Journal of Mathematical Economics 60 (2015) 43–48

Contents lists available at ScienceDirect

Journal of Mathematical Economics

journal homepage: www.elsevier.com/locate/jmateco

Worst-case expected utility ☆

Shiri Alon

The Department of Economics, Bar Ilan University, Ramat-Gan, Israel

ARTICLE INFO

Article history:
Received 2 October 2014
Received in revised form
12 April 2015
Accepted 24 June 2015
Available online 2 July 2015

Keywords:
SEU
CEU
MEU
Uncertainty aversion
Neo-additive capacity
Unforeseen contingencies

ABSTRACT

The paper presents a model in which a decision maker, having a preference relation over purely subjective acts, slightly deviates from the Subjective Expected Utility decision rule, exhibiting an uncertainty averse behavior *á*-la Schmeidler (1989). The resulting representation is as if the decision maker adds to the formulation of the problem one new state, representing the occurrence of some unforeseen event. Each Savage act is extended to the new, endogenous state by assigning this state with the worst consequence the act obtains on all other, primitive states. On the extended decision problem a Subjective Expected Utility rule is applied. The representation thus expresses the common practice of a 'worst-case scenario' assumption as means to cope with unforeseen contingencies. The model is a special case of the neo-additive capacities model of Chateauneuf, Eichberger and Grant.

1. Introduction

The paper presents an axiomatic model of decisions in which the decision maker is assumed to slightly deviate from the Subjective Expected Utility paradigm of Savage (1952, 1954), and where all deviations are being driven by uncertainty aversion. The decision maker is modeled by a binary relation over Savage acts, and the representation characterized is as if the decision maker adds a new state to the decision problem and extends each act to this new state. The new, endogenous state may be interpreted as signifying 'some unforeseen event occurs'. In the representation caution is exhibited in that the consequence assumed for each act on the endogenous state is its worst consequence over all primitive (foreseen) states. On the extended decision problem, containing the new state, a Subjective Expected Utility rule is applied. The representation thus expresses the common practice of assuming a worst-case scenario in the face of the unknown, and the model is accordingly called Worst-Case Expected Utility (shortened to Worst-Case EU in the sequel). The model demonstrates that uncertainty averse behavior can emerge from a seemingly small deviation from Subjective Expected Utility.

The setup employed in the paper consists of a rich set of consequences and an unconstrained (possibly finite) state-space. A Subjective Expected Utility decision maker in such a setup may be characterized by a Tradeoff Consistency assumption that holds unconditionally (see Kobberling and Wakker, 2003). The decision maker in the current model, on the other hand, is assumed to satisfy a weakened form of consistency that holds whenever all acts considered obtain their worst consequence on the same event. Otherwise the decision maker follows an Uncertainty Aversion assumption á-la Schmeidler (1989). To compare, the form of consistency assumed here is stronger than the type required for a non-additive expected utility representation, as in Schmeidler (1989). Put together, the consistency axiom of the model and Uncertainty Aversion identify a preference of the decision maker to hedge his or her worst-case scenario, by averaging the worst consequence.

The model is a special case of the neo-additive capacities model of Chateauneuf et al. (2007), which is placed in a purely subjective framework as the one employed here. The neo-additive model characterizes preferences that simultaneously overweigh best and worst consequences. The two models are closely related, and the differences originate from the fact that the fundamental feature of a Worst-Case EU decision maker is uncertainty aversion, interpreted as a cautious response to unforeseen events, whereas within the neo-additive model the decision maker may exhibit uncertainty aversion with respect to some events, and uncertainty attraction with respect to others. Accordingly, the Worst-Case EU is characterized by violating Subjective Expected Utility only in an uncertainty averse fashion, while in the neo-additive model these violations depend on the acts involved, specifically on their sharing of events on which best and worst consequences are obtained.

☆ The author is thankful to David Schmeidler, Tzachi Gilboa, Gabi Gayer, Fabio Maccheroni and Peter Wakker for helpful comments. Support of the ISF through Grant #1188/14 is thankfully acknowledged.

E-mail address: Shiri.Alon-Eron@biu.ac.il.

http://dx.doi.org/10.1016/j.jmateco.2015.06.017
0304-4068/

Since the neo-additive model is a special case of the non-additive expected utility model first introduced by Schmeidler (1989) (henceforth abbreviated as CEU, for Choquet Expected Utility), so is the Worst-Case EU decision rule. More precisely, the Worst-Case EU rule, unlike the neo-additive rule, is a special case of the CEU rule with a convex non-additive probability,[1] thus also of the Maxmin Expected Utility rule of Gilboa and Schmeidler (1989) (abbreviated as MEU).

In fact, the Worst-Case Expected Utility model can be formulated as an MEU rule with respect to an ε-contaminated set of priors, where the contaminating set of priors is the entire simplex of probabilities, and ε is endogenous. Such an ε-contamination rule was axiomatized by Kopylov (2006),[2] however Kopylov's characterization is done in an Anscombe–Aumann framework (Anscombe and Aumann, 1963) which includes exogenous probabilities, while the characterization given here is purely subjective. In addition, the view taken here, which is expressed in the different axioms employed, is of the model as a simple, uncertainty averse departure from Subjective Expected Utility.

The Worst-Case EU model is related to models that entertain two associated state spaces. In Jaffray and Wakker (1994) and in Mukerji (1997) one of the two state spaces assumed is the primitive payoff-relevant state space on which acts are defined, so that each payoff-relevant state is assigned a single consequence. The other state space is an underlying space, on which there is complete probabilistic knowledge through an additive probability measure. Each state in the probabilizable space maps to a set of payoff-relevant states, thus additive probabilistic evaluations on the first space become non-additive when translated to the second. Jaffray and Wakker show how this structure leads to a preference relation over belief functions, and axiomatize a decision rule over belief functions. Mukerji identifies epistemic foundations which lead to a non-additive probability assessment in this setup, and then discusses the links to a CEU decision rule. Mukerji suggests a cautious mapping of acts from the payoff-relevant space to the underlying, probabilizable space, where each underlying state is assigned the worst consequence obtained on the payoff-relevant states. Gilboa and Schmeidler (1994) show that a CEU decision rule over a state space is equal to an additive representation over a corresponding 'grand' state space. Each state in the 'grand' space represents an event from the original, primitive state space, and acts are extended to the larger state space by assuming their worst consequence on every event. Gilboa and Schmeidler interpret this duality as a sign that the primitive state space is misspecified, and does not contain all actual states in its formulation.

The Worst-Case EU model is a special case of the two-tiered state space. Compared to Jaffray and Wakker (1994) and Mukerji (1997), the underlying, probabilizable state space contains only one state in addition to the payoff-relevant space—the endogenously derived state, which maps to all payoff-relevant states. In the Gilboa and Schmeidler (1994) setup, it is a special case by translating the non-additive probability over the primitive state space to an additive measure over the grand state space, so that only the states which represent either singletons or the entire state space are assigned a non-zero probability. Among these three last papers, only the one by Jaffray and Wakker contains axiomatization. Their axiomatization, however, is given in a setup different from that of the current paper, as they assume a primitive knowledge structure of a two-tiered state-space.

[1] That is, the non-additive probability ν satisfies, for every pair of events E and F, $\nu(E \cup F) \geq \nu(E) + \nu(F) - \nu(E \cap F)$.

[2] Kopylov also axiomatizes the general case of ε-contamination, see both Kopylov (2006) and Kopylov (2008). See also Nishimura and Ozaki (2006) for an axiomatization of ε-contamination with an exogenous ε.

Other, more distantly related models, suppose existence of unforeseen contingencies as part of their setup, by allowing acts to obtain a *set* of consequences on each state. Works of Nehring (1999), Ghirardato (2001), Jaffray and Jeleva (2011), and Vierø (2009) are consistent with the view that a formulated state may in fact represent a *set* of states which the decision maker cannot tell apart (this is usually referred to as *coarse contingencies*). An even more general approach is taken by Karni and Viero (2013), that models a decision maker who gradually becomes aware of new components of the decision problem, consequences and states alike, and responds to those in a reverse Bayesianism manner. In a different, extensive-form setup that allows for history-dependent decision rules, Grant and Quiggin (2013) model a decision maker with growing awareness, who may realize, based on past experience, that there are possible contingencies he or she is unaware of. Grant and Quiggin suggest decision rules that apply to such bounded awareness problems, offering a cautious decision rule under such circumstances. As opposed to these unforeseen contingencies models, the model presented in this paper is placed in a standard Savage setup, where the very possibility of unforeseen contingencies is derived axiomatically from preferences, apparent only in the representation. At the same time, the multi-consequence and growing awareness models mentioned accommodate a richer description of unforeseen elements and of the decision maker's attitude toward them.

Lastly I mention that in the context of preferences over lotteries pessimistic departures from expected utility as described here were axiomatized by Gilboa (1988) and by Jaffray (1988), and departures that may be both pessimistic and optimistic were characterized by Cohen (1992) (further elaboration may be found in Chateauneuf et al., 2007).

The Worst Case EU representation is aimed to be as simple as possible but still accommodate the question of unforeseen contingencies. Within the representation, the procedure of assigning to the endogenous, 'some unforeseen event occurs' state the worst consequence of each act is but one possibility, which is a simple expression of a cautious conduct. Other modes of extension of acts may be considered, such as assigning to the endogenous state a globally worst consequence, constant for all extended acts. This kind of extension would naturally lead to a standard Subjective Expected Utility rule. However, assigning the endogenous state with the same consequence for all acts would mean that what happens on this state is independent of the act chosen. By contrast, the choice made here reflects an interpretation that the consequence in case an unforeseen event occurs depends on the act, and the endogenous state does not represent a total catastrophe where all acts collapse to the same worst consequence. Put differently, for each alternative faced, the worst-case scenario *of this alternative* is considered.

The paper is organized as follows. In the next section there is a description of the notation and standard axioms. Section 3 contains an explanation of the Tradeoff Consistency axiom adopted and a formulation of the Uncertainty Aversion assumption of the model. The section concludes with the representation theorem of the model, and an equivalent MEU representation. Proofs are given in Section 4.

2. Notation and basic axioms

Let S denote the set of *states of nature*, endowed with a sigma-algebra of *events*, Σ. Let X be a nonempty set of *consequences*, and $\mathcal{F}$ the set of *acts*, which are taken to be finite-valued mappings from S to X, measurable w.r.t. Σ (i.e., simple acts).

The decision maker's preference relation over acts is denoted by $\succsim$, with $\sim$ and $\succ$ being its symmetric and asymmetric components. With the usual slight abuse of notation, x sometimes denotes the