Omega 64 (2016) 1–12

Contents lists available at ScienceDirect

Omega

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

Entropy-optimal weight constraint elicitation with additive multi-attribute utility models ☆

Gert van Valkenhoef [a,*], Tommi Tervonen [b]

[a] Department of Epidemiology, University Medical Center Groningen, University of Groningen, The Netherlands
[b] Evidera Ltd, London, United Kingdom

ARTICLE INFO

Article history:
Received 11 March 2015
Accepted 27 October 2015
Available online 1 November 2015

Keywords:
Multicriteria
Decision making/process
Probability
Information theory

ABSTRACT

We consider the elicitation of incomplete preference information for the additive utility model in terms of linear constraints on the weights. Eliciting incomplete preferences using holistic pair-wise judgments is convenient for the decision maker, but selecting the best pair-wise comparison is difficult. We propose a framework for comparing holistic preference elicitation questions based on their expected information gain, and introduce a procedure for approximating the optimal question. We extend the basic approach to generate reference alternatives that differ on only a few attributes, and to determine when further preference information is unlikely to reduce decision uncertainty. We present results from computational experiments that assess the performance of the procedure and assess the impact of limiting the number of attributes on which the reference alternatives differ. The tests show that the proposed method performs well, and when implemented in a decision support system it may substantially improve on-line elicitation using pair-wise comparisons.

1. Introduction

We consider multi-attribute decision problems where a set of alternatives indexed with $I=\{1,\dots,m\}$ are evaluated based on their uncertain measurements on a set of attributes indexed with $J=\{1,\dots,n\}$. Based on the evaluations, the DM either needs to rank the alternatives, or to choose a small subset (possibly of size one) of best alternatives. The alternatives' attribute measurements x_j^i, $\forall (i,j)\in I\times J$, have a joint density $f_X(x)$. We assume that the DM's preferences are representable with an additive multi-attribute utility function u that consists, $\forall j\in J$, of partial utility functions $u_j:\mathbb{R}\to[0,1]$ and their scaling factors (weights) w_j:

$$u(x^i,w)=\sum_{j\in J} w_j u_j(x_j^i), \tag{1}$$

where x^i is the vector of attribute measurements for alternative $i\in I$. The model (1) has two types of preference information that need to be elicited from the DM: partial utility functions u_j that simultaneously model the DM's risk attitude and the attractiveness of the attribute scale levels, and the weights w_j that express relative importance of the attribute scale swings, i.e. trade-offs [19].

There are various textbook methods for eliciting the partial utility functions and the weights. For example, the partial utility functions u_j can be obtained with the standard gamble method and the weights w_j with the swing weighting method [19]. However, research has shown that utility and weight elicitation are cognitively demanding and prone to behavioral biases (see e.g. [4,10,25,41]). Various techniques have been developed to reduce these biases and to consequently make the additive model applicable in practical decision aiding settings. For example, instead of eliciting exact and complete weight information, preference information about the weights can be elicited in an incomplete format, which leads to having linear weight constraints [27,29]. However, there is little evidence that eliciting such weight constraints directly would not be affected by the same behavioral biases that affect the direct methods for exact weight elicitation. Obtaining the constraints indirectly by asking the DM to provide preference information in the form of holistic pair-wise judgments over a given pair of reference alternatives with deterministic evaluations (e.g. x^1 is preferred over x^2 or vice versa) has been proposed as a practical technique for lowering the DM's cognitive burden [8,15,16].

On the other hand, instead of true utility functions, one can apply simpler value functions that do not model the DM's risk attitude but solely quantify the attractiveness of deterministic outcome levels. Uncertainty in the attribute measurements and incomplete information on the weights can then be analyzed to describe the possible decision outcomes. In the SMAA methodology

☆This manuscript was processed by Associate Editor Yeh.

* Corresponding author.
E-mail addresses: g.h.m.van.valkenhoef@rug.nl (G. van Valkenhoef), tommi.tervonen@evidera.com (T. Tervonen).

http://dx.doi.org/10.1016/j.omega.2015.10.014
0305-0483/

[21,20,31] such analyses result in rank probabilities for each decision alternative. These are also known as rank acceptabilities and they are usually estimated through Monte Carlo simulation [32]. The SMAA approach has been successfully applied for decision support in various domains (e.g. [11,20,34,35,38]).

The usefulness of the simulation approach is related to the capability of the rank probabilities to discriminate the decision alternatives, which in turn depends on the linear weight constraints resulting from the answers to the pair-wise questions. The number of questions the DM needs to answer should be as small as possible to minimize the DM's cognitive load. This paper develops an entropy-based method for optimizing information gain from pair-wise holistic elicitation questions.

1.1. Background

Choosing the most appropriate pair of actual or fictitious reference alternatives for indirect preference elicitation is not a trivial task, and different approaches have been developed for tackling the problem. Rios-Insua and Mateos [28] defined conditions for reducing the efficient discrete set of alternatives after observing a single pair-wise choice, but they did not consider further which pair of alternatives to use in the next elicitation iteration. The question-response process and the optimal choice sequence (policy) was analyzed by Holloway and White [12]. They assumed deterministic attribute measurements, a finite set of alternatives, and linear partial utility functions. Iyengar et al. [13] presented a heuristic for choosing the pair-wise question to ask so that the resulting weight space is cut approximately in half. However, none of these works took into account uncertain attribute measurements or the information theoretic basis of the decision; such an entropy-based method for problems with discrete outcome sets was introduced by Abbas [1].

Entropy quantifies the amount of uncertainty over a set of events [6]. In multi-attribute choice problems the events are the different alternatives that can obtain the first rank, that is, there are $|I|$ possible events. In ranking problems the events are the different possible rankings over the set of alternatives. Entropy allows us to compare the level of uncertainty associated with different imprecise preferences, and therefore select the elicitation question that maximizes the information gain. Entropy has been used in a related context for constructing a full joint outcome distribution when the marginal probabilities and pair-wise correlations are elicited from experts [3]. The objective of using entropy in specifying the outcome distribution is different: rather than eliciting preferences to minimize entropy so that the decision recommendation becomes more certain, it is used to select the maximum entropy (most uncertain) distribution compatible with the given information. We do not consider the construction of the full joint outcome distribution in this paper, and emphasize that our use of entropy as a metric of overall decision uncertainty is quite different from the maximum entropy approach to the selection of probability distributions.

1.2. Contribution

Our contribution in this paper is two-fold. First, we develop an entropy-based framework for reasoning about holistic preference elicitation questions in multi-attribute ranking problems with continuous measurement distributions. We consider additive utility models where the partial utility functions, apart from the scaling factors (weights) are elicited beforehand. We extend the previous work of Abbas [1] by considering a discrete set of decision alternatives with uncertain attribute measurements modeled through a joint probability distribution. Thus, the set of possible outcomes is uncountable as we do not integrate uncertainty into a utility measure, but the actual set of decision alternatives is finite. We believe that the current work is the first to consider this problem setting, and therefore the previously mentioned approaches are not applicable. Neither are the various questioning methods developed for interactive multi-objective optimization (see e.g. [30]). The second contribution is the development of a greedy technique that chooses the myopically optimal question in each elicitation iteration. We numerically investigate the procedure's performance with artificial problems, and present an application of the technique on benefit-risk analysis of antithrombolytic drugs.

The remainder of this paper is structured as follows. In Section 2, we describe the theory of entropy optimal weight constraint elicitation and the specific case of pair-wise questions. Extensions that consider restricted sets of reference alternatives and a stopping criterion are discussed in Section 3. The computational experiments and their results are presented in Section 4, and an example is analyzed in Section 5. Section 6 ends the paper with a discussion.

2. Entropy-optimal weight elicitation

In this section, we first outline our general framework for entropy optimal weight constraint elicitation. We then show how the relevant quantities can be approximated using Monte Carlo simulation, and how the framework can be applied to find approximately entropy-optimal pair-wise elicitation questions. In what follows, we will refer to riskless utility (value) when talking about utility, although the developed theory mostly also applies to true utility functions that model the DM's risk attitude. Incomplete information on the weights is represented through a density $f_W(w)$ that is non-zero only within the $(n-1)$-simplex

$$W_n = \left\{ w \in \mathbb{R}^n \,|\, w \geq 0, \sum_{j \in J} w_j = 1 \right\}. \tag{2}$$

The ranges of u_j in (1) are defined through hypothetical alternatives x^{worst} and x^{best} that have all the measurements at the worst and the best levels in x, respectively, and thus $\forall j \in J : u_j(x_j^{\text{worst}}) = 0, u_j(x_j^{\text{best}}) = 1$. In case x is not bounded, the worst and best values can be chosen e.g. as hull of the 95% confidence intervals [35]. We consider the case of a single DM, and model incomplete information about her preferences with a density $f_W(w)$ that is uniform within a *feasible weight space* W' formed by restricting W_n with linear constraints.

Starting from the current feasible weight space $W' \subseteq W_n$, an elicitation question Q with an answer set $A(Q)$ will reduce the feasible weight space to a sub-region $W'' \in A(Q)$. For example, when $n=2$, the first question could be which one of two alternatives with partial utility vectors $u^1 = (1,0)$, $u^2 = (0,1)$ the DM (weakly) prefers. Possible answers to this question are $A(Q) = \{\{w \in W_2 \,|\, w_1 \geq w_2\}, \{w \in W_2 \,|\, w_2 \geq w_1\}\}$. In general, the answer set will be complete ($\bigcup_{W'' \in A(Q)} W'' = W'$), but the answers do not need to be disjoint. The answer set could be uncountable, e.g. when the DM is asked to pick a single number from a real valued scale (contingent valuation), but we consider only countable answer sets in this paper (conjoint valuation).

Depending on the problem structure, the choice of the elicitation question can have a large impact on how well the decision alternatives are discriminated with the answer. The question should maximize the expected information gain, which we define by taking into account the objective of applying the decision model, e.g. to rank the alternatives. Fig. 1 illustrates this in a simple problem with two alternatives A and B and two attributes. The partial utilities are distributed as $u_1(x_1^A) \sim \mathcal{U}(0, 0.3)$, $u_2(x_2^A) \sim \mathcal{U}(0.7, 1)$, $u_1(x_1^B) \sim \mathcal{U}(0.3, 0.6)$,