A Note on States and Acts in Theories of Decision under Uncertainty

Burkhard C. Schipper*

February 19, 2016

Abstract

The purpose of this note is to discuss some problems with the notion of states and acts in theories of decision under uncertainty. We argue that any single-person decision problem under uncertainty can be cast in a canonical state space in which states describe the consequence of every action. Casting decision problems in the canonical state space opens up the analysis to the question of how decisions depend on the state space. It also makes transparent that the set of all acts is problematic because it must contain incoherent acts that assign to some state a consequence that is inconsistent with the internal structure of that state. We discuss three approaches for dealing with incoherent acts, each unsatisfactory in that it limits what can be meaningfully revealed about beliefs with choice experiments.

Keywords: Decision theory, State spaces, States of nature, Acts, Subjective expected utility, Ambiguity, Ellsberg paradox.

JEL-Classifications: D80, D81.

1 States in Theories of Decision under Uncertainty

In many theories of decision under uncertainty such as Subjective Expected Utility (Savage 1954, Anscombe and Aumann, 1963), Choquet expected utility (Schmeidler, 1989), Maxmin Expected Utility (Gilboa and Schmeidler, 1989) etc., states of nature are used to list possible resolutions of uncertainty. A decision maker is uncertain which action would yield which consequences. A state is "a description of the world so complete that, if true and known, the consequences of every action would be known" (Arrow, 1971, p.

*Department of Economics, University of California, Davis, Email: bcschipper@ucdavis.edu

45).[1] States should be mutually exclusive, the set of states should be exhaustive, and the occurrence of states should not be affected by the decision makers actions (see Machina, 2003). In the "practice" of doing decision theory, states are often represented by colors of balls in an urn (like in the Ellsberg paradox) or are taken as an abstract set satisfying some mathematical structure convenient for deriving representation results. Yet in the context of decision theory, for an element of a set to qualify as a state, it should describe a possible resolution of uncertainty. Making this precise would require a language with which possible resolutions of uncertainty can be described. The vocabulary of the language should be rich enough to describe possible actions, possible consequences and their possible relationships. If a state could not be described with such a language, then it is not clear what exactly it is supposed to represent in the decision context. Nevertheless decision theory remained largely silent about the internal structure of states.[2] It seems that the internal structure of states is not perceived to be an issue except for some special problems like Newcomb's paradox, which gave rise to an extensive literature on causal decision theory in philosophy (see for instance Gibbard and Harper, 1978).

In order to make the internal structure of states precise, we could introduce a formal language with propositions such as "action a leads to consequence c", propositional connectives, and appropriate axioms and inferences rules, and then characterize states as maximal consistent sets of formulas similar to the construction of the canonical model in modal logic (Chellas, 1980).[3] Since economists are typically not trained in formal logic, it may be more useful for the discussion to describe the internal structure of states in a straightforward way by using a kind of canonical space of states of nature discussed in Gibbard and Harper (more precisely even in their working paper version from 1976) and attributed to Jeffrey. Given a set A of primitive actions and a set C of consequences, the canonical space of states of nature is simply $S \equiv C^A$. I.e., in line with Arrow's quote each state specifies the consequences of every action. For instance, if the set of primitive actions is $A = \{\text{own the firm}, \text{ don't own the firm}\}$ and the set of consequences is $C = \{\$100, \$0\}$, then the state space S is given by the columns in Figure 1.

Although a canonical state space describes exhaustively all possible resolutions of uncertainty for the context of the decision and thus is typically large, it is still a space of "small worlds" in the sense of Savage as it is constructed for the set of actions and set of consequences relevant to the decision maker in the particular decision context.

[1] I first saw this quote in Karni (2015) in his interesting discussion of the notion of states in theories of decision under uncertainty.

[2] The internal structure of states is also taken more seriously in interactive epistemology or epistemic game theory. Yet, states of the world in game theory are different from states of nature in decision theory as former also describe hierarchies of beliefs while this seems to be unnecessary for single-person decision theory.

[3] In the context of decision theory, such a construction wouldn't make necessarily use of belief operators since states in (single-person) decision theory just describe physical uncertainties and not beliefs of the decision maker.

Figure 1: Example with Two Actions and Consequences

		States			
		s_1	s_2	s_3	s_4
Actions	own	100	100	0	0
	don't own	100	0	100	0

Any single-person decision theoretic context using states and acts could be cast more explicitly in a canonical state space. It has been used by Schmeidler and Wakker (1987), Karni and Schmeidler (1991), Gilboa (2009), and Karni and Vierø (2013, 2015), and is advocated for by Karni (2015). The advantage is that it makes the analyst aware of implications of focusing on particular subsets of states in an ad hoc way as often done in decision theory. Moreover, it makes explicit that the set of all acts defined on exhaustive state spaces must contain incoherent acts. These points are illustrated in the following sections.

2 State Spaces Do Matter

To illustrate the point, consider probably the most famous paradox to Subjective Expected Utility, the three-color Ellsberg paradox (Ellsberg, 1961) shown in Figure 2. Typically the discussion focuses on just the three states, represented by the colors of the balls in the urn. There is an urn with 30 red balls and 60 balls that are either blue or green. The decision maker has to choose once between actions A and B (Choice 1) and once between actions C and D (Choice 2). It is then argued that it is plausible that a decision maker chooses A over B but D over C. Action A appears to be more favorable than B because the decision maker knows that there are 30 red balls yielding \$100 when choosing A but only an unknown number of green balls yielding \$100 when choosing B. D appears to be more favorable than C because the decision maker knows that there are 60 green or blue balls yielding \$100 when choosing D but only an unknown number of balls (above 30) yielding \$100 when choosing C. Yet, this "choice reversal" violates Savage's Sure Thing Principle because the only difference between the choices is the payoff in state s_3, which is the same within each choice but differs across choices. This gave rise to a sizable and growing literature on decision making under ambiguity and Knightian uncertainty.

Note that choices between those actions do not reveal whether or not the decision maker represents the problem as in Figure 2. It is not absurd to think that any participants to an experiment may entertain a tiny amount of doubt about what the experimenter tells him/her. To be able to depict all possible resolutions of uncertainty for these

Figure 2: Ellsberg Paradox

		States		
		s_1	s_2	s_3
	Actions	Red	Green	Blue
Choice 1	A	\$100	\$0	\$0
	B	\$0	\$100	\$0
Choice 2	C	\$100	\$0	\$100
	D	\$0	\$100	\$100
	# Balls	30	60	

four actions that some decision maker may entertain, we embed the Ellsberg paradox in the canonical state space depicted in Figure 3. The first three states are simply the states considered previously in Figure 2. E.g., the state associated with the red colored ball describes that action A yields consequence \$100, action B yields consequence \$0, action C yields consequence \$100 and action D consequence \$0.

Figure 3: Ellsberg in the Canonical State Space

		State															
		s_1	s_2	s_3	s_4	s_5	s_6	s_7	s_8	s_9	s_{10}	s_{11}	s_{12}	s_{13}	s_{14}	s_{15}	s_{16}
	Actions	Red	Green	Blue													
Choice 1	A	\$100	\$0	\$0	\$0	\$100	\$0	\$0	\$0	\$100	\$0	\$100	\$100	\$0	\$100	\$100	\$100
	B	\$0	\$100	\$0	\$0	\$0	\$100	\$0	\$0	\$100	\$100	\$0	\$100	\$100	\$0	\$100	\$100
Choice 2	C	\$100	\$0	\$100	\$0	\$0	\$0	\$100	\$0	\$0	\$100	\$0	\$100	\$100	\$100	\$0	\$100
	D	\$0	\$100	\$100	\$0	\$0	\$0	\$0	\$100	\$0	\$0	\$100	\$0	\$100	\$100	\$100	\$100
	Balls #	30	60														
		30	30	30	0	0	0	0	0	0	0	1	0	0	0	0	0

When all possible resolutions of uncertainties are considered in the canonical state spaces, the typical "choice reversal" of the Ellsberg paradox becomes rationalizable by Subjective Expected Utility. E.g., put equal weight on the first three "original" states (let's say 30 balls of each color) and some tiny ε-weight on state 11 (e.g., just one brown ball) for $\varepsilon > 0$ but arbitrarily small. Then the expected values from the four actions A, B, C, and D are about \$34, \$33, \$66, and \$67, respectively. Hence, A yields a higher expected value than B but C yields a lower expected value than D.

To sum up, the example shows that unless the analyst is willing to bet her/his entire wealth on that the decision maker's representation of the decision problems corresponds exactly to the analyst's representation, she/he should be prepared to be fundamentally wrong about whether or not certain choice behavior corresponds to certain "axioms" of a decision theory. Conclusions about choice behavior may depend dramatically on the representation of the decision problem.[4] Any "behavioral implications" are relative to the primitives of the model which themselves are not directly implied by the very behavior but assumed by the analyst. The primitives provide the analyst with the context in which to interpret behavior.

Casting the analysis explicitly in the canonical state space opens up the primitives for analysis. Such an analysis could be carried out systematically on the collection of all canonical state spaces defined as follows: For a given set actions A and set of consequences C, define a collection of canonical state spaces $\mathcal{S}$ by $S' \in \mathcal{S}$ if and only if there are a nonempty subset of actions $A' \in 2^A \setminus \{\emptyset\}$ and a nonempty subset of consequences $C' \in 2^C \setminus \{\emptyset\}$ such that $S' \equiv {C'}^{A'}$. It is easy to see that the collection of states space $\mathcal{S}$ forms a lattice where the order among spaces is defined by set inclusions of actions and consequences. That is, $S'' \equiv {C''}^{A''}$ is richer than $S' \equiv {C'}^{A'}$ if and only if $A'' \supseteq A'$ and $C'' \supseteq C'$. The so defined lattice of canonical state spaces is an example of an unawareness structure (Heifetz, Meier, and Schipper, 2006). Schipper (2013) presents Subjective Expected Utility where preferences can depend on state spaces in a lattice of spaces.

3 Incoherencies in the Set of All Acts

Theories of decision making under uncertainty such as Subjective Expected Utility, Choquet Expected Utility, Maxmin Expected Utility etc. typically define acts as functions from the state space to the set of consequences. I.e., an act is a function $f : S \longrightarrow C$.[5] The decision maker's preferences are then defined on the set of *all* acts. We will argue that this is problematic.

For the sake of clarity, consider the initial example with two actions, two consequences, and hence four states of Figure 1. In the lower part of Figure 4 we list all acts. There are 16 acts. Each row represents an act that specifies a consequence for each state.

[4] Previously this has been rather timidly alluded to in Gilboa (2009) Section 13.3.4. Results by Gilboa and Schmeidler (1994) and Ghirardato and Le Breton (2000, buried in Appendix A) may also be interpreted in this vein although they have been viewed till now mainly as being of technical interest. More recently, Grabiszewski (2015) and Grant et al. (2015) present more forcefully results along these lines.

[5] For simplicity, we restrict our discussion to "pure" acts. Anscombe-Aumann acts are defined by $f : S \longrightarrow \Delta(C)$ where $\Delta(C)$ denotes the set of probability measures on C.

Figure 4: All Acts in the Example with Two Actions and Consequences

	States			
Actions	s_1	s_2	s_3	s_4
own	\$100	\$100	\$0	\$0
don't own	\$100	\$0	\$100	\$0
Acts				
f_1	\$100	\$100	\$100	\$100
f_2	\$100	\$100	\$100	\$0
f_3	\$100	\$100	\$0	\$100
f_4	\$100	\$0	\$100	\$100
f_5	\$0	\$100	\$100	\$100
f_6	\$100	\$100	\$0	\$0
f_7	\$100	\$0	\$0	\$100
f_8	\$0	\$0	\$100	\$100
f_9	\$0	\$100	\$100	\$0
f_{10}	\$100	\$0	\$100	\$0
f_{11}	\$0	\$100	\$0	\$100
f_{12}	\$100	\$0	\$0	\$0
f_{13}	\$0	\$100	\$0	\$0
f_{14}	\$0	\$0	\$100	\$0
f_{15}	\$0	\$0	\$0	\$100
f_{16}	\$0	\$0	\$0	\$0

incoherent act
coherent act
incoherent outcome

Not surprisingly the primitive actions "own the firm" and "don't own the firm" can be represented by acts f_6 and f_{10}, respectively. Other acts assign different primitive actions to different states. For instance, act f_2 can be interpreted as a "contract" that would make the decision maker own the firm in states in s_1 and s_2 and not own the firm in states s_3 and s_4.

The problem is that some of these acts are incoherent in the sense that they may describe a consequence at some state that is inconsistent with the internal structure of that state. For instance, the act f_1 assigns consequence \$100 to state s_4 even though no primitive action could ever lead to consequence \$100 at that state s_4. It is not entirely clear what this act is supposed to represent and how a decision makers is supposed to be presented with a choice over such acts. There are many more such acts that we all mark in orange. The outcomes of acts that are inconsistent with the internal structure of the state we marked in red. Coherent acts that in each state assign an a consequence consistent with the consequences available at that state are marked in green.[6]

Note that the problem is not specific to the example. It is easy to see that the set of all acts defined on the canonical state space must include incoherent acts. In fact the benefit of the canonical state space is to make this inconsistency explicit. If we were to define all acts on an abstract state space as it is usually done in decision theory, we wouldn't even notice the incoherency of some acts since abstract states are silent about their internal structure (i.e., what actions lead to which consequences). That is, the incoherency of some acts in the set of all acts does not go away if an abstract state space is to represent something meaningfully about the decision context. It is just not visible with an abstract state space. Note further that the problem persists even if we start with a larger set of primitive actions or a larger set of consequences. There will be just more states in the canonical state space but the set of *all* acts must contain incoherent acts. For instance, in the above example, perhaps there is an un-modeled primitive action that at state s_4 could lead to consequence \$100. But if such an action would be represented explicitly in the canonical state space (which it should given that uncertainties are to be specified exhaustively), then we would end up with a larger state space. The set of all acts on this space would again contain some incoherent acts, etc. Note in particular that every constant acts must be incoherent, which is rather unfortunate given the important role such acts play in identifying risk preferences.

The problem of incoherencies in the set of all acts may have been alluded to already in an example by Savage (1954, p. 25) although the discussion there is far from clear. Aumann (1971) in a letter to Savage pointed out that "one could construct nonsensical acts such as 'You get sunshine if it rains, and rain otherwise." Apart from earlier mentioned references such as Gilboa (2009, p. 116), Chambers and Hayashi (2015), and Karni (2015), there is surprisingly not much discussion about the issue in the literature. We believe that when formulating decision theoretic problems with the canonical state space, the issue becomes more transparent.

4 Three Ways of Dealing with Incoherent Acts

I think there are at least three approaches to deal with incoherent acts. Each of them is unsatisfactory in the sense that it limits what can be meaningfully revealed about a decision maker's belief from choices.

[6]The problem of incoherent acts let some to claim that choice over acts on the canonical state space is generally not observable and specifically that in this example the decision maker has only one observable binary comparison between primitive actions (Gilboa, 2009, p. 116, Chambers and Hayashi, 2015). Yet, this is clearly not the case as the decision maker could choose at least among the four coherent acts and thus has six observable binary comparisons.

Ad Hoc Restrictions on the State Space

We could restrict in an ad hoc way the analysis to a suitable subset of states of the canonical state space so that all acts are coherent on this restricted space. In particular, given the canonical state space $S \equiv C^A$, consider in an ad hoc subset of states $S' \subset S$ defined by[7]

$$S' := \{s \in S : \forall c \in C \ \exists a \in A \text{ s.t. } s(a) = c\}.$$

This is the set of states in which *every* consequence occurs upon some action. Now restrict the domain of acts to such an ad hoc state space. Then acts are never incoherent as for any state (in the ad hoc state space S') it can "contract" an primitive action that yields the consequence that it assigns to this state. Choice between such acts has a clear meaning. Yet, while arguably such an approach is often assumed in decision theory (for instance when considering the Ellsberg paradox), it leaves us with decision theories that are "incomplete" (and as we have seen in Section 2 potentially not robust). They fall short of being able to reveal beliefs about all resolutions of uncertainty that may arise in the decision context. Implicitly states in $S \setminus S'$ are assumed to be null. But this remains an ad hoc assumption as behavior interpreted with such a restricted model does not imply that states in $S \setminus S'$ are null.

Restricting to the Set of All Coherent Acts

We could restrict the set of acts to coherent acts only and thereby rule out any incoherent acts. In particular, for any state $s \in S$, define

$$C(s) = \{c \in C : \exists a \in A \ \text{ s.t. } \ s(a) = c\}.$$

The set $C(s)$ is the set of consequences that in state s could be reached by some action. A *coherent act* is a function[8]

$$f : S \longrightarrow \bigcup_{s \in S} C(s) \text{ s.t. } f(s) \in C(s).$$

That is, for any state $s \in S$ the coherent act f assigns only a consequence that is actually available at s. Coherent acts are consistent with the internal structure of states. The benefit of this approach is again that now choice between such coherent acts has a clear

[7] Note that since $S \equiv C^A$, we can consider each state as a function from the set of primitive actions to the set of consequences.

[8] For simplicity, we just focus on "pure" acts. To allow for Anscombe-Aumann acts, we let

$$f : S \longrightarrow \bigcup_{s \in S} \Delta(C(s)) \text{ s.t. } f(s) \in \Delta(C(s)),$$

where $\Delta(C(s))$ denotes the set of probability measures on $C(s)$.

meaning. The drawback is that the set of coherent acts is considerably smaller than the set of all acts and thus provides much less structure to reveal beliefs from the decision maker's behavior. (For instance, in the example with two actions and two consequences of Figure 4 only 4 out of 16 acts are coherent.) That's why theories of decision making typically require the set all acts.

When restricting to the set of all coherent acts on the canonical state space, some states are now forced to be trivially null. To see this, lets define first what is a null event. For any two coherent acts f, g and event $E \subseteq S$, define the composite act $f_E g$ by

$$f_E g(s) = \begin{cases} f(s) & \text{if } s \in E \\ g(s) & \text{otherwise} \end{cases}$$

(For any two coherent acts and any event, the composite act is also coherent.) Define state s to be (Savage-)null if for any coherent acts f, g, h we have $f_{\{s\}}g \sim h_{\{s\}}g$. Then any state s for which $|C(s)| = 1$ is trivially null because all coherent acts must agree on the unique consequence in s. Yet, the framework forces it to be null rather than null being implied by behavior. One may argue that beliefs for states with just one consequence are behaviorally irrelevant since the decision maker must be indifferent between any acts conditional on such a state (because all acts yield the same consequence in that state). Thus, no matter what probability we would attach to such states, it is not relevant to decisions using expected utilities. While this is true for Subjective Expected Utility, it is not true in general for generalizations such as Choquet Expected Utility, Maxmin Expected Utility, etc. To see this, consider the example of Figure 5. There are three states and two actions. Both actions yield -\$36 in state s_3. Thus, we may want to considered this state to be "irrelevant". The middle part shows a table with two different probability distributions, p and q, on the state space while the lower part shows the corresponding probability distributions on the space $\{s_1, s_2\}$ only. (I.e., distribution p' preserves the relative likelihood ratio of p for all states where p' is defined; similarly for q'.)

Figure 6 presents the expected values for both actions and probability distributions p, q, p', and q'. Clearly, when focusing on expected utility, it wouldn't matter whether we make decisions with p or p' (or alternatively q or q'). However, under Maxmin Expected Utility the decision maker uses the minimum of expected value over p and q. Hence, she would choose a over b when forming beliefs over the entire space $\{s_1, s_2, s_3\}$. If instead she forms corresponding beliefs p' and q' over states $\{s_1, s_2\}$ only (i.e., excludes the "irrelevant" state s_3), then she would choose b over a. Hence, a state that is irrelevant under expected utility because any act yields the same consequence in that state may become behaviorally relevant under generalizations of expected utility. Thus, restricting to coherent acts only still limits what can be possibly revealed about a decision maker's belief.

Restricting to the set of coherent acts is reminiscent of rare approaches to decision theory with state-dependent consequence sets although the motivation of those theories is different (see Fishburn, 1970, Section 13.2, following Pratt, Raiffa and Schlaifer, 1964,

Figure 5: Example with a States with a Single Consequence

	States		
Actions	s_1	s_2	s_3
a	\$132	\$0	-\$36
b	\$36	\$60	-\$36
	Prob measures on $\{s_1, s_2, s_2\}$		
p	1/3	1/3	1/3
q	1/4	1/2	1/4
	Corresponding prob measures on nontrival states $\{s_1, s_2\}$		
p'	1/2	1/2	
q'	1/3	2/3	

Figure 6: Relevance of a State with a Single Consequence

	Actions	
	a	b
E_p	32	20
E_q	24	30
Min	24	20
$E_{p'}$	66	48
$E_{q'}$	44	52
Min	44	48

as well as Hammond, 1999). All these approaches require that $|C(s)| \geq 2$ for all s and thus still have to (implicitly) restrict to an ad hoc space of states as discussed earlier if they want to completely avoid incoherent acts.

Simply Allow for Incoherence

Finally, we may simply allow for incoherent acts. The question then is how to design choices among such acts and how to interpret such choices. Intuitively, an incoherent act is not unthinkable. Think of a contract written in an incoherent way (e.g., containing a contradiction). Moreover, choices between incoherent acts are not necessarily unobservable in reality. Arguable many contracts that we choose in reality may be inconsistently specified in some way or another. That's may be one of the reasons why some contractual relationships end up in courts of law as there can be disagreement about the consequences of incoherent contracts in various states. So allowing for incoherent acts may make decision theory more realistic. While these arguments are somewhat compelling at the intuitive level, it is less clear how to analyze explicitly inconsistences in a satisfactory way.

When Aumann (1971) wrote to Savage regarding what appears to be incoherent acts, Savage himself seemed to be "prepared to live with them until something better comes along" (Savage, 1971). He regarded "it as fanciful but not as nonsense to say 'You experience sunshine if it rains, and rains otherwise" and claimed that he "can contemplate the possibility that the lady dies medically and yet is restored in good health to her husband." It appears that 45 years after this exchange of letters nothing better came along. It might be possible to express such incoherent acts in logical frameworks that allow for inconsistencies such as impossible possible worlds models (e.g. Rantala, 1982). It is far less clear though how would one always operationalize choice experiments involving incoherent acts. The fact that decision theory is mostly silent on this issue and tacitly accepts incoherent acts reveals that it grounded less on the revealed preference paradigm than often claimed (e.g. Gul and Pesendorfer, 2008).

References

[1] Anscombe, F.J. and R. Aumann (1963). A definition of subjective probability, Annals of Mathematics and Statistics 34, 199-205.

[2] Aumann, R. (1971). Letter from Robert Aumann to Leonard Savage, 8 January 1971, Dreze, J.H. (ed.), Essays on Economic Decisions under Uncertainty, Cambridge University Press, Cambridge, 1987, pp. 76–78.

[3] Arrow, K. (1971). Essays in the theory of risk-bearing, Markham Publishing, Chicago.

[4] Chambers, C. and T. Hayashi (2015) Reverse Bayesianism: A Comment, mimeo.

[5] Chellas, B.F. (1980). Modal logic: An introduction, Cambridge University Press.

[6] Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms, Quarterly Journal of Economics 75, 643–669.

[7] Fishburn, P. (1970). Utility theory for decision making, John Wiley & Sons.

[8] Ghirardato, P. and M. Le Breton (2000). Choquet rationality, Journal of Economic Theory 90, 277–285.

[9] Gibbard, A. and W.L. Harper (1978). Counterfactuals and two kinds of expected utility, in: Hooker, Leach, and McClennen, (eds.), Foundations and Applications of Decision Theory, Vol. 1, D. Reidel, 125–62.

[10] Gibbard, A. and W.L. Harper (1976). Counterfactuals and two kinds of expected utility, working paper.

[11] Gilboa, I. (2009). Theory of decision under uncertainty, Cambridge University Press.

[12] Gilboa, I. and D. Schmeidler (1994). Additive representations of non-additive measures and the Choquet integral, Annals of Operations Research 52, 43–65.

[13] Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with a non-unique prior, Journal of Mathematical Economics 18, 141–153.

[14] Grabiszewski, K. (2015). On the rejectability of the subjective expected utility theory, mimeo.

[15] Grant, S., Kline, J., Meneghel, I., Quiggin, J., and R. Tourky (2015). A theory of robust experiments for choice under uncertaity, mimeo.

[16] Gul, F. and W. Pesendorfer (2008). The case for mindless economics, in: Caplin, A. and A. Schotter (eds.), The foundations of positive and normative economics: A handbook, Oxford University Press, pp. 3–39.

[17] Hammond, P. (1999). Subjectively expected state-independent utility on state-dependent consequence domains, in: Machina, M.J. and B. Munier (eds.) Beliefs, interactions, and preferences in decision making, Kluwer Academic,7–21.

[18] Heifetz, A., Meier, M. and B.C. Schipper (2006). Interactive unawareness, Journal of Economic Theory 130, 78–94.

[19] Karni, E. (2015). States of nature and the nature of states, mimeo.

[20] Karni, E. and D. Schmeidler (1991). "Utility theory with uncertainty", in: Handbook of Mathematical Economics, Vol. 4, Hildenbrand, W. and H. Sonnenschein (eds.), Elsevier, 1763–1831.

[21] Karni, E. and M.-L. Vierø (2015). Probabilistic sophistication and reverse Bayesianism, Journal of Risk and Uncertainty 50, 189–208.

[22] Karni, E. and M.-L. Vierø (2013). "Reverse Bayesianism": A choice-based theory of growing awareness, American Economic Review 103, 2790–2810.

[23] Machina, M. (2003). States of the world and the state of decision theory, in: Meyer, D. (ed.), The economics of risk, Chapter 2, W.E. Upjohn Institute for Employment Research.

[24] Pratt, J.W., Raiffa, H., and R. Schlaifer (1964). The foundations of decision under uncertainty: An elementary exposition, Journal of the American Statistical Association 59, 353–375.

[25] Rantala, V. (1982). Impossible worlds semantics and logical omniscience, Acta Philosophica Fennica 35, 18–24.

[26] Savage, L.J. (1954). The foundations of statistics, John Wiley & Sons.

[27] Savage, L. J. (1971). Letter from Leonard Savage to Robert Aumann, 27 January 1971, in: Dreze, J.H. (ed.), Essays on Economic Decisions under Uncertainty, Cambridge University Press, Cambridge, 1987, pp. 78–81.

[28] Schmeidler, D. (1989). Subjective probability and expected utility without additivity, Econometrica 57, 571–587.

[29] Schmeidler, D. and P. Wakker (1987). Expected utility and mathematical expectation, in: The New Palgrave: A Dictionary of Economics, 1st edition, Eatwell, J., Milgate, M., and P. Newman (eds.), Macmillan Press.

[30] Schipper, B.C. (2013). Awareness-dependent subjective expected utility, International Journal of Game Theory 42, 725-753.