International Journal of Approximate Reasoning 51 (2010) 1100–1113

Contents lists available at ScienceDirect

International Journal of Approximate Reasoning

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

Sensitivity of decisions with imprecise utility trade-off parameters using boundary linear utility

Malcolm Farrow [a,*], Michael Goldstein [b]

[a] School of Mathematics and Statistics, Newcastle University, Newcastle upon Tyne NE1 7RU, UK
[b] Department of Mathematical Sciences, Durham University, South Road, Durham DH1 3LE, UK

ARTICLE INFO

Article history:
Available online 22 August 2010

Keywords:
Robust decisions
Imprecise utilities
Utility hierarchies
Mutual utility independence
Boundary linear utility
Sensitivity analysis

ABSTRACT

In earlier work we have developed methods for analysing decision problems based on multi-attribute utility hierarchies, structured by mutual utility independence, which are not precisely specified due to unwillingness or inability of an individual or group to agree on precise values for the trade-offs between the various attributes. Our analysis is based on whatever limited collection of preferences we may assert between attribute collections. In this paper we introduce three methods to assess the robustness of our selected decision. Two of these use the metric generated by the weights of the boundary linear utility. The third applies directly to the space of trade-off parameters.

1. Introduction

In two earlier papers we have developed a methodology for decision analysis with multi-attribute utilities which does not require the specification of precise trade-offs between different risks. Multi-attribute utilities may be imprecisely specified, due to an unwillingness or inability on the part of a client to specify fixed risk trade-offs or because of disagreement within a group with responsibility for the decision.

In [3] we introduced our approach to constructing imprecise multi-attribute utility hierarchies. We described the structure which we use, which is based on a utility hierarchy with utility independence at each node, and explained the notion of imprecise utility trade-offs for such a hierarchy, based on limited collections of stated preferences between outcomes. This leads to a set R of possible trade-off specifications $\underline{\theta}$. In this context we used a concept of Pareto optimality to reduce the set of alternatives. These methods and some associated theory are summarised in Section 2 of this paper.

We are particularly concerned with problems where the number of alternatives among which we must choose is large. Many real decision problems, for example in experimental design, have very large spaces of possible choices. Relaxing the requirement for precise trade-off specification reduces our ability to eliminate choices by dominance and can leave us with a large class of choices, none of which is dominated by any other over the whole range of possible trade-offs allowed by the imprecise specification. We are therefore faced with the need for methods to reduce the decision space which are tractable even when the decision space is very large and there is a complicated multi-attribute utility structure to consider. Such methods should respect the range R of trade-offs and favour choices which perform well, in some sense, compared to others over the whole of R. In [4] we described ways to reduce the class of alternatives that we must consider, by eliminating choices which are "ε-dominated" and combining choices which are "ε-equivalent". We explored the effects of different values of ε and of different parts of the hierarchy to see when and why choices are eliminated.

* Corresponding author.
E-mail addresses: Malcolm.Farrow@newcastle.ac.uk (M. Farrow), Michael.Goldstein@durham.ac.uk (M. Goldstein).

0888-613X/$ - see front matter
doi:10.1016/j.ijar.2010.08.002

To choose a single choice d^* from our reduced list, we can use the boundary linear utility approach described in [3], or choose the choice which is the last to be eliminated as we increase the value of our ε criterion as described in [4]. We can then find the set D^* of choices which are "almost equivalent" to d^* and perhaps use secondary considerations to choose among them. We review boundary linear utility in Section 3 of this paper.

In Section 4 we describe methods, based on the boundary linear utility, for exploring the sensitivity of possible choices to variation in the utility trade-offs. This helps us to find a decision which, as far as possible, is a good choice over the whole range of possible trade-offs.

The practical implementation of our approach is illustrated throughout by an example concerning the introduction of a new course module at a university, which we first described in [4].

2. Mutually utility independent hierarchies and imprecise utility trade-offs

2.1. Mutually utility independent hierarchies

In [3] we proposed a general class of multi-attribute utility functions. This uses the concept of mutual utility independence among sets of attributes in order to impose a structure on the utility function. Attributes $\underline{Y} = (Y_1, \dots, Y_k)$ are *utility independent* of the attributes $\underline{Z} = (Z_1, \dots, Z_r)$ if conditional preferences over lotteries with differing values of $\underline{Y}$ but fixed values, $\underline{z}$, of $\underline{Z}$, do not depend on the particular choice of $\underline{z}$. Attributes $\underline{X} = (X_1, \dots, X_s)$ are *mutually utility independent* if every subset of $\underline{X}$ is utility independent of its complement. If attributes $\underline{X}$ are mutually utility independent, then the utility function for $\underline{X}$ must be given by the *multiplicative form*

$$1 + kU(\underline{X}) = \prod_{i=1}^{s} [1 + ka_i U_i(X_i)] \tag{1}$$

or the *additive form*

$$U(\underline{X}) = \sum_{i=1}^{s} a_i U_i(X_i) \tag{2}$$

(see [7]) where $U_i(X_i)$ is a conditional utility function for attribute X_i, namely an evaluation of the utility of X_i for fixed values of the other attributes. The coefficients in (1) and (2) are the *trade-off parameters*. The a_i reflect the relative importance of the attributes. For a more detailed discussion of the role of k, see Chapter 6 of [7]. In the case where $s = 2$, k reflects the degree to which rewards may be regarded as complementary, if $k > 0$, or as substitutes, if $k < 0$.

The assumption of mutual utility independence, which many people would often be prepared to make, is enough in itself to reduce the problem to one of considering a finite number of parameters.

Keeney and Raiffa [7] also describe the idea of a hierarchy of utilities, as follows. We form an overall multi-attribute utility from marginal utilities for the various attributes by a hierarchical structure in which, at each node, several utilities are merged into a combined utility. This combined utility is merged with others at a node in the next level until, finally, one overall utility function is formed. If, at each node, we have mutual utility independence for the utilities combined at that node, then we term such a utility function a *Mutually Utility Independent Hierarchic (MUIH)* utility. Thus, in a MUIH utility, at each node we combine utilities using either (1) or (2). An example of such a hierarchy is shown in Fig. 1.

Our hierarchical structure allows us to relax the requirement for overall mutual utility independence by allowing the user to specify utility independence just at the nodes of the hierarchy and, of course, the user can choose this structure.

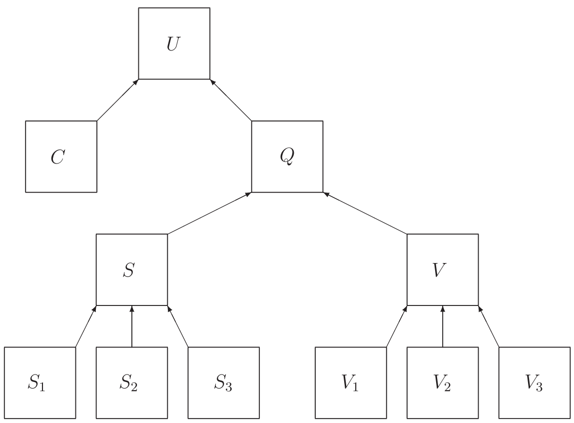

Fig. 1. Utility hierarchy for the course design example.