Omega 64 (2016) 13–23

Contents lists available at ScienceDirect

Omega

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

Modelling uncertainty in stochastic multicriteria acceptability analysis ☆

Ian N. Durbach [a,b,*], Jon M. Calder [a]

[a] Department of Statistical Sciences, University of Cape Town, Rondebosch 7701, South Africa
[b] African Institute for Mathematical Sciences, 6-8 Melrose Road, Muizenberg 7945, South Africa

ARTICLE INFO

Article history:
Received 27 February 2015
Accepted 29 October 2015
Available online 6 November 2015

Keywords:
Decision making/process
Decision support systems
Multicriteria
Risk
Sensitivity analysis

ABSTRACT

This paper considers problem contexts in which decision makers are unable or unwilling to assess trade-off information precisely. A simulation experiment is used to assess (a) how closely a rank order of alternatives based on partial information and stochastic multicriteria acceptability analysis (SMAA) can approximate results obtained using full-information multi-attribute utility theory (MAUT) with multiplicative utility, and (b) which characteristics of the decision problem influence the accuracy of this approximation. We find that fairly good accuracy can be achieved with limited preference information, and is highest if either quantiles and probability distributions are used to represent uncertainty.

1. Introduction

When facilitating decisions it sometimes happens that some inputs to the preference model either cannot be assessed at all or can only be assessed within relatively large bounds of uncertainty (e.g. [1,12,13]). This can happen for a number of reasons: a lack of time, a politically sensitive problem context, or a lack of decision maker (DM) involvement, for example. Whatever the reason, in these cases the DM is unable or unwilling to express him or herself with the degree of precision required by conventional decision aids. We call decision problems which must be addressed under such conditions "low-involvement" decisions. The question is what, if any, decision support can be provided in such situations.

Stochastic multicriteria acceptability analysis (SMAA [25,19]) is a family of decision models that can be used with arbitrarily precise preference information. It addresses low-involvement decision-making by providing information about the types of preferences (if any) that would lead to the selection of each alternative.

In this paper we use a simulation experiment to evaluate the ability of SMAA to approximate results obtained using multi-attribute utility theory (MAUT) where preferences are represented by a multiplicative utility function. In particular, we ask how closely results computed from a key output from SMAA (the acceptability index) can approximate those obtained using MAUT. In doing so we hope to provide a broad indication of the losses that are possible if facilitators choose to use a low-involvement decision aid such as SMAA rather than compelling DMs to be more precise in their assessment of certain types of preference information – for example, using more detailed problem structuring. We also wish to test the robustness of the SMAA approach to various aspects of the decision process: the size of the decision problem, the way attribute evaluations are distributed, the underlying preference functions, the accuracy of assessed information, the amount of preference information gathered, and the way in which the acceptability index is constructed.

In addition to the conventional SMAA model, which uses probability distributions to represent uncertainty in the attribute evaluations, we also introduce and evaluate a number of 'simplified' models which make use of summarised measures of uncertainty instead of a full probability distribution [6]. By assessing the accuracy of both conventional and simplified uncertainty models under a range of different conditions we hope to provide motivation for the use of simplified models in appropriate circumstances. A similar approach has been used in Durbach and Stewart [9] to assess the effect of using simplified uncertainty formats in general decision-making, and we employ a similar simulation structure in the current paper.

Our view is that in nearly all cases it is preferable to resolve preferential uncertainty through discussion and problem structuring rather than by employing more 'lenient' decision aids, because of the additional insight and opportunities for learning. We focus on those circumstances in which the DM is unable or unwilling to participate fully in this process. In using a simulation experiment, we acknowledge that we can only evaluate the extent to which

☆This manuscript was processed by Associate Editor Doumpos.

* Corresponding author at: Department of Statistical Sciences, University of Cape Town, Rondebosch 7701, South Africa.
E-mail address: ian.durbach@uct.ac.za (I.N. Durbach).

http://dx.doi.org/10.1016/j.omega.2015.10.015
0305-0483/

using SMAA rather than MAUT might impact results. We cannot evaluate critical issues like whether the reduced time spent on problem structuring in SMAA is "worth" the reduction in decision quality, or the degree to which the problem structuring process, through the insight it generates for the DM, is useful as an end in itself. Simulation results are unable to provide general conclusions on the viability of different methods, but provide inputs to such discussion by identifying the potential trade-offs in accuracy that are implied when using a simplified model. Ultimately accuracy must be weighed against other factors to determine which decision model may be most appropriate for a problem.

The remainder of the paper is structured as follows. Section 2 provides a review of the relevant background literature and notation. Section 3 describes the structure of our simulation experiment. Section 4 presents the results in a direct fashion, delaying a more detailed discussion of these until Section 5, which also provides theoretical justification for some key findings using results from applied probability theory. Section 6 discusses implications of the simulation results for the use of SMAA as a prescriptive decision aid, and concludes.

2. Notation and background

Consider a decision problem consisting of I alternatives $\{a_1, a_2, \dots, a_I\}$ evaluated on J attributes $\{c_1, c_2, \dots, c_J\}$. Let Z_{ij} be a random variable denoting the attribute evaluation of a_i on c_j, and U be a multi-attribute utility function mapping the attribute evaluations of alternative a_i (denoted $\mathbf{Z}_i$) to a real value using a weight vector $\mathbf{w}$. A joint density function $f_X(\mathbf{Z})$ governs the generation of the Z_{ij} in the space $X \subseteq R^{I \times J}$, and a second joint density function $g(\mathbf{w})$ governs the generation of imprecise or unknown weights in the weight space W. Total lack of knowledge is usually represented by a uniform distribution in W. If restrictions have been placed on W we denote the feasible weight space by W'.

The original SMAA method [20] analysed the combinations of attribute weights that result in each of a set of alternatives being selected when using an additive utility function. Subsequently, a number of SMAA variants have been developed. These differ in terms of the preference model used and thus the type of preference information that is imprecisely known, but are all based upon Monte Carlo simulation from distributions which govern unknown preference parameters (and attribute evaluations). For example, SMAA variants are available for value function [20,17], outranking [10], reference point [21,5], prospect theory [18], Choquet integral [2], and AHP [7] methods. Comprehensive reviews are given by Tervonen and Figueira [25] and Lahdelma and Salminen [19].

Given a particular weight vector $\mathbf{w}$, the global utility of each alternative can be computed and a rank ordering of alternatives obtained. SMAA-2 [17] is based on simulating a large number of random weight vectors from $g(\mathbf{w})$ and observing the proportion and distinguishing features of weight vectors which result in each alternative obtaining a particular rank r (usually the "best" rank, $r=1$), using an additive value function model. Let the set of weight vectors that result in alternative a_i obtaining rank r be denoted by W_i^r. SMAA is based on an analysis of these sets of weights using a number of descriptive measures, the most important of which are:

Acceptability indices: The rank-r acceptability index b_i^r measures the proportion of all simulation runs, i.e. weight vectors, that make alternative a_i obtain rank r. A cumulative form of the acceptability index called the R-best ranks acceptability index is defined as $\mathcal{B}_i^R = \sum_{r=1}^{R} b_i^r$ and measures the proportion of all weight vectors for which alternative a_i appears anywhere in the best R ranks. In the discussion in Section 5 we make use of ordered acceptability indices, where we denote the alternative with the k-th largest rank-r acceptability index by $a_{(k)}^r$, and its acceptability index by $b_{(k)}^r$.
Central weight vectors: The central weight vector $\mathbf{w}_i^c$ is defined as the expected center of gravity of the favourable weight space W_i^1. It gives a concise description of the "typical" preferences supporting the selection of a particular alternative a_i, and in practice is computed from the empirical (element-wise) averages of all weight vectors supporting the selection of a_i as the best alternative.

The exact number of Monte Carlo iterations that are required to achieve a given accuracy is discussed in Tervonen and Lahdelma [26]. To estimate the acceptability index within δ of the true value with 95% confidence, one requires $1.96^2/4\delta^2$ iterations – so that 10 000 iterations will usually be sufficient to achieve error bounds of 1%.

Uncertain attribute evaluations are conventionally treated in SMAA using probability distributions, with each simulation run drawing values at random from these distributions. Adapting SMAA to use other uncertainty formats, however, is generally straightforward, as described in Durbach and Davis [6]. Each uncertain attribute is simply replaced by a number of lower-level attributes which capture the uncertainty in the evaluations on that attribute, using one of many possible simplified uncertainty formats. In this paper, we test the effect of using three different formats: expected values; variances; and quantiles. This transforms the decision problem into one having the same appearance as a deterministic decision problem, which can be treated by any of the existing SMAA models with some minor modifications. We sometimes refer to these collectively as "SMAA models" although it should be clear that these are approximations of external uncertainty built on top of the same SMAA approach. The models are described in more detail in Section 3. Table 1 provides a summary of notation used in the paper.

3. Description of the simulation study

The general structure of the simulation[1] is summarised in Fig. 1 and implements the following basic steps:

1. Form a hypothetical problem context, generating the relevant attribute evaluations.
2. Apply a multiplicative MAUT model to derive idealised or 'true' utilities and thus find the 'true' rank ordering of alternatives.
3. Calculate summarised measures of uncertainty based on the generated data and incorporating observational errors.
4. Run different models using SMAA and the inputs from step 3, and then compare the model results against the 'true' utilities and rank order obtained from step 2.

3.1. Generating attribute evaluations

We consider a decision problem involving I alternatives evaluated over J attributes. External uncertainty about attribute performance is captured by simulating attribute evaluations for alternative a_i on attribute c_j from a gamma distribution with mean μ_{ij}, standard deviation σ_{ij}, and skewness ξ_{ij}. We denote this distribution $f_{ij}(Z_{ij})$. Each mean μ_{ij} is drawn randomly from a uniform distribution between 0.2 and 0.8 for all alternatives on all attributes. Means are then standardised across all the attributes to lie on the unit hypersphere i.e.

[1] All codes used to run the simulations described in this section are openly available from http://dx.doi.org/10.5281/zenodo.30523.