), the second with (11). But, by construction of p, $p_i(s_i) = \sum_{s \in s_i} p(s) = p(s_i)$. Thus, by the uniqueness of the probability and the uniqueness, up to positive affine transformation, of the utility in the representation of $\preceq_i$, there exists a positive real number α_i and a real number β_i such that, for each $s_j \in S_i, x \in X$: $\sum_{s \in s_j} p(s) \cdot \gamma(s, x) \cdot u(x) = \alpha_i \cdot p_i(s_j) \cdot u(x) + \beta_i p_i(s_j) = \alpha_i \cdot p(s_j) \cdot u(x) + \beta_i \cdot p(s_j)$. So (9) is satisfied.

Uniqueness. The uniqueness properties of p follow from Theorem 1 and A13, and the uniqueness properties of u from the von Neumann-Morgestern theorem applied to constant acts. Moreover, as noted above, U is unique up to simple positive affine transformation. So, for another representation p', u', γ', there are a, b, a' and b'_s for each $s \in S$, with a and a' positive, such that, for all $s \in S, x \in X$, $p'(s) =$

$p(s)$, $u'(x) = a \cdot u(x) + b$ and $p(s) \cdot \gamma'(s, x) \cdot u'(x) = a' \cdot p(s) \cdot \gamma(s, x) \cdot u(x) + b'_s$. Substituting for u and solving for γ' as in the proof of Theorem 2, one obtains $\gamma'(s, x) = \frac{a'}{a}\gamma(s, x) - \frac{b \cdot a'}{a} \cdot \frac{\gamma(s,x)}{u'(x)} + \frac{b'_s}{p(s) \cdot u'(x)}$, and thus the required result, with $c = \frac{a'}{a}$ and $d_s = \frac{b_s}{p(s)}$ when $p(s) \neq 0$ and $d_s = 0$ when $p(s) = 0$. □

Acknowledgements The author would like to thank Philippe Mongin, three anonymous referees and the associate editor for their valuable comments, as well as the audiences of 38th European Mathematical Psychology Group Conference (September 2007) and the Decision and Uncertainty Workshop (May 2008).

References

Anscombe, F. J., & Aumann, R. J. (1963). A definition of subjective probability. *The Annals of Mathematical Statistics, 34*, 199–205.

Arrow, K. J. (1974). Optimal insurance and generalized deductibles. *Scandinavian Actuarial Journal, 1*, 1–42.

Cook, P. J., & Graham, D. A. (1977). The demand for insurance and protection: The case of irreplaceable commodities. *The Quarterly Journal of Economics, 91*, 143–156.

Dekel, E., Lipman, B. L., & Rustichini, A. (2001). Representing preferences with a unique subjective state space. *Econometrica, 69*, 891–934.

Drèze, J. H. (1987). *Essays on economic decisions under uncertainty*. Cambridge: Cambridge University Press.

Drèze, J. H., & Rustichini, A. (2004). State-dependent utility and decision theory. In S. Barberà, P. J. Hammond, & C. Seidl (Eds.), *Handbook of utility theory* (Vol. 2). Dordrecht: Kluwer.

Fishburn, P. C. (1970). *Utility theory for decision making*. New York: Wiley.

Hammond, P. J. (1998). Subjective expected utility. In S. Barberà, P. J. Hammond, & C. Seidl (Eds.), *Handbook of utility theory* (Vol. 1). Dordrecht: Kluwer.

Harrison, G. W., List, J. A., & Towe, C. (2007). Naturally occurring preferences and exogenous laboratory experiments: A case study of risk aversion. *Econometrica, 75*, 433–458.

Heifetz, A., Meier, M., & Schipper, B. (2006). Interactive unawareness. *Journal of Economic Theory, 130*, 78–94.

Hill, B. (2007). Dynamic awareness. Technical Report 881/2007, GREGHEC.

Hill, B. (2008). Towards a "sophisticated" model of belief dynamics. Part I: The general framework. *Studia Logica, 89*(1), 81–109.

Hill, B. (2009). Three analyses of sour grapes. In T. Grüne-Yanoff & S. O. Hansson (Eds.), *Preference change: Approaches from philosophy, economics and psychology*. Berlin: Springer, Theory and Decision Library A.

Kahneman, D., Wakker, P. P., & Sarin, R. (1997). Back to Bentham? Explorations of experienced utility. *Quarterly Journal of Economics, 112*, 375–405.

Karni, E. (1983). Risk aversion for state-dependent utility functions: Measurement and applications. *International Economic Review, 24*, 637–647.

Karni, E. (1985). *Decision making under uncertainty*. Cambridge: Harvard University Press.

Karni, E. (1993a). A definition of subjective probabilities with state-dependent preferences. *Econometrica, 61*, 187–198.

Karni, E. (1993b). Subjective expected utility theory with state dependent preferences. *Journal of Economic Theory, 60*, 428–438.

Karni, E. (1996). Probabilities and beliefs. *Journal of Risk and Uncertainty, 13*, 249–262.

Karni, E. (2007). Foundations of Bayesian theory. *Journal of Economic Theory, 132*, 167–188.

Karni, E. (2008). A theory of Bayesian decision making. Working Paper.

Karni, E., & Mongin, P. (2000). On the determination of subjective probability by choices. *Management Science, 46*, 233–248.

Karni, E., & Schmeidler, D. (1993). On the uniqueness of subjective probabilities. *Economic Theory, 3*, 267–277.

Karni, E., Schmeidler, D., & Vind, K. (1983). On state dependent preferences and subjective probabilities. *Econometrica, 51*, 1021–1032.

MacCrimmon, K. R., & Wehrung, D. A. (1990). Characteristics of risk taking executives. *Management Science, 36*, 422–435.

Savage, L. J. (1954). *The foundations of statistics* (2nd ed., 1971). New York: Dover.

Weber, E. U., Blais, A.-R., & Betz, E. (2002). A domain-specific risk-attitude scale: Measuring risk perceptions and risk behaviours. *Journal of Behavioral Decision Making, 15*, 263–290.