

Imprecise Swing Weighting for Multi-Attribute Utility Elicitation Based on Partial Preferences

Matthias C. M. Troffaes

Durham University (United Kingdom)

MATTHIAS.TROFFAES@DURHAM.AC.UK

Ullrika Sahlin

Lund University (Sweden)

ULLRIKA.SAHLIN@CEC.LU.SE

Abstract

We describe a novel approach to multi-attribute utility elicitation which is both general enough to cover a wide range of problems, whilst at the same time simple enough to admit reasonably straightforward calculations. We allow both utilities and probabilities to be only partially specified, through bounding. We still assume marginal utilities to be precise. We derive necessary and sufficient conditions under which our elicitation procedure is consistent. As a special case, we obtain an imprecise generalization of the well known swing weighting method for eliciting multi-attribute utility functions. An example from ecological risk assessment demonstrates our method.

Keywords: utility; partial preference; consistency; uniqueness; multi-attribute; elicitation; imprecise; robust; swing weighting.

1. Introduction

In many decision problems where outcomes feature multiple attributes, additive multi-attribute utility functions are a popular choice due to their simplicity (Clemen and Reilly, 2001). They split the joint utility function into a weighted sum of marginal utility functions. Elicitation of the joint can then be split into two elicitation procedures: one for each of the marginals, and one for the weights.

A reoccurring issue is the precision of the attribute weights. Indeed, whilst marginal utility functions on separate attributes are often quite easily elicited, the way in which these attributes should be weighed against each other is much harder to quantify precisely. Such decision problems appear in different applications (Yemshanov et al., 2013; Hermerén et al., 2014). So, even if an additive form can be assumed, the weights themselves might still be subject to imprecision due to incomplete preferences between multi-attribute lotteries. Hermerén et al. (2014) suggest different types of value uncertainty and conclude that more work is needed to understand the cause of this uncertainty, in order to understand how to treat it. The only applied example we identified is Yemshanov et al. (2013), who considered uncertainty in weights by a multidimensional efficiency frontier analysis, which treat each attribute separately. Here we are interested in the elicitation of these weights.

Swing weighting (von Winterfeldt and Edwards, 1986) is a simple and popular method for eliciting the weights of an additive multi-attribute utility function. Unfortunately, the standard treatment of swing weighting uses ‘scores’ which are not usually directly interpreted in terms of preferences over lotteries, giving it the impression of a heuristic rather than a normative method. Moreover, swing weighting forces completeness of preferences between multi-attribute lotteries.

In this paper, we generalise swing weighting so bounds on the weights of the joint utility function can be elicited normatively. This is an important step to further widen the applicability of utility theory in problems where the consequences of decisions have multiple aspects that cannot be easily

weighed against each other. As an extra bonus, we also derive a normative interpretation of the standard swing weighting procedure. The resulting problems, when both probabilities and utilities are allowed to be imprecise, require quadratic programming, for which standard algorithms exist.

We are of course aware that decision theory has been generalised to deal with arbitrary partial preferences in their full generality (Seidenfeld et al., 1995). However, such theories can be technically difficult to work with due to the fact that they lead to non-convex sets of utilities and probabilities. Various special cases have been studied that do allow convex analysis to be used for elicitation, modelling, and inference (Williams, 1975, 2007; Levi, 1980; Walley, 1991). These works do not explicitly try to deal with multiple attributes. The contribution of this paper can be seen as a practical approach towards multi-attribute decision problems where marginal utilities are still precise, but where we wish to be a bit more cautious about modelling preferences across attributes. It can be seen as a simple generalisation of Walley (1991) to the multi-attribute case.

The idea of generalising swing weighting to allow for partial preferences is not new either; see for instance Mustajoki et al. (2005); Riabacke et al. (2009); Gomes et al. (2011); Riabacke et al. (2012); Danielson et al. (2014) and references therein. Those works generally focus on reducing the cognitive requirements on decision makers, and propose specific models for eliciting attribute weights, but without relating the elicitation directly to preferences between multi-attribute gambles. Instead, in this paper, we develop a general mathematical framework for eliciting attribute weights in a directly operational way through preferences between multi-attribute gambles. We thereby generalise the interval swing weighting method proposed by Mustajoki et al. (2005) (at least in the cases where the reference attribute is either the worst or the best attribute). The theory that we develop can be adapted to a wide range of situations, and possibly could accommodate cognitive limitations in a more flexible way, although we will not fully explore this in this paper.

The paper is structured as follows. Section 2 introduces the notation and explains the assumptions that we make throughout the paper. Section 3 briefly describes how marginal utility functions can be elicited, and serves as an introduction to the idea of utility elicitation. Section 4 reviews the standard swing weighting procedure, and provides a simple normative interpretation of swing weighting in terms of lotteries. Section 5 generalises the swing weighting procedure to allow imprecise weights, and section 6 identifies necessary and sufficient conditions for this elicitation procedure to be consistent. Section 7 provides a fully worked example of our method, using an example from ecological risk assessment. Section 8 concludes the paper.

2. Notation and Assumptions

Let $\mathcal{R} := \mathcal{A}_1 \times \dots \times \mathcal{A}_n$ be a finite set of rewards, each reward $r = (a_1, \dots, a_n)$ comprising of n attributes. A *lottery* ℓ on \mathcal{R} is a probability mass function over \mathcal{R} , and is interpreted as a random reward with precisely known probabilities. The set of all lotteries over \mathcal{R} is denoted by $L(\mathcal{R})$.

Note that at this stage, we are not yet interested in modelling uncertainty. Rather, we will use lotteries in order to elicit a subject's attitudes towards rewards. For modelling uncertainty, one might consider *horse lotteries*, which for our purpose would be mappings from some finite possibility space Ω to $L(\mathcal{R})$. This follows the traditional approach (Anscombe and Aumann, 1963; Seidenfeld et al., 1995). In this paper, we do not consider horse lotteries, and focus purely on the utility aspect. That said, in section 7, we will demonstrate how uncertainty can be incorporated in an example.

So, we wish to model our preferences between lotteries over our multi-attribute rewards. A *utility function* on \mathcal{R} is any function $U: \mathcal{R} \rightarrow \mathbb{R}$. We lift U to $L(\mathcal{R})$ in the usual way:

$$U(\ell) := \sum_{r \in \mathcal{R}} \ell(r)U(r). \quad (1)$$

Note that U satisfies

$$U(\alpha\ell_1 + (1 - \alpha)\ell_2) = \alpha U(\ell_1) + (1 - \alpha)U(\ell_2) \quad (2)$$

for all $\alpha \in [0, 1]$. The standard approach assumes that our preferences form a complete preorder \succeq on $L(\mathcal{R})$ and can be represented through a utility function U , where

$$\ell_1 \succeq \ell_2 \iff U(\ell_1) \geq U(\ell_2) \quad (3)$$

for all ℓ_1 and $\ell_2 \in L(\mathcal{R})$. This representation can be directly motivated from some simple assumptions on \succeq (Herstein and Milnor, 1953).

However, in many applications, preferences between rewards are inherently incomplete, in the sense that there may be lotteries between which we cannot state any preference. We will assume that our preferences form a preorder \succeq on $L(\mathcal{R})$ (so we drop completeness), and can be represented through a *set* \mathcal{U} of utility functions $U: L(\mathcal{R}) \rightarrow \mathbb{R}$:

$$\ell_1 \succeq \ell_2 \iff \forall U \in \mathcal{U}: U(\ell_1) \geq U(\ell_2) \quad (4)$$

for all ℓ_1 and $\ell_2 \in L(\mathcal{R})$. Elicitation is then concerned with finding a procedure for identifying \mathcal{U} .

In cases where rewards are comprised of multiple attributes, in standard utility theory, it is customary to split the elicitation problem into two parts:

1. Elicit *marginal utility functions* $U_i: \mathcal{A}_i \rightarrow \mathbb{R}$ for each $i \in \{1, \dots, n\}$.
2. Assume that the joint utility function can be written as a particular function of the marginal utility functions, and elicit the parameters of that function.

The simplest of these joint forms is the *additive form*:

$$U(a_1, \dots, a_n) = \sum_{i=1}^n k_i U_i(a_i) \quad (5)$$

Again, this form can be directly motivated from some simple assumptions on \succeq (Keeney and Raiffa, 1993). Although those assumptions are quite restrictive and are easily criticised, the simplicity of the additive form, having only n parameters, make it one of the most commonly used models for multi-attribute utility in practical applications.

3. Elicitation of Marginal Utility

To introduce the idea of utility elicitation, and for the sake of completeness, we mention a simple standard method for eliciting the marginal utility functions U_i (Clemen and Reilly, 2001):

1. Identify a worst reward \underline{a}_i and a best reward \bar{a}_i in \mathcal{A}_i .
2. For every other reward a_i in \mathcal{A}_i , find $\alpha(a_i)$ so that the subject is indifferent between (i) getting a_i with certainty and (ii) getting \underline{a}_i with probability $1 - \alpha(a_i)$ or \bar{a}_i with probability $\alpha(a_i)$:

$$a_i \simeq (1 - \alpha(a_i))\underline{a}_i \oplus \alpha(a_i)\bar{a}_i \quad (6)$$

where \oplus denotes the combination of rewards into lotteries, so $(1 - \alpha)r_1 \oplus \alpha r_2$ is the lottery ℓ with $\ell(r_1) = 1 - \alpha$, $\ell(r_2) = \alpha$, and $\ell(r) = 0$ for all other rewards. We also denoted indifference by \simeq : $\ell_1 \simeq \ell_2 \iff (\ell_1 \succeq \ell_2 \text{ and } \ell_2 \succeq \ell_1)$.

3. Set $U_i(\underline{a}_i) := 0$, $U_i(\bar{a}_i) := 1$, and $U_i(a_i) := \alpha(a_i)$ for every other reward a_i in \mathcal{A}_i .

Naturally, an interesting question relates to how we can relax this elicitation procedure to allow for incomplete preferences in the marginals. As we shall see in section 5, allowing incompleteness in both marginals and in the weights introduces non-linear constraints. So, for practical reasons, in this paper, we only investigate incompleteness in the weights, and assume marginals to be fully precise.

4. Swing Weighting

For eliciting the weights k_i in the joint utility function of eq. (5), various methods exist, but a simple and effective method is *swing weighting* (von Winterfeldt and Edwards, 1986):

1. Score the following $n + 1$ rewards:

reward	score
$r_0 := (\underline{a}_1, \underline{a}_2, \dots, \underline{a}_n)$	0
$r_1 := (\bar{a}_1, \underline{a}_2, \dots, \underline{a}_n)$	s_1
$r_2 := (\underline{a}_1, \bar{a}_2, \dots, \underline{a}_n)$	s_2
\vdots	\vdots
$r_n := (\underline{a}_1, \underline{a}_2, \dots, \bar{a}_n)$	s_n

where the worst score is 0 (always assigned to the worst reward), the best score is 100, and the other scores indicate the “relative improvement” from the worst reward.

2. Set

$$k_i := \frac{s_i}{\sum_{i=1}^n s_i}. \quad (7)$$

Note that this formula hinges on the assumption that all marginal utility functions are renormalized to the $[0, 1]$ interval—this is the case if we use the marginal method as described earlier.

Although we find swing weighting a straightforward and mathematically elegant method for eliciting the weights, what is missing is an interpretation directly in terms of preferences over lotteries. In fact, this is very easy to do, but much to our surprise it is not mentioned anywhere in the literature as far as we could find:

1. Consider again the rewards r_0, \dots, r_n as constructed above. Clearly r_0 is the worst reward.
2. Identify the best of these rewards. Without loss of generality, we may assume that this is r_n (we can always permute the order of the attributes if need be).
3. For all $i \in \{1, \dots, n\}$, find α_i such that

$$r_i \simeq (1 - \alpha_i)r_0 \oplus \alpha_i r_n. \quad (8)$$

(Note that $\alpha_n = 1$.)

4. Set

$$k_i := \frac{\alpha_i}{\sum_{i=1}^n \alpha_i} \quad (9)$$

It is easy to see that this choice of k_i is the only choice that is compatible with eq. (8). So, we can interpret the swing weighting scores directly in terms of probabilities, which we find more appealing. This also puts the method on a firm normative basis.

5. Imprecise Swing Weighting

A common criticism against the swing weighting method (and, in fact, also against the marginal method that we presented) is that all lotteries considered in the elicitation involve extremes only. We therefore adapt the swing weighting method to allow for more flexible comparisons, not just focusing on extremes.

Remember that we are dropping the completeness assumption, and therefore that we are interested in identifying a set \mathcal{U} of utility functions, rather than a single utility function. To do so, we will view all weights k_i as parameters (so we have n parameters), and we will represent \mathcal{U} through a collection of constraints on these parameters:

1. Consider *any* joint rewards r_0, \dots, r_n such that for all $j \in \{1, \dots, n-1\}$ we have that

$$r_0 \preceq r_j \preceq r_n \quad (10)$$

2. For all $j \in \{1, \dots, n-1\}$, find the largest $\underline{\alpha}_j$ and smallest $\bar{\alpha}_j$ such that

$$(1 - \underline{\alpha}_j)r_0 \oplus \underline{\alpha}_j r_n \preceq r_j \preceq (1 - \bar{\alpha}_j)r_0 \oplus \bar{\alpha}_j r_n \quad (11)$$

3. Let u_j denote the vector of marginal utilities for r_j , i.e. if $r_j = (a_1, \dots, a_n)$ then $u_j = (U_1(a_1), \dots, U_n(a_n))$. Let k denote the vector (k_1, \dots, k_n) . With this notation, impose

$$\forall j \in \{1, \dots, n-1\}: \quad (u_j - (1 - \underline{\alpha}_j)u_0 - \underline{\alpha}_j u_n) \cdot k \geq 0 \quad (12a)$$

$$\forall j \in \{1, \dots, n-1\}: \quad (u_j - (1 - \bar{\alpha}_j)u_0 - \bar{\alpha}_j u_n) \cdot k \leq 0 \quad (12b)$$

$$1 \cdot k = 1 \quad (12c)$$

The last constraint is simply another way of writing $\sum_{i=1}^n k_i = 1$, and fixes the multiplicative scaling of the joint utility function.

To see that the other two constraints indeed represent the elicited preferences, note that eq. (11) is equivalent to

$$(1 - \underline{\alpha}_j)U(r_0) + \underline{\alpha}_j U(r_n) \leq U(r_j) \leq (1 - \bar{\alpha}_j)U(r_0) + \bar{\alpha}_j U(r_n) \quad (13)$$

and note that $U(r_j) = u_j \cdot k$.

These inequalities are quadratic in the marginal utilities and in the weights. However, if the marginal utilities are precise, then we have a simple set of linear constraints on the weights k_j .

Naturally, we also recover swing weighting as a special case. In the imprecise case however it is important to realise that we cannot always take the rewards as in the standard swing weighting method. We already argued that this might be a bad idea due to the focus on extremes, however it may also cause a problem because the method requires that there is a single best attribute—we may not have such best attribute if we allow for incompleteness.

6. Consistency and Uniqueness

The procedure that we described works for any choice of rewards r_j . Naturally, a good choice of rewards r_j should ensure that the constraints obtained admit a solution for all possible choices of $0 \leq \underline{\alpha}_j \leq \bar{\alpha}_j \leq 1$. In fact, we also would like this solution to be unique in the precise case (i.e. when $\underline{\alpha}_j = \bar{\alpha}_j$ for all j), so that we can at least in principle allow a complete elicitation of preferences if possible. Both of these desiderata are satisfied if:

- (i) $u_0 \leq u_n$, and
- (ii) the system

$$\forall j \in \{1, \dots, n-1\}: \quad (u_j - (1 - \alpha_j)u_0 - \alpha_j u_n) \cdot k = 0 \quad (14a)$$

$$1 \cdot k = 1 \quad (14b)$$

has a unique solution, regardless our choice of $\alpha_1, \dots, \alpha_{n-1} \in [0, 1]$.

Note that $u_0 \leq u_n$ guarantees that $(u_j - (1 - \alpha_j)u_0 - \alpha_j u_n)$ is a decreasing function of α_j , so in this case it is ensured that, say, if we solve eqs. (12b) and (12c) with equalities everywhere, then the inequality in eq. (12a) is automatically satisfied; in other words, eq. (12) is consistent.

We will henceforth assume that $u_0 \leq u_n$, and focus on the uniqueness of the solution of eq. (14).

Theorem 1 Consider any $\alpha_1, \dots, \alpha_{n-1} \in [0, 1]$. If the matrix

$$\begin{bmatrix} u_1 - (1 - \alpha_1)u_0 - \alpha_1 u_n \\ u_2 - (1 - \alpha_2)u_0 - \alpha_2 u_n \\ \vdots \\ u_{n-1} - (1 - \alpha_{n-1})u_0 - \alpha_{n-1} u_n \\ 1 \end{bmatrix} \quad (15)$$

has full rank, then eq. (14) has a unique solution.

The following theorem provides much quicker check for uniqueness, in case u_0 is constant (note that u_0 being constant is a standard feature of the usual swing weighting procedure).

Theorem 2 Consider any $\alpha_1, \dots, \alpha_{n-1} \in [0, 1]$. Assume that u_0 is constant, and that the vectors $(u_1, \dots, u_{n-1}, 1)$ are linearly independent. Let λ_j be the coefficients that decompose u_n as a linear combination of $(u_1, \dots, u_{n-1}, 1)$, i.e.

$$u_n = \lambda_n + \sum_{j=1}^{n-1} \lambda_j u_j \quad (16)$$

Then eq. (14) has a unique solution if and only if

$$\sum_{j=1}^{n-1} \alpha_j \lambda_j \neq 1 \quad (17)$$

In particular, when $\lambda_1 \leq 0, \dots, \lambda_{n-1} \leq 0$, then eq. (14) has a unique solution, regardless our choice of $\alpha_1, \dots, \alpha_{n-1} \in [0, 1]$.

Proof We need to show that the matrix

$$\begin{bmatrix} u_1 - (1 - \alpha_1)u_0 - \alpha_1 u_n \\ u_2 - (1 - \alpha_2)u_0 - \alpha_2 u_n \\ \vdots \\ u_{n-1} - (1 - \alpha_{n-1})u_0 - \alpha_{n-1} u_n \\ 1 \end{bmatrix} \quad (18)$$

has full rank. Because u_0 is constant, and so is the final row, this matrix has full rank if and only if

$$\begin{bmatrix} u_1 - \alpha_1 u_n \\ u_2 - \alpha_2 u_n \\ \vdots \\ u_{n-1} - \alpha_{n-1} u_n \\ 1 \end{bmatrix} \quad (19)$$

has full rank.

Because the $(u_1, \dots, u_{n-1}, 1)$ are linearly independent, we can write u_n as a linear combination of these vectors:

$$u_n = \lambda_n + \sum_{j=1}^{n-1} \lambda_j u_j \quad (20)$$

So, our matrix can be written as

$$\begin{aligned} \begin{bmatrix} u_1 - \alpha_1 u_n \\ u_2 - \alpha_2 u_n \\ \vdots \\ u_{n-1} - \alpha_{n-1} u_n \\ 1 \end{bmatrix} &= \begin{bmatrix} u_1 - \alpha_1 \left(\lambda_n + \sum_{j=1}^{n-1} \lambda_j u_j \right) \\ u_2 - \alpha_2 \left(\lambda_n + \sum_{j=1}^{n-1} \lambda_j u_j \right) \\ \vdots \\ u_{n-1} - \alpha_{n-1} \left(\lambda_n + \sum_{j=1}^{n-1} \lambda_j u_j \right) \\ 1 \end{bmatrix} \\ &= \begin{bmatrix} 1 - \alpha_1 \lambda_1 & -\alpha_1 \lambda_2 & \dots & -\alpha_1 \lambda_{n-1} & -\alpha_1 \lambda_n \\ -\alpha_2 \lambda_1 & 1 - \alpha_2 \lambda_2 & \dots & -\alpha_2 \lambda_{n-1} & -\alpha_2 \lambda_n \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ -\alpha_{n-1} \lambda_1 & -\alpha_{n-1} \lambda_2 & \dots & 1 - \alpha_{n-1} \lambda_{n-1} & -\alpha_{n-1} \lambda_n \\ 0 & 0 & \dots & 0 & 1 \end{bmatrix} \begin{bmatrix} u_1 \\ u_2 \\ \vdots \\ u_{n-1} \\ 1 \end{bmatrix} \end{aligned} \quad (21)$$

which has full rank if both matrices on the right hand side have full rank. The second matrix has full rank by assumption. The first matrix has full rank if and only if

$$\begin{bmatrix} 1 - \alpha_1 \lambda_1 & -\alpha_1 \lambda_2 & \dots & -\alpha_1 \lambda_{n-1} \\ -\alpha_2 \lambda_1 & 1 - \alpha_2 \lambda_2 & \dots & -\alpha_2 \lambda_{n-1} \\ \vdots & \vdots & \ddots & \vdots \\ -\alpha_{n-1} \lambda_1 & -\alpha_{n-1} \lambda_2 & \dots & 1 - \alpha_{n-1} \lambda_{n-1} \end{bmatrix} \quad (23)$$

has full rank. This matrix can be written as

$$I - \alpha \lambda^T \quad (24)$$

where $\alpha = (\alpha_1, \dots, \alpha_{n-1})$ and $\lambda = (\lambda_1, \dots, \lambda_{n-1})$. This has full rank if and only if its determinant is non-zero. We now use Sylvester's determinant identity:

$$\det(I - \alpha \lambda^T) = \det(1 - \lambda^T \alpha) = 1 - \sum_{j=1}^{n-1} \alpha_j \lambda_j \quad (25)$$

We arrive at the desired result. ■

This theorem applies for instance if we use the joint rewards as in standard swing weighting:

reward	u_j
$r_0 := (\underline{a}_1, \underline{a}_2, \dots, \underline{a}_n)$	$(0, 0, \dots, 0)$
$r_1 := (\bar{a}_1, \underline{a}_2, \dots, \underline{a}_n)$	$(1, 0, \dots, 0)$
$r_2 := (\underline{a}_1, \bar{a}_2, \dots, \underline{a}_n)$	$(0, 1, \dots, 0)$
\vdots	\vdots
$r_n := (\underline{a}_1, \underline{a}_2, \dots, \bar{a}_n)$	$(0, 0, \dots, 1)$

Note that $u_0 \leq u_n$, as required. Also, u_0 is constant (zero), and $(u_1, \dots, u_{n-1}, 1)$ are linearly independent: the theorem applies. Because

$$u_n = 1 - \sum_{j=1}^{n-1} u_j, \quad (26)$$

it follows that $\lambda_j = -1$ for all $j \in \{1, \dots, n-1\}$. The condition for uniqueness is satisfied.

We can also consider the case where u_n is constant:

Theorem 3 Consider any $\alpha_1, \dots, \alpha_{n-1} \in [0, 1]$. Assume that u_n is constant, and that the vectors $(u_1, \dots, u_{n-1}, 1)$ are linearly independent. Let λ_j be the coefficients that decompose u_0 as a linear combination of $(u_1, \dots, u_{n-1}, 1)$, i.e.

$$u_0 = \lambda_n + \sum_{j=1}^{n-1} \lambda_j u_j \quad (27)$$

Then eq. (14) has a unique solution if and only if

$$\sum_{j=1}^{n-1} (1 - \alpha_j) \lambda_j \neq 1 \quad (28)$$

In particular, when $\lambda_1 \leq 0, \dots, \lambda_{n-1} \leq 0$, then eq. (14) has a unique solution, regardless our choice of $\alpha_1, \dots, \alpha_{n-1} \in [0, 1]$.

The proof is almost identical to the proof of the previous theorem, and hence is left as an exercise to the reader. We will use this variant in the example below.

7. Example

We now provide a fully worked example to see the theory at work. In addition to imprecise utilities, we will also admit imprecise probabilities.

Following [Bohman and Edsman \(2013\)](#), we are interested in an ecological management decision, namely the eradication of an invasive species that has been observed in a water system. The following management decisions were identified:

- I Do nothing.
- II Mechanical removal.
- III Drain the system on water and remove of individuals by hand.
- IV Drain the system of water, dredge and sieve the masses to identify and remove individuals.
- V Use a decomposable biocide in combination with drainage to increase the biocide concentration.
- VI Increase pH in combination with drainage and removal by hand.

The decision comprises several attributes. Each decision was scored according to attributes identified as relevant by a group of experts. For each of these attributes, a discrete scale ranging from 1 to 4 was constructed, where 1 corresponds to the worst outcome, and 4 corresponds to the best outcome. We will interpret these scores as marginal utility functions.

Besides the attributes, the experts also bounded the probability that the method is successful in eradication. Note that we are using hypothetical values here, these values are not actual expert judgements, and only serve to demonstrate the methodology. Note also that in the actual problem, there is considerable uncertainty about whether the invasive species is present at all. For simplicity, in this example, we assume that the alien species is present with certainty.

The following table lists all attributes considered, as well as the interpretation of the scores for each attribute and for each management decision, and the expert assessments for the attribute scores, in case of success:

Attribute	Worst (score 1)	Best (score 4)	Decision d					
			I	II	III	IV	V	VI
Biotic impact	High	Low	4	4	3	3	2	1
Longevity of impacts	Long	Short	4	4	3	3	1	2
Experience	Little	High	4	3	1	4	1	1
Feasibility	Difficult	Easy	4	4	2	3	1	2
Cost	High	Low	4	4	3	1	2	3

In case of failure to eradicate the invasive species, the scores for biotic impact and longevity of impacts drop to their worst values:

Attribute	Worst (score 1)	Best (score 4)	Decision d					
			I	II	III	IV	V	VI
Biotic impact	High	Low	1	1	1	1	1	1
Longevity of impacts	Long	Short	1	1	1	1	1	1
Experience	Little	High	4	3	1	4	1	1
Feasibility	Difficult	Easy	4	4	2	3	1	2
Cost	High	Low	4	4	3	1	2	3

Bounds on the probability of successful eradication of the species under the different management decisions are:

Probability	Decision d					
	I	II	III	IV	V	VI
\underline{p}_d	0	0.05	0.3	0.4	1.0	0.7
\overline{p}_d	0	0.25	0.5	0.7	1.0	0.8

The joint expected utility of decision d can be written as:

$$\sum_{j=1}^n k_j (\theta U_{1j}(a_{jd}) + (1 - \theta)U_{2j}(a_{jd})) \tag{29}$$

where U_{1j} are the marginal utilities as listed in the first table, and U_{2j} are the marginal utilities as listed in the second table (both after rescaling to 0–1).

Because the decision affects the probability of successful management (i.e. we have act-state dependence), we will treat the problem using interval dominance.

For eliciting the weights k_j of the joint utility function, we will use a variant of swing weighting, and we will consider the following joint rewards (directly expressed in terms of marginal utilities, rescaled to 0–1):

rewards
$u_0 := (2/3, 1, 1, 1, 1)$
$u_1 := (1, 2/3, 1, 1, 1)$
$u_2 := (1, 1, 2/3, 1, 1)$
$u_3 := (1, 1, 1, 2/3, 1)$
$u_4 := (1, 1, 1, 1, 2/3)$
$u_5 := (1, 1, 1, 1, 1)$

These rewards are more natural from an ecological risk perspective compared to the rewards considered by the original swing weighting method: they consider only small changes from a normal state, instead of extremes, and are thus easier to compare (regardless of any imprecision in preferences).

Note that $u_0 \leq u_5$ as required for consistency. Also note that u_5 is constant, so we can apply theorem 3. We see that

$$u_0 = 14/3 - \sum_{j=1}^4 u_j. \quad (30)$$

Consequently all $\lambda_j = -1$ in theorem 3, and so the condition for uniqueness is always satisfied.

We consider biotic impact to be the most important attribute, so clearly we have that $r_0 \preceq r_j \preceq r_5$ for all $j \in \{0, \dots, 5\}$. We also assess that

$$0.8r_0 \oplus 0.2r_5 \preceq r_1 \preceq 0.7r_0 \oplus 0.3r_5 \quad (31)$$

$$0.5r_0 \oplus 0.5r_5 \preceq r_2 \preceq 0.4r_0 \oplus 0.6r_5 \quad (32)$$

$$0.3r_0 \oplus 0.7r_5 \preceq r_3 \preceq 0.1r_0 \oplus 0.9r_5 \quad (33)$$

$$0.2r_0 \oplus 0.8r_5 \preceq r_4 \preceq 0.1r_0 \oplus 0.9r_5 \quad (34)$$

With these assessments, eq. (12) becomes

$$((1, 2/3, 1, 1, 1) - 0.8(2/3, 1, 1, 1, 1) - 0.2) \cdot k \geq 0 \quad (35a)$$

$$((1, 1, 2/3, 1, 1) - 0.5(2/3, 1, 1, 1, 1) - 0.5) \cdot k \geq 0 \quad (35b)$$

$$((1, 1, 1, 2/3, 1) - 0.3(2/3, 1, 1, 1, 1) - 0.7) \cdot k \geq 0 \quad (35c)$$

$$((1, 1, 1, 1, 2/3) - 0.2(2/3, 1, 1, 1, 1) - 0.8) \cdot k \geq 0 \quad (35d)$$

$$((1, 2/3, 1, 1, 1) - 0.7(2/3, 1, 1, 1, 1) - 0.3) \cdot k \leq 0 \quad (35e)$$

$$((1, 1, 2/3, 1, 1) - 0.4(2/3, 1, 1, 1, 1) - 0.6) \cdot k \leq 0 \quad (35f)$$

$$((1, 1, 1, 2/3, 1) - 0.1(2/3, 1, 1, 1, 1) - 0.9) \cdot k \leq 0 \quad (35g)$$

$$((1, 1, 1, 1, 2/3) - 0.1(2/3, 1, 1, 1, 1) - 0.9) \cdot k \leq 0 \quad (35h)$$

$$1 \cdot k = 1 \quad (35i)$$

So, for each decision, we need to minimize and maximize the joint utility expressed in eq. (29), subject to the above constraints and subject to $\underline{p}_d \leq \theta \leq \bar{p}_d$. The constraints are all linear, and the objective function is quadratic, hence this is a quadratic programming problem. Because θ itself only appears linearly and is constrained separately, it suffices to consider only the extreme values for θ . Consequently, for each decision, we must merely solve two linear programs: one for $\theta = \underline{p}_d$ and one for $\theta = \bar{p}_d$.

Using `scipy` (Jones et al., 2001–), we find the following bounds:

Decision	Lower Utility	Upper Utility
I	0.25	0.37
II	0.23	0.47
III	0.18	0.31
IV	0.38	0.57
V	0.14	0.17
VI	0.11	0.17

Options I, III, V, and VI are dominated by option IV so should not be considered. Either option II (mechanical removal) or IV (drain the system of water, dredge and sieve), could be considered.

For the sake of completeness, we also present the bounds on the attribute weights, resulting from eq. (35):

Attribute	Lower Weight	Upper Weight
Biotic impact	0.36	0.43
Longevity of impacts	0.26	0.33
Experience	0.15	0.21
Feasibility	0.04	0.12
Cost	0.04	0.08

8. Conclusions

In this paper, we provided an imprecise generalisation of the swing weighting method for eliciting multi-attribute utility functions. The proposed method enables us to cover a wider range of problems where preference can only be partially specified, whilst at the same time still admitting straightforward calculations. We studied the consistency of the elicitation procedure, and found simple conditions under which consistency is always guaranteed. We demonstrated our method using a real example concerning the management of an invasive species featuring substantial uncertainty in the management outcomes and ambiguity in the preferences over different impacts. In this example, we allowed both utilities and probabilities to be only partially specified, through bounding.

We do note that our approach is still limited in that we will assume that all marginal utility functions are precise. Relaxing this is possible but leads to fully non-linear optimisation, and more work is needed to identify whether such treatment can be feasible in practice. Naturally, another limitation is that we only discussed *additive* multi-attribute utility functions.

Another open end is that we have assumed that our preferences over horse lotteries are representable by a convex set of weights along with a convex set of probability mass functions. Whilst such representation is appealing mathematically (inference becomes a quadratic programming problem), it would be interesting to have an axiomatic treatment from first principles (as in [Seidenfeld et al. \(1995\)](#)) identifying the conditions under which such treatment is feasible.

Acknowledgements

U. Sahlin has been funded by the Swedish research council FORMAS project number 219-2013-1271. The research presented in this paper is a contribution to the strategic research area Biodiversity and Ecosystems in a Changing Climate, BECC, funded by the Swedish government.

References

- F. J. Anscombe and R. J. Aumann. A definition of subjective probability. *Annals of Mathematical Statistics*, 34(1):199–205, Mar. 1963.
- P. Bohman and L. Edsman. Marmorkräftan i Märstaån. Riskanalys och åtgärdsförslag. Aqua Reports 2013:17, Sveriges lantbruksuniversitet, Drottningholm, 2013.
- R. T. Clemen and T. Reilly. *Making Hard Decisions*. Duxbury, 2001.
- M. Danielson, L. Ekenberg, A. Larsson, and M. Riabacke. Weighting under ambiguous preferences and imprecise differences in a cardinal rank ordering process. *International Journal of Computational Intelligence Systems*, 7:105–112, 2014. doi:[10.1080/18756891.2014.853954](https://doi.org/10.1080/18756891.2014.853954).
- L. F. A. M. Gomes, L. A. D. Rangel, and M. d. R. Leal Junior. Treatment of uncertainty through the interval smart/swing weighting method: a case study. *Pesquisa Operacional*, 31(3):467–485, 12 2011. ISSN 0101-7438. doi:[10.1590/S0101-74382011000300004](https://doi.org/10.1590/S0101-74382011000300004).
- G. Hermerén, I. Brinck, J. Persson, and N.-E. Sahlin. *Value uncertainty and value instability in decision-making*, pages 100–110. Liber Amicorum Pascal Engel. University of Geneva, 2014.
- I. N. Herstein and J. Milnor. An axiomatic approach to measurable utility. *Econometrica*, 21(2): 291–297, Apr. 1953.
- E. Jones, T. Oliphant, P. Peterson, et al. SciPy: Open source scientific tools for Python, 2001–. URL <http://www.scipy.org/>. [Online; accessed 2017-02-20].
- R. L. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. Cambridge University Press, 1993. ISBN 0-521-44185-4.
- I. Levi. *The Enterprise of Knowledge. An Essay on Knowledge, Credal Probability, and Chance*. MIT Press, Cambridge, 1980.
- J. Mustajoki, R. P. Hämäläinen, and A. Salo. Decision support by interval SMART/SWING: Incorporating imprecision in the SMART and SWING methods. *Decision Sciences*, 36(2):317–339, 2005. ISSN 1540-5915. doi:[10.1111/j.1540-5414.2005.00075.x](https://doi.org/10.1111/j.1540-5414.2005.00075.x).
- M. Riabacke, M. Danielson, L. Ekenberg, and A. Larsson. *A Prescriptive Approach for Eliciting Imprecise Weight Statements in an MCDA Process*, pages 168–179. Springer Berlin Heidelberg, 2009. ISBN 978-3-642-04428-1. doi:[10.1007/978-3-642-04428-1_15](https://doi.org/10.1007/978-3-642-04428-1_15).
- M. Riabacke, M. Danielson, and L. Ekenberg. State-of-the-art prescriptive criteria weight elicitation. *Advances in Decision Sciences*, 2012. doi:[10.1155/2012/276584](https://doi.org/10.1155/2012/276584).
- T. Seidenfeld, M. J. Schervish, and J. B. Kadane. A representation of partially ordered preferences. *The Annals of Statistics*, 23:2168–2217, 1995.
- D. von Winterfeldt and W. Edwards. *Decision Analysis and Behavioral Research*. Cambridge University Press, Cambridge, 1986.
- P. Walley. *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, London, 1991.

- P. M. Williams. Notes on conditional previsions. Technical report, School of Math. and Phys. Sci., Univ. of Sussex, 1975.
- P. M. Williams. Notes on conditional previsions. *International Journal of Approximate Reasoning*, 44(3):366–383, 2007. doi:[10.1016/j.ijar.2006.07.019](https://doi.org/10.1016/j.ijar.2006.07.019).
- D. Yemshanov, F. H. Koch, Y. Ben-Haim, M. Downing, F. Sapio, and M. Siltanen. A new multi-criteria risk mapping approach based on a multiattribute frontier concept. *Risk Analysis*, 33(9): 1694–1709, 2013. ISSN 1539-6924. doi:[10.1111/risa.12013](https://doi.org/10.1111/risa.12013).