

Efficient Pairwise Preference Elicitation Allowing for Indifference

Juergen Branke^a, Salvatore Corrente^b, Salvatore Greco^{b,c}, Walter Gutjahr^d

^a*Warwick Business School, The University of Warwick, Coventry, CV4 7AL, United Kingdom*

^b*Department of Economics and Business, University of Catania, Corso Italia, 55, 95129 Catania, Italy*

^c*Operations and Systems Management, Portsmouth Business School, Portsmouth PO1 3DE, United Kingdom*

^d*University of Vienna, 1090 Wien, Austria*

Abstract

Many methods in Multi-Criteria Decision Analysis for choice problems rely on eliciting pairwise preference information in their attempt to efficiently identify the most preferred solution out of a larger set of solutions. That is, they repeatedly ask the decision maker which of two solutions is preferred, and then use this information to reduce the number of possibly preferred solutions until only one remains. However, if the solutions have a very similar value to the decision maker, he/she may not be able to accurately decide which solution is preferred. This paper makes two main contributions. First, it extends Robust Ordinal Regression to allow a user to declare indifference in case the values of the two solutions do not differ by more than some personal threshold. Second, we propose and compare several heuristics to pick pairs of solutions to be shown to the decision maker in order to minimize the number of interactions necessary.

Keywords: Multi-Criteria Decision Analysis, pairwise preference elicitation, indifference, efficient information collection, robust ordinal regression

1. Introduction

Identifying the most preferred solution out of a given set of alternatives may be a challenging task for a decision maker (DM) if there are many

alternatives and criteria. Multi-Criteria Decision Analysis (MCDA) supports users in this task, and a variety of methods have been proposed. Robust Ordinal Regression (ROR) [24] elicits pairwise preference information from the DM, and uses this information together with an assumed underlying preference model to derive additional preference relations and thereby enrich the preference ordering. A key advantage of ROR is that it simultaneously takes into account *all* value functions compatible with the DM’s preference information.

In practice, it may be difficult for a DM to specify which of two solutions is more preferred if these solutions have a very similar value (note that they may be quite different, but have different strengths and weaknesses so that overall, they constitute a similar value to the DM). In such cases, forcing a DM to make a choice is likely to lead to errors, which can then lead to inconsistencies in the preference information such as intransitivity of preferences (the user may state that a is preferred over b , b is preferred over c , but c is preferred over a). This, in turn, may lead to an empty set of compatible value functions and abortion of the method. A few papers have proposed mechanisms to deal with inconsistent preference information, and usually suggest to discard the oldest preference information (e.g., [35]).

In our paper, we extend ROR to allow the DM to declare indifference between two solutions in case the values of the two solutions are similar. In particular, we assume that the DM has a personal internal (unknown to our method) “precision threshold” δ^T , and will declare indifference if the difference in value between the two solutions is within δ^T , i.e., $|U(a) - U(b)| \leq \delta^T$, where $U(a)$ denotes the utility or value of alternative a . This means that we are supposing that the preference relation \succsim over the set of considered alternatives is a semiorder, that is, \succsim is strongly complete, Ferrers transitive and semitransitive (see [16]). Our extended version of ROR will attempt to learn the DM’s precision threshold along with their value function. This approach is easier on the DM and avoids the complications of forcing the DM to specify a preference even if the alternatives have very similar values.

Let us underline that the meaning of the threshold δ^T used in the paper is completely different from that one mentioned in Roy et al. [40]. Indeed, in the mentioned paper, the thresholds are used to represent arbitrariness, imprecision, uncertainty and ill-determination on the performances of the considered alternatives on the criteria at hand. In our context, the threshold is related to the nature of the preference relation \succsim that is supposed to be a semiorder permitting to represent the indecisiveness of the DM in choosing

between two alternatives [33].

We propose two basic algorithms for this scenario, one that attempts to identify all solutions acceptable to the DM (note that because of the DM's indifference threshold, it may happen that there are several solutions that are equally acceptable to the DM), and one that attempts to identify at least one acceptable solution.

Furthermore, in this paper we propose several heuristics to pick the pairs of solutions to be shown to the DM with the goal to minimize the expected number of interactions necessary to identify the most preferred solution. To test ROR with indifference information and these heuristics, we propose a configurable benchmark generator that allows to generate a wide range of artificial benchmark problems.

The paper is structured as follows. We start in Section 2 with an overview on related work. Section 3 defines some fundamental concepts when allowing for indifference information. The actual algorithms implementing a version of ROR that allows for indifference information can be found in Section 4. The heuristics to speed up convergence are explained in Section 5, followed by empirical evaluation in Section 6. The paper concludes with a summary and some ideas for future work in Section 7.

2. Related work

2.1. Multi-Criteria Decision Analysis and value functions

MCDA (see [15, 23]) is a methodology to support complex decisions in which a plurality of criteria have to be taken into consideration. In the MCDA perspective, the basic elements of a decision problem are the set of alternatives $A = \{a, b, \dots\}$, that can be either finite or infinite (in this paper we consider finite set A , only), and the set of criteria $G = \{g_1, \dots, g_m\}$ by which the alternatives from A are evaluated [38]. Without loss of generality, each criterion $g_j \in G$ can be considered as a real-valued function $g_j : A \rightarrow \mathcal{I}_j \subseteq \mathbb{R}$, such that for all $a, b \in A$, $g_j(a) \geq g_j(b)$ if and only if a is at least as good as b on criterion g_j and $\mathcal{I}_j = \{g_j(a) : a \in A\}$ ¹. In the following, for the sake of simplicity, sometimes we identify the criteria g_j with their indices j . MCDA deals with the following three basic decision problems [39]:

¹Therefore, \mathcal{I}_j is the set of performances of alternatives in A with respect to criterion $g_j \in G$.

- the *choice* problem, in case the best alternative or a small set composed of the best alternatives have to be selected;
- the *sorting* problem, in case each alternative has to be assigned to some pre-defined and ordered categories;
- the *ranking* problem, in case the alternatives have to be ordered from the best to the worst.

In this paper we shall consider only the choice problem. When for a and b from A , $g_j(a) \geq g_j(b)$ for all criteria $g_j \in G$ one says that a (weakly) dominates b , and in this case, there is no doubt that a is comprehensively at least as good as b . However, in practical problems, dominance is rarely observed and, therefore, to compare at a comprehensive level alternatives from A it is necessary to aggregate the evaluations assigned to the alternatives by criteria from G . Very often a value function representing the overall desirability of alternatives $a \in A$ by means of a synthesis of their evaluations on criteria $g_j \in G$ is considered. Formally a value function $U : \prod_{j=1}^m \mathcal{I}_j \rightarrow \mathbb{R}$, assigns a value to each alternative from A such that for any $a, b \in A$, a is at least as good as b ($a \succsim b$) if and only if $U(g_1(a), \dots, g_m(a)) \geq U(g_1(b), \dots, g_m(b))$.

Very often the value function U is assumed to be additive [30, 49], that is, for all $a \in A$

$$U(g_1(a), \dots, g_m(a)) = \sum_{j=1}^m u_j(g_j(a)), \quad (1)$$

with $u_j(g_j(a))$ non-decreasing for all $g_j \in G$.

A special case of the additive form of the value function U , very widely used in the literature and in the applications, is the weighted sum, that is

$$U(g_1(a), \dots, g_m(a)) = \sum_{j=1}^m w_j g_j(a), \quad (2)$$

with $w_j \geq 0$ for all $g_j \in G$ and $\sum_{j=1}^m w_j = 1$.

In some decision problems it seems reasonable to consider interaction between criteria. Suppose that evaluations of different criteria are expressed on a common scale (several methods have been proposed for this, see e.g. [3, 14]). In this case, we can apply a value function expressed as a weighted sum (2). Consider a decision problem related to a choice of a car taken from [2]. If the DM is interested in a sports car, the criteria maximum speed and

acceleration have surely a relevant weight. However, in general cars with a high maximum speed have also a good acceleration and there is therefore the risk of an over-evaluation of those cars. Thus, we can say that there is redundancy between the criteria of maximum speed and acceleration. This can be explained by saying that the weight of the two criteria maximum speed and acceleration considered together is smaller than the sum of the weight of the two criteria considered separately. On the other hand, still for a DM interested in sports cars, price is not a criterion as relevant as maximum speed. However, since in general a car with a good maximum speed is expensive, a car having both a good maximum speed and a good price is particularly appreciated. Thus we can say that there is synergy between the criteria of maximum speed and price. This can be explained by saying that the weight of the two criteria maximum speed and price considered together is greater than the sum of the weight of the two criteria considered separately. To model redundancy and synergy between criteria a widely used model is the Choquet integral [10]. It is based on a representation of weight of criteria in terms of a capacity, which assigns to each subset of criteria $T \subseteq G$, the value $\mu(T)$ corresponding to the total weight of criteria from T . The weights $\mu(T), T \subseteq G$, are not required to be additive, which means that for any $S, R \subseteq G$ such that $S \cap R = \emptyset$, one can have $\mu(S \cup T) = \mu(S) + \mu(T)$ in case there is neither redundancy nor synergy between S and T , but also $\mu(S \cup T) < \mu(S) + \mu(T)$ in case of redundancy, and $\mu(S \cup T) > \mu(S) + \mu(T)$ in case of synergy. Formally, denoting by 2^G the power set of G being the set of all subsets of G , the capacity μ is a set function $\mu : 2^G \rightarrow [0, 1]$ that satisfies the following properties:

- 1a) $\mu(\emptyset) = 0$ and $\mu(G) = 1$ (boundary conditions),
- 2a) $\forall S \subseteq T \subseteq G, \mu(S) \leq \mu(T)$ (monotonicity condition).

Given $a \in A$ and a capacity μ on 2^G , the Choquet integral (for application of Choquet integral to MCDA see [21]), as above explained, gives the analogous of the weighted sum in case of additive weights, and it is defined as follows:

$$C_\mu(a) = \sum_{i=1}^m [g_{(i)}(a) - g_{(i-1)}(a)] \mu(\{g_j \in G : g_j(a) \geq g_{(i)}(a)\}), \quad (3)$$

where $g_{(\cdot)}$ reorders the criteria so that $g_{(1)}(a) \leq \dots \leq g_{(m)}(a)$, and $g_{(0)}(a) = 0$. A useful concept dealing with the Choquet integral is the Möbius representation of the capacity μ being the function $m : 2^G \rightarrow \mathbb{R}$ [37, 43] such that, for all $S \subseteq G$

$$\mu(S) = \sum_{T \subseteq S} m(T). \quad (4)$$

Equation (4) is equivalent to

$$m(S) = \sum_{T \subseteq S} (-1)^{|S-T|} \mu(T) \quad (5)$$

which permits to easily compute the Möbius representation.

The properties 1a) and 2a) can be reformulated also in terms of the Möbius representation [9] as follows:

$$1b) \quad m(\emptyset) = 0, \quad \sum_{T \subseteq G} m(T) = 1,$$

$$2b) \quad \forall i \in G \text{ and } \forall R \subseteq G \setminus \{i\}, \quad m(\{i\}) + \sum_{T \subseteq R} m(T \cup \{i\}) \geq 0.$$

The Möbius representation permits to express the Choquet integral in a linear form as a weighted sum of minimum values given to the considered alternative $a \in A$ by all subsets of criteria T from G [20],

$$C_\mu(a) = \sum_{T \subseteq G} m(T) \min_{i \in T} g_i(a). \quad (6)$$

Observe that the terms $m(T) \min_{i \in T} g_i(a)$ in (6) can be interpreted as representing the contribution of the interaction between criteria from $T \subseteq G$ to the overall evaluation $C_\mu(a)$ of alternative $a \in A$. In order to handle a simpler decision model, in many applications it seems reasonable to take into account interactions between no more than k criteria. This is obtained by putting $m(T) = 0$ if $|T| > k$ and the corresponding capacities are called k -additive [22]. Very frequently 2-additive capacities are used so that the Möbius representation and the corresponding Choquet integral become

$$\mu(S) = \sum_{i \in S} m(\{i\}) + \sum_{\{i,j\} \subseteq S} m(\{i,j\}), \quad \forall S \subseteq G \quad (7)$$

$$C_\mu(a) = \sum_{i \in G} m(\{i\})g_i(a) + \sum_{\{i,j\} \subseteq G} m(\{i,j\})\min\{g_i(a), g_j(a)\}. \quad (8)$$

Moreover, properties 1b) and 2b) reduce to

$$1c) \quad m(\emptyset) = 0, \quad \sum_{i \in G} m(\{i\}) + \sum_{\{i,j\} \subseteq G} m(\{i,j\}) = 1,$$

$$2c) \quad \begin{cases} m(\{i\}) \geq 0, \quad \forall i \in G, \\ m(\{i\}) + \sum_{j \in T} m(\{i,j\}) \geq 0, \quad \forall i \in G \text{ and } \forall T \subseteq G \setminus \{i\}, T \neq \emptyset. \end{cases}$$

2.2. Ordinal Regression and Robust Ordinal Regression

Each decision model requires specification of some parameters. For example, if one adopts the additive value function (1), it is necessary to define the parameters related to the formulation of the marginal value functions $u_j(g_j(a)), j = 1, \dots, m$. Eliciting directly the parameters of the adopted decision model from the DM in real-world decision making situations requires a technical background from the DM that in general is not reasonable to expect. Moreover, even in the case that the DM has such a background, a high cognitive effort is required so that, in any case, the preference information elicited in this way is not reliable enough. Consequently, within MCDA, many methods have been proposed to determine the parameters characterizing the considered decision model in an indirect way, i.e., inducing the values of parameters in the adopted decision model, from some holistic preference comparisons of alternatives given by the DM. This indirect preference elicitation does not require any technical background and requires a quite mild cognitive effort. The use of such an indirect preference information is the basis of the ordinal regression paradigm.

The ordinal regression methodology (for a survey see [44]) was proposed by Jacquet-Lagrèze and Siskos [27] that introduced the UTA (UTILités Additives) method which aims at inferring one or more additive value functions in the form (1) from some preference information expressed by the DM in terms of pairwise comparisons of some from alternatives A , so that, for $a, b \in A$:

- $a \succsim_{DM} b$ if for the DM alternative a is at least as good as alternative b ,

- $a \succ_{DM} b$ if for the DM alternative a is preferred to alternative b (that is $a \succsim_{DM} b$ and $\text{not}(b \succsim_{DM} a)$),
- $a \sim_{DM} b$ if for the DM alternative a is indifferent to alternative b (that is $a \succsim_{DM} b$ and $b \succsim_{DM} a$).

Technically, ordinal regression requires to solve an optimization problem of the type:

$$\begin{aligned}
& \max \varepsilon, \text{ s.t.} \\
& \left. \begin{aligned}
& U(a) \geq U(b) \text{ if } a \succsim_{DM} b, \text{ with } a, b \in A, \\
& U(a) \geq U(b) + \varepsilon \text{ if } a \succ_{DM} b, \text{ with } a, b \in A, \\
& U(a) = U(b) \text{ if } a \sim_{DM} b, \text{ with } a, b \in A,
\end{aligned} \right\} E^{HPC} \\
& \left. \begin{aligned}
& \text{constraints related to the analytical form of the} \\
& \text{adopted value function}
\end{aligned} \right\} E^{AF}
\end{aligned} \quad \left. \vphantom{\begin{aligned} \max \varepsilon, \text{ s.t.} \end{aligned}} \right\} E^{OR} \quad (9)$$

The nature of the optimization problem (9) depends on the form of the value function. Generally it is a linear programming problem. This is the case of additive value functions (1), weighted sum value functions (2) and of the value function in the form of the Choquet integral (3).

If the set of constraints E^{OR} is feasible and $\varepsilon^* > 0$, with $\varepsilon^* = \max \varepsilon$, then there exists at least one value function of the considered form that is compatible with the preference information supplied by the DM. On the contrary, if there is no compatible value function, one can select the value function U on the considered class that minimizes the sum of deviation errors or the number of ranking errors in the sense of Kendall or Spearman distance [27]. When there is a compatible value function, in general, there is a plurality of value functions of the selected class (e.g. weighted sum, additive value function, Choquet integral) that permits to represent the preference information supplied by the DM. The methods proposed within the ordinal regression paradigm select one of these value functions according to some predefined principle. One has to admit that to take into consideration only one of the many compatible value functions is always arbitrary to some extent. Because of this observation, Greco et al. [24] proposed to take into account the whole set \mathcal{U} of value functions compatible with the preference information supplied by the DM through a methodology called Robust Ordinal Regression (for some surveys see [12, 13, 25]) that is based on the necessary and possible preference relations \succsim^N and \succsim^P defined as follows:

- $a \succsim^N b$ iff $U(a) \geq U(b)$ for all $U \in \mathcal{U}$, with $a, b \in A$,

- $a \succsim^P b$ iff $U(a) \geq U(b)$ for at least one $U \in \mathcal{U}$, with $a, b \in A$.

Necessary and possible preferences \succsim^N and \succsim^P on A are characterized by the following properties [19]:

- \succsim^N is a partial preorder on A ,
- $\succsim^N \subseteq \succsim^P$,
- $a \succsim^N b$ and $b \succsim^P c \Rightarrow a \succsim^P c, \forall a, b, c \in A$,
- $a \succsim^P b$ and $b \succsim^N c \Rightarrow a \succsim^P c, \forall a, b, c \in A$,
- $a \succsim^N b$ or $b \succsim^P a, \forall a, b \in A$.

To compute the necessary and possible preference relations one has to consider the following sets of constraints,

$$\left. \begin{array}{l} U(b) \geq U(a) + \varepsilon \\ E^{OR}, \end{array} \right\} E^N(a, b) \qquad \left. \begin{array}{l} U(a) \geq U(b) \\ E^{OR}. \end{array} \right\} E^P(a, b)$$

We have that:

- $a \succsim^N b$ if $E^N(a, b)$ is infeasible or if $\varepsilon^N \leq 0$ where $\varepsilon^N = \max \varepsilon$ subject to $E^N(a, b)$,
- $a \succsim^P b$ if $E^P(a, b)$ is feasible and $\varepsilon^P > 0$ where $\varepsilon^P = \max \varepsilon$ subject to $E^P(a, b)$.

Another methodology that takes into consideration the whole set of value functions compatible with the preference expressed by the DM is a recent development of Stochastic Multiattribute Acceptability Analysis (SMAA, [31, 32]) proposed in [29], where for each pair of alternatives $a, b \in A$ one computes the probability that a is preferred to b measured in terms of the hypervolume of the set of compatible value functions U such that $U(a) > U(b)$, taking the hypervolume of \mathcal{U} equal to 1. With the same approach one can compute also the probability that each alternative $a \in A$ attains a given ranking position. The hypervolumes considered in this extension of SMAA are approximated by means of a sampling procedure in the polyhedron of the value functions compatible with the preference information supplied by

the DM based on the Hit-And-Run procedure (HAR; [45, 46]). For an extension of this SMAA methodology to the case of value functions expressed as Choquet integral see [3].

In case of a value function expressed as a Choquet integral with respect to a 2-additive capacity, the linear programming problem related to ordinal regression becomes [34]:

$$\begin{array}{l}
\max \varepsilon, \text{ s.t.} \\
\left. \begin{array}{l}
C_\mu(a) \geq C_\mu(b) \text{ if } a \succsim_{DM} b, \\
C_\mu(a) \geq C_\mu(b) + \varepsilon \text{ if } a \succ_{DM} b, \\
C_\mu(a) = C_\mu(b) \text{ if } a \sim_{DM} b,
\end{array} \right\} E^{HPC} \\
\left. \begin{array}{l}
m(\{\emptyset\}) = 0, \sum_{i \in G} m(\{i\}) + \sum_{\{i,j\} \subseteq G} m(\{i,j\}) = 1, \\
m(\{i\}) \geq 0, \forall i \in G, \\
m(\{i\}) + \sum_{j \in T} m(\{i,j\}) \geq 0, \forall i \in G \text{ and } \forall T \subseteq G \setminus \{i\}, T \neq \emptyset.
\end{array} \right\} E^{AF}
\end{array} \left. \vphantom{\begin{array}{l} E^{HPC} \\ E^{AF} \end{array}} \right\} E^{OR}$$

According to [8], indifference information can be formulated using an indifference threshold, so that a minimal difference between the evaluations of alternatives a and b based on the adopted value function U does not result in a preference, but an indifference between the two alternatives. Formally, the following constraint transformations are performed:

$$\begin{array}{ll}
U(a) - U(b) = 0 & \Rightarrow |U(a) - U(b)| \leq \delta^T, \quad \text{if } a \sim_{DM} b, \\
U(a) > U(b) & \Rightarrow U(a) > U(b) + \delta^T, \quad \text{if } a \succ_{DM} b, \\
U(a) \geq U(b) & \Rightarrow U(a) \geq U(b) - \delta^T, \quad \text{if } a \succsim_{DM} b,
\end{array}$$

with $\delta^T > 0$ being a predefined indifference threshold and $a, b \in A$. The indifference threshold δ^T increases the soundness of the preference information supplied by the DM, because it avoids to force a choice between alternatives a and b when the DM perceives only a negligible advantage of one alternative over the other. Let us also remember that preference structures with indifference thresholds, called semiorders, have been proposed in [33] and largely investigated in the literature (e.g. [17, 36]), so that it appears quite natural to use them also in the ordinal regression approach.

2.3. Impact of indifference threshold

The following experiment helps us to better understand the relationship between the value of the indifference threshold and the percentage of pairs

of alternatives being indifferent. We considered 100 alternatives evaluated on three and five criteria and assumed the preferences of the DM are representable by a value function expressed as a weighted sum. We randomly sampled

- the evaluations of 100 alternatives on the considered criteria in the interval $[0, 100]$,
- the weights assigned to the considered criteria.

Then we could compute the minimal value and the maximal value of the threshold corresponding to some given number of pairs of alternatives being indifferent. We repeated this 1000 times based on different random numbers, for the case of three and five objectives. The results are summarized in Table 1. As can be seen, already for relatively small values of the threshold, a relatively large number of pairs of solutions become indifferent. For example, with three criteria, on average, a value of the threshold just slightly greater than 1% of the scale length is enough to have 150 indifferent pairs of alternatives, and with five criteria, again on average, already a threshold slightly greater than 0.8% is sufficient. This demonstrates the importance of taking indifference into account, because otherwise, there is the risk to misrepresent a DM's preference statements of pairs of alternatives for which there is no clear advantage of one over the other. Further evidence in this sense is given by Figure 1 showing for each number of pairs of indifferent alternatives the corresponding average threshold obtained in the same sequence of 1000 random data sets used for Table 1.

Table 1: Average values and standard deviations of the minimal and maximal threshold δ corresponding to ci couples of indifferent alternatives in case of 100 alternatives, with 3 criteria and 5 criteria (the values are expressed in percentage of the scale length)

		ci 20		ci 80		ci 100		ci 150		ci 200		ci 300		ci 400		ci 500	
		min δ	max δ	min δ	max δ	min δ	max δ	min δ	max δ	min δ	max δ	min δ	max δ	min δ	max δ	min δ	max δ
3 criteria	<i>mean</i>	0.139	0.034	0.559	0.086	0.699	0.104	1.053	0.146	1.410	0.188	2.119	0.27	2.827	0.352	3.537	0.43
	<i>std</i>	0.146	0.035	0.566	0.087	0.706	0.105	1.060	0.147	1.417	0.189	2.116	0.271	2.834	0.353	3.543	0.43
5 criteria	<i>mean</i>	0.109	0.028	0.438	0.065	0.549	0.078	0.822	0.11	1.095	0.139	1.642	0.195	2.192	0.258	2.743	0.316
	<i>std</i>	0.114	0.029	0.443	0.066	0.555	0.079	0.827	0.11	1.101	0.139	1.648	0.196	2.198	0.258	2.748	0.317

2.4. Procedures for optimizing preference elicitation

Several approaches have been proposed in the literature to optimize the procedure of preference elicitation. Some approaches suggest specific decision rules to rank alternatives on the basis of an incomplete ordinal information. For example [5] proposes to consider the ranking corresponding to




Figure 1: Average value of the indifference threshold δ on the y axis; given number of pairs of indifferent alternatives on x axis; 100 alternatives with 1000 random data sets. The scale for value function U is $[0, 100]$.

the weighted sum with equal weights or rank ordered centroid, showing that the latter gives a good approximation of the DM’s preferences. Other rules based on the use of incomplete information are discussed in [28, 41, 42]. Other procedures try to propose the DM queries that permit to cut maximally the space of the parameters in the selected preference model (very often weights of a value function expressed in terms of a weighted sum (2) (see e.g. [18, 26, 47])). Other approaches try to propose queries for which it is maximal the expected reduction of the entropy in the space of the value functions compatible with the preferences expressed by the DM [1, 48]. Other procedures propose questions based on the minimization of the maximum regret being the worst-case loss when recommending an alternative a instead of another alternative b ([7, 50]; for an application of this approach to a value function expressed as a Choquet integral see [6]). Finally, another approach recently proposed, considers heuristics based on necessary preferences obtained by robust ordinal regression and pairwise preference indices and ranking acceptability indices obtained by SMAA [11]. Our contribution is in the same spirit of this last approach, with an important difference being that we focus on choice problem rather than on ranking problem.

3. Robust Ordinal Regression allowing for indifference

As pointed out in Section 2.2, in practice, DMs are not able to make arbitrarily precise judgements when comparing pairs of solutions. Thus, we need

to assume that every DM has a personal, unknown “precision threshold” δ^T , and is indifferent between two alternatives if and only if $|U(x) - U(y)| \leq \delta^T$. We want to extend ROR to allow the DM to actually declare indifference in such cases, avoiding errors from a forced preference decision. Since δ^T is not known, the new ROR approach allowing for indifference, that we call RORi, has to learn the possible δ^T along with the possible utility functions U .

In the following, we define a number of concepts that will be helpful in characterizing solutions under the RORi framework. We also, where possible, explain how these characteristics may be checked. For certain classes of user preference models, the mathematical programs that we provide below are simple Linear Programs and can be evaluated efficiently. This is true for example for linear preference models, the two-additive Choquet integral, the Cobb-Douglas utility model or the monotonic additive utility model. If not stated otherwise, we shall assume this situation and will therefore call our mathematical programs “LPs” even if they can be generalized to the nonlinear case.

Let A denote the set of available alternatives and P the set of pairwise preferences and indifferences elicited from the DM so far, where a relation $(a \succ b) \in P$ means the DM has stated that a is preferred over b , and $(a \sim b) \in P$ represents an indifference statement. We assume A to be finite.

Definition 1 (Possibly acceptable). *A solution x is called possibly acceptable if there is at least one (U, δ) compatible with user preference information P such that $\forall y \in A \setminus \{x\} : U(x) \geq U(y) - \delta$.*

A solution x is possibly acceptable if and only if the following mathematical program **LP1**(x) returns a solution value $\varepsilon > 0$:

$$\left. \begin{array}{l} \max \varepsilon, \text{ s.t.} \\ U(x) \geq U(y) - \delta \quad \forall y \in A \setminus \{x\} \\ U(a) \geq U(b) + \delta + \varepsilon \quad \forall (a \succ b) \in P \\ |U(a) - U(b)| \leq \delta \quad \forall (a \sim b) \in P \\ \delta \geq 0 \\ E^{AF} \end{array} \right\}$$

Therein, the decision variables are the parameters defining the utility U , plus the variables δ and ε . Note that the third constraint is equivalent to $-\delta \leq U(a) - U(b) \leq \delta$. Hence, by the assumption that the user preference model can be linearized, we get an LP problem.

Definition 2 (Necessarily acceptable). A solution x is called necessarily acceptable if for all (U, δ) compatible with user preference information P and $\forall y \in A \setminus \{x\} : U(x) \geq U(y) - \delta$.

It is not so easy to determine whether a solution x is necessarily acceptable. We have to make sure that there does not exist any triple (U, δ, y) such that $U(x) < U(y) - \delta$. But for this, we have to compare x with each $y \in A \setminus \{x\}$ individually using a separate LP. For specifying how this can be done, we start by defining the following binary relation: Solution x is called *possibly preferred* over y if and only if there is a (U, δ) compatible with user preference information such that $U(x) > U(y) + \delta$. This is the case if and only if the following **LP2**(x, y) returns a solution value $\varepsilon > 0$:

$$\left. \begin{array}{l} \max \varepsilon, \text{ s.t.} \\ U(x) \geq U(y) + \delta + \varepsilon \\ U(a) \geq U(b) + \delta + \varepsilon \quad \forall (a \succ b) \in P \\ |U(a) - U(b)| \leq \delta \quad \forall (a \sim b) \in P \\ \delta \geq 0 \\ E^{AF} \end{array} \right\}$$

For abbreviation, let us write $\text{LP2}(x, y) > 0$ if $\text{LP2}(x, y)$ returns a solution value $\varepsilon > 0$, and $\text{LP2}(x, y) \leq 0$ if it returns a solution value $\varepsilon \leq 0$ or is infeasible. Obviously, a solution x is *necessarily acceptable* if and only if for each $y \in A \setminus \{x\}$, solution y is not possibly preferred over x . This is the case if and only if $\text{LP2}(y, x) \leq 0$ for all $y \in A \setminus \{x\}$ (there is no compatible utility function that strictly prefers y).

Of course it can happen that neither x is possibly preferred over y nor vice versa which brings to the following definition:

Definition 3 (Necessarily indistinguishable). Two solutions x and y are necessarily indistinguishable if and only if for all (U, δ) compatible with user preference information P it holds that $|U(x) - U(y)| \leq \delta$.

This can again be checked by using the above LP2: x and y are necessarily indistinguishable if and only if both $\text{LP2}(x, y) \leq 0$ and $\text{LP2}(y, x) \leq 0$.

Definition 4 (Possibly best). A solution x is called possibly best if there is at least one (U, δ) compatible with user preference information P such that $\forall y \in A \setminus \{x\} : U(x) \geq U(y)$.

In other words, solution x is possibly best if the following mathematical program $\mathbf{LP3}(x)$ returns a solution value $\varepsilon > 0$:

$$\begin{array}{l} \max \varepsilon, \text{ s.t.} \\ \left. \begin{array}{l} U(x) \geq U(y) \quad \forall y \in A \setminus \{x\} \\ U(a) \geq U(b) + \delta + \varepsilon \quad \forall (a \succ b) \in P \\ |U(a) - U(b)| \leq \delta \quad \forall (a \sim b) \in P \\ \delta \geq 0 \\ E^{AF} \end{array} \right\} \end{array}$$

Definition 5 (Necessarily best). *A solution x is called necessarily best if for all (U, δ) compatible with user preference information P and $\forall y \in A \setminus \{x\} : U(x) \geq U(y)$.*

Similarly to necessarily acceptable, determining whether a solution x is necessarily best is computationally demanding.

Definition 6 (Possibly globally preferred). *A solution x is called possibly globally preferred if there is at least one (U, δ) compatible with user preference information P such that $\forall y \in A \setminus \{x\} : U(x) > U(y) + \delta$.*

The following $\mathbf{LP4}(x)$ determines whether a solution x is possibly globally preferred, which is the case if and only if $\mathbf{LP4}(x)$ returns a solution value $\varepsilon > 0$:

$$\begin{array}{l} \max \varepsilon, \text{ s.t.} \\ \left. \begin{array}{l} U(x) \geq U(y) + \delta + \varepsilon \quad \forall y \in A \setminus \{x\} \\ U(a) \geq U(b) + \delta + \varepsilon \quad \forall (a \succ b) \in P \\ |U(a) - U(b)| \leq \delta \quad \forall (a \sim b) \in P \\ \delta \geq 0 \\ E^{AF} \end{array} \right\} \end{array}$$

Definition 7 (Necessarily globally preferred). *A solution x is called necessarily globally preferred if for all (U, δ) compatible with user preference information P and $\forall y \in A \setminus \{x\} : U(x) > U(y) + \delta$.*

Again, checking whether x is necessarily globally preferred requires separate comparisons with each $y \in A \setminus \{x\}$.

To define our goal, we assume the DM has an underlying true utility function U^T and true precision threshold δ^T (please note that we do not assume the DM is aware of it, only that it exists). Then, the above definitions can be adapted and a solution x is considered

- **(truly) acceptable**, if $U^T(x) \geq U^T(y) - \delta^T$ for all $y \in A$,
- **(truly) best**, if $U^T(x) \geq U^T(y)$ for all $y \in A$.

Two solutions x, y are called **(truly) indistinguishable**, if $|U^T(x) - U^T(y)| \leq \delta^T$.

The selection problem addressed in this paper may then have two possible goals:

- **Goal A:** Identify **all** solutions truly acceptable to the DM, i.e., all x such that there is no y with $U^T(y) > U^T(x) + \delta^T$. This gives the DM the maximal choice of solutions to choose from, which is particularly useful if there are other, secondary criteria that are not explicitly modelled.
- **Goal B:** Identify **at least one** solution truly acceptable to the DM, i.e., at least one x such that there is no y with $U^T(y) > U^T(x) + \delta^T$. While this may generate a smaller set of acceptable solutions compared to Goal A, it is expected to converge faster since it can discard solutions more quickly as long as at least one acceptable solution remains.

In the following, we will present ways to reach those goals efficiently.

4. Algorithms

RORi iteratively presents pairs of alternatives (x, y) to the DM and asks whether the DM prefers x , prefers y , or is indifferent. At any point in time, the algorithm maintains a list $S \subseteq A$ of candidate solutions which is guaranteed to include a solution acceptable to the DM based on the information elicited so far. By eliciting more and more information, the size of this set reduces until there is only one solution left or it is clear that all remaining solutions are necessarily indistinguishable, and the algorithm terminates.

In the following, we present two general algorithmic frameworks that iteratively elicit preference information and guarantee that they will terminate with an identification of exactly the set of all acceptable solutions (Algorithm 1) or a subset of the acceptable solutions including the best alternative (Algorithm 2), corresponding to the above two goals.

Obviously, the size of S in each iteration and the number of preference elicitations required before the algorithm terminates depend on the pairs of solutions shown to the DM. So, the natural question arises exactly which pair of solutions should be shown to the DM in each iteration such that the expected number of required preference elicitations is minimized. This will be discussed in Section 5.

4.1. Identifying all truly acceptable solutions

Algorithm 1 shows the general algorithmic framework that will identify **all** solutions that are *acceptable* to the DM, i.e., all x such that there is no y with $U^T(y) > U^T(x) + \delta^T$.

Algorithm 1 Finding all truly acceptable solutions

- 1: Identify the set of possibly acceptable solutions S using LP1 for each solution.
 - 2: Search for a pair (x, y) ($x, y \in S$) of solutions that have not yet been shown to the DM and that are not necessarily indistinguishable (tested by LP2).
 - 3: If such a pair is found, show it to the DM, update the constraints accordingly, and go to Step 1.
 - 4: Otherwise stop and return set S .
-

Proposition 4.1. *If the true utility function U^T is contained in the set \mathcal{U} of utility functions defined by the preference model E^{AF} , then Algorithm 1 terminates with the set of all truly acceptable solutions.*

Proof. Let $\mathcal{U}^{(k)}$ denote the set of all pairs (U, δ) such that (U, δ) is compatible with the preference information $P^{(k)}$ present at the beginning of the k -th iteration of the algorithm: $(U, \delta) \in \mathcal{U}^{(k)}$ if and only if $U(a) > U(b) + \delta$ $\forall (a \succ b) \in P^{(k)}$ and $|U(a) - U(b)| \leq \delta$ $\forall (a \sim b) \in P^{(k)}$. We always have $(U^T, \delta^T) \in \mathcal{U}^{(k)}$, since the DM never declares a preference information that is inconsistent with his true (U, δ) . Let now x be acceptable, i.e., $U^T(x) \geq U^T(y) - \delta^T$ for all $y \in A$. Then, because $(U^T, \delta^T) \in \mathcal{U}^{(t)}$, x is possibly acceptable in any iteration $k = 1, \dots, t$ and thus passes the test in Step 1 of Algorithm 1 in all iterations. So, in particular, x is contained in the final solution set $S^{(t)}$.

Conversely, assume that $y \in A$ is *not* acceptable, i.e., there is a $w \in A$ such that $U^T(w) > U^T(y) + \delta^T$. Let x denote a truly best solution (which must exist in view of the finiteness of A), i.e. a solution with $U^T(x) \geq U^T(z)$ for all $z \in A$. In particular,

$$U^T(x) \geq U^T(w) > U^T(y) + \delta^T. \quad (10)$$

Equation (10) implies together with

$$\left. \begin{array}{l} U^T(a) > U^T(b) + \delta \quad \forall (a \succ b) \in P^{(k)} \\ |U^T(a) - U^T(b)| \leq \delta \quad \forall (a \sim b) \in P^{(k)} \end{array} \right\}$$

that the constraints of $\text{LP2}(x, y)$ are satisfied for some $\varepsilon > 0$, and so we have $\text{LP2}(x, y) > 0$. The termination condition of Algorithm 1 entails that in the last iteration t , either (i) (x, y) has already been shown to the DM before, or (ii) $\text{LP2}(x, y) \leq 0$ and $\text{LP2}(y, x) \leq 0$.

Since $\text{LP2}(x, y) > 0$ as demonstrated above, (x, y) must have been shown to the DM in a previous iteration $k < t$. In view of (10), the DM has stated $x \succ y$, so the constraint $U(x) > U(y) + \delta$ must have been added to the set of constraints. As a consequence, y is not possibly acceptable in iteration t anymore and is therefore not contained in the final solution set. \square

4.2. Identifying at least one truly acceptable solution

Identifying all acceptable solutions as described in the previous subsection provides the DM with some choice. However, we may be able to converge quicker and to a smaller set if we aim at identifying the true best rather than all acceptable solutions. Unfortunately, identifying the DM's true best solution may not always be possible. If there are other solutions that are within the DM's precision threshold, then we may not be able to conclude which solution is best, simply because the DM would not be able to rank these solutions. The following Algorithm 2 thus finds a subset of acceptable solutions that is guaranteed to include the true best solution. The only difference between Algorithm 2 and Algorithm 1 is that in each iteration, the current set S is determined to include all possibly best solutions rather than all possibly acceptable solutions.

Proposition 4.2. *If the true utility function U^T is contained in the set \mathcal{U} of utility functions defined by the preference model E^{AF} , then Algorithm 2 terminates with a set of solutions S with the following properties:*

Algorithm 2 Finding at least one acceptable solution

- 1: Identify the set of possibly best solutions S using LP3 for each solution.
 - 2: Search for a pair (x, y) ($x, y \in S$) of solutions that have not yet been shown to the DM and that are not necessarily indistinguishable (tested by LP2).
 - 3: If such a pair is found, show it to the DM, update the constraints accordingly, and go to Step 1.
 - 4: Otherwise stop and return set S .
-

(a) *Each truly best solution is contained in S .*

(b) *The solutions in S are pairwise truly indistinguishable:*

$$|U^T(x) - U^T(y)| \leq \delta^T \text{ for all } x, y \in S.$$

Proof. To show (a), let x be truly best. With $\mathcal{U}^{(t)}$ defined analogously to the proof of Proposition 1, we have again $(U^T, \delta^T) \in \mathcal{U}^{(t)}$ in the final iteration t . Therefore, solution x is possibly best in iteration t and thus passes the test in Step 1 of Algorithm 2 in this iteration. As a consequence, x is contained in the final set S . To show (b), let us consider two different solutions $x, y \in S$. According to the termination condition of Algorithm 2, either (x, y) has already been shown to the DM before, or

$$\text{LP2}(x, y) \leq 0 \text{ and } \text{LP2}(y, x) \leq 0. \quad (11)$$

In the first case where (x, y) has already been shown in some previous iteration $k < t$, the DM can only have stated indifference, since if he/she would have stated (say) $x \succ y$, then the constraint $U(x) > U(y) + \delta$ would have been added to the set of constraints, with the consequence that y would not have passed the test in Step 1 of iteration t , being not possibly best anymore. So there remains only the case that iteration t produces (11). In this case, however, x and y are necessarily indistinguishable based on the information present in iteration t , i.e., $|U(x) - U(y)| \leq \delta \quad \forall (U, \delta) \in \mathcal{U}^{(t)}$. In particular, because of $(U^T, \delta^T) \in \mathcal{U}^{(t)}$, this implies $|U^T(x) - U^T(y)| \leq \delta^T$. \square

In the call of LP3 by Algorithm 2, the constraints of the first type, $U(x) \geq U(y) \quad \forall y \in A \setminus \{x\}$, have to be satisfied for each solution y different from x . A natural question is whether in the subsequent iterations k , we can successively restrict the set of solutions y against which this constraint has to be checked to those solutions $y \in A \setminus \{x\}$ that have remained in the

solution set in the previous iteration $k - 1$. The following result provides a positive answer to this question.

Proposition 4.3. *If in the call of LP3 by the k -th iteration of Algorithm 2, the constraint $U(x) \geq U(y) \forall y \in A \setminus \{x\}$ is replaced by the constraint $U(x) \geq U(y) \forall y \in S^{(k-1)} \setminus \{x\}$, where $S^{(\ell)}$ denotes the set of solutions as produced in iteration ℓ of the algorithm, then the feasible set of LP3 remains unchanged, and therefore the same set of possibly best solutions x is produced in this iteration.*

Proof. Let $A_{\mathcal{U}}^*$ denote the set of possibly best solutions with respect to preference information \mathcal{U} , and let $\mathcal{U}^{(k)}$ be defined analogously to the proof of Proposition 1. Then $S^{(k-1)} = A_{\mathcal{U}^{(k-1)}}^*$. First, we show that for each preference information \mathcal{U} on A , for each fixed $(U, \delta) \in \mathcal{U}$ and for each fixed $x \in A$, the condition

$$U(x) \geq U(y) \quad \forall y \in A_{\mathcal{U}}^* \setminus \{x\} \quad (12)$$

implies

$$U(x) \geq U(y) \quad \forall y \in A \setminus \{x\}. \quad (13)$$

Indeed, assume that (12) holds, but that (13) is violated for some y , i.e., $U(x) < U(y)$. Let $z^* \in \arg \max_{z \in A \setminus \{x\}} U(z)$. Then $U(y) \leq U(z^*)$, hence $U(x) < U(z^*)$, and z^* is obviously possibly best with respect to \mathcal{U} , that is, $z^* \in A_{\mathcal{U}}^* \setminus \{x\}$. This contradicts (12).

To show the statement of the Proposition, we consider now an iteration k and some fixed $x \in A$. In iteration k , the constraints of LP3 restrict the feasible set to $(U, \delta) \in \mathcal{U}^{(k)}$. Since the constraints are extended in each iteration of Algorithm 2, we have $\mathcal{U}^{(k)} \subseteq \mathcal{U}^{(k-1)}$. Thus, in the call of LP3 in iteration k for the given x , each solution (U, δ) that is feasible with respect to the second and the third type of constraints, also satisfies the condition $(U, \delta) \in \mathcal{U}^{(k-1)}$. Consequently, if for such a pair (U, δ) , the inequality

$$U(x) \geq U(y) \quad \forall y \in A_{\mathcal{U}^{(k-1)}}^* \setminus \{x\} \quad (14)$$

corresponding to (12) holds, we can apply the auxiliary result from above with $\mathcal{U} = \mathcal{U}^{(k-1)}$ and conclude that (13) is valid, and therefore (U, δ) is feasible for LP3. Conversely, each (U, δ) that is feasible for LP3 in iteration k to the given x satisfies (13) and therefore also the weaker condition (14). Thus, the feasible sets of the two LP variants are identical. \square

5. Heuristics to speed up convergence

In the above, we have simply restricted the questions (or pairs of solutions) to be shown to the DM to those pairs that are guaranteed to provide new information. In the following, we consider some heuristics to pick a pair among all the “reasonable” pairs of solutions. Denote by \mathcal{S} the set of pairs of solutions in $(S \times S)$ that have not yet been shown to the DM and that are not necessarily indistinguishable at any stage in the interaction process (this set gets smaller and smaller over the course of the interaction, but we use the same symbol in all iterations in order to simplify notation).

The benchmark will be the following simple approach H1.

H1 Pick a pair randomly among all the pairs in \mathcal{S} .

The following heuristic H2 always maintains the solution c identified by the DM as the best so far, and compares it to the solution in S that could possibly be preferred over c by the largest margin across all compatible value functions. This method has been proposed in [4].

H2 The first pair of solutions is picked randomly from \mathcal{S} . Let us denote by c the solution preferred by the DM (one at random in case of indifference). In subsequent iterations, present the pair $(x, c) \in \mathcal{S}$ with $x = \operatorname{argmax} LP2(x, c)$. If the DM prefers x , the best so far solution c is set to x .

H3 picks pairs of solutions that can have maximal contradiction in values based on the set of compatible value functions.

H3 Pick the pair $(x, y) \in \mathcal{S}$ as $(x, y) = \operatorname{argmax}\{LP2(x, y) + LP2(y, x)\}$.

H4 uses the most discriminative value function in the spirit of [35] as a best guess of the DM’s true value function, and then shows the DM the presumed best and second best solution. The hope is that if the DM confirms the presumed ordering, it would also be possible to remove many other solutions from S .

We define the most discriminative value function as the one that maximizes the following LP5, allowing for the maximum slack in each constraint:

max ε , s.t.

$$\left. \begin{array}{l} U(a) \geq U(b) + \delta + \varepsilon \quad \forall (a \succ b) \in P \\ |U(a) - U(b)| \leq \delta - \varepsilon \quad \forall (a \sim b) \in P \\ \delta \geq 0 \\ E^{AF} \end{array} \right\}$$

H4 Identify the most discriminative value function (U^d, δ^d) by solving LP5 above. Then, show the DM the pair of solutions consisting of the best and second best solution according to U^d .

The common idea behind H5, H6 and H7 is to look one step ahead. For each possible pair $(x, y) \in \mathcal{S}$, these heuristics consider the three possible responses of the DM: $x \succ y, x \prec y, x \sim y$. H5 minimizes the number of remaining possibly acceptable solutions in the worst case, whereas H6 minimizes the expected number of the remaining possibly acceptable solutions. H7 looks at the expected resulting entropy, where entropy of a decision situation has been defined in [48] and can be calculated based on a number of value functions (U, δ) compatible with preference information. In particular, the entropy E given a set of solutions S and a finite sample set of utility functions \mathcal{U} is defined as

$$E(S, \mathcal{U}) = - \sum_{x \in S} p_x \log p_x \quad (15)$$

where p_x is the fraction of utility functions in \mathcal{U} for which solution x is best.

To calculate the expected value in H6 and H7, it is necessary to assign probabilities to the three possible responses of the DM. We do this by generating 10,000 random utility functions compatible with the elicited preference information using the algorithm by [45, 46]. Then the probability that the DM will select $x \succ y, x \prec y$, or $x \sim y$ is estimated by the fraction of utility functions that would result in the respective decision.

H5 For each pair $(x, y) \in \mathcal{S}$

- **P1** Add the constraint $U(x) \geq U(y) + \delta + \varepsilon$ to LP1 and determine the number of solutions that can be still possibly acceptable solutions, denoted by PP_{xy}^\succ ,

- **P2** Add the constraint $U(y) \geq U(x) + \delta + \varepsilon$ to LP1 and determine the number of solutions that can be still possibly acceptable solutions, denoted by $PP_{xy}^<$,
- **P3** Add the constraint $|U(y) - U(x)| \leq \delta$ to LP1 and determine the number of solutions that can be still possibly acceptable solutions, denoted by PP_{xy}^{\sim} .

Pick the pair (x, y) that minimizes the number of possibly acceptable solutions in the worst case, i.e., $(x, y) = \operatorname{argmin}_{(x,y)} \max\{PP_{xy}^>, PP_{xy}^<, PP_{xy}^{\sim}\}$

H6 For each pair $(x, y) \in \mathcal{S}$

- **P1** Add the constraint $U(x) \geq U(y) + \delta + \varepsilon$ to LP1 and check how many solutions can be still possibly best solutions, denoted by $PP_{xy}^>$,
- **P2** Add the constraint $U(y) \geq U(x) + \delta + \varepsilon$ to LP1 and check how many solutions can be still possibly best solutions, denoted by $PP_{xy}^<$,
- **P3** Add the constraint $|U(y) - U(x)| \leq \delta$ to LP1 and check how many solutions can be still possibly best solutions, denoted by PP_{xy}^{\sim} ,
- **P4** Generate a set of utility functions (U, δ) compatible with the preference information. Determine what fraction of those utility functions would prefer x , prefer y , or be indifferent $(p_{xy}^>, p_{xy}^<, p_{xy}^{\sim})$.

Pick the pair (x, y) that minimizes the expected number of possibly best solutions, i.e., $(x, y) = \operatorname{argmin}_{(x,y)} \{p_{xy}^> \cdot PP_{xy}^> + p_{xy}^< \cdot PP_{xy}^< + p_{xy}^{\sim} \cdot PP_{xy}^{\sim}\}$.

H7 For each pair $(x, y) \in \mathcal{S}$

- **P1** Add the constraint $U(x) \geq U(y) + \delta + \varepsilon$ to LP3 and check how many solutions can be still be possibly best solutions, denoted by $PB_{xy}^>$,
- **P2** Add the constraint $U(y) \geq U(x) + \delta + \varepsilon$ to LP3 and check how many solutions can be still possibly best solutions, denoted by $PB_{xy}^<$,
- **P3** Add the constraint $|U(y) - U(x)| \leq \delta$ to LP3 and check how many solutions can be still possibly best solutions, denoted by PB_{xy}^{\sim} ,
- **P4** Generate a set of utility functions $(U, \delta) \in \mathcal{U}$ compatible with the preference information. Determine what fraction of those utility functions would prefer x , prefer y , or be indifferent, $(p_{xy}^>, p_{xy}^<, p_{xy}^{\sim})$. Also calculate the entropy of the resulting decision situations $(E_{xy}^>, E_{xy}^<, E_{xy}^{\sim})$.

Pick the pair (x, y) that minimizes the expected entropy of the resulting decision situation, i.e., $(x, y) = \operatorname{argmin}\{p_{xy}^{\succ} \cdot E_{xy}^{\succ} + p_{xy}^{\prec} \cdot E_{xy}^{\prec} + p_{xy}^{\sim} \cdot E_{xy}^{\sim}\}$.

Note that H5, H6 and H7 as described above require to solve 3 LPs for every pair of solutions still in \mathcal{S} that has not yet been shown to the DM, which makes the algorithm rather slow. In the computational experiments reported below, we have thus replaced the solution of the LPs by a numerical approximation based on the 10,000 random compatible utility functions that we require anyway to estimate probabilities in H6 and H7. In particular, for each considered pair of solutions and possible DM preference ($x \succ y$, $x \prec y$, or $x \sim y$), we remove the utility functions that are no longer compatible with the presumed new preference information and simply check how many solutions are still acceptable or best for at least one of the remaining preference functions. A limited number of experiments indicated that this is computationally much faster and even works slightly better, presumably because it ignores cases where a solution would only be acceptable or best by a tiny fraction of compatible utility functions.

6. Empirical results

6.1. Experimental setup

For our experiments, we need an artificial decision maker, and we assume the two-additive Choquet preference model as the DM’s internal model. Our artificial decision maker answers to preference elicitations in terms of pairwise comparisons according to this underlying utility function (U^T, δ^T) , that is of course otherwise unknown to the elicitation algorithms. For every experiment, we average over 50 problem instances, using the same problem instances across algorithms. First, we generate random weights for the different attributes, resulting in a linear utility function of the artificial DM. For each problem instance, we randomly generate the desired number of mutually non-dominated alternatives in real-valued attribute space. Then, the artificial DM’s discrimination threshold δ^T is determined such that we obtain the desired percentage of possible pairwise comparisons that will result in an indifference statement. For this, we start with $\delta^T = 0$ and increase δ^T step by step until we have the desired percentage of indifference statements. The percentage of indifferences among all pairs of solutions is a parameter of the benchmark problem and can be controlled.

Table 2: Average number and standard error of required DM interactions until the algorithm terminates with a subset of all acceptable solutions including the best one, for $N = 100$ and $P = \{20, 80, 150\}$ indifferent pairs and $m = 3$ criteria.

P	H0	H1	H2	H3	H4	H5	H6	H7
20	99	13.98	16.64	13.94	16.78	12.28	10.96	18.0
	± 0	± 0.34	± 0.62	± 0.41	± 0.52	± 0.30	± 0.29	± 0.55
80	99	13.96	16.32	13.74	16.56	12.28	10.56	17.52
	± 0	± 0.34	± 0.63	± 0.43	± 0.52	± 0.31	± 0.31	± 0.61
150	99	13.5	16.3	13.66	16.56	12.08	10.54	17.16
	± 0	± 0.38	± 0.62	± 0.43	± 0.52	± 0.33	± 0.29	± 0.64

As an additional benchmark to demonstrate the benefit of using RORi and our heuristics to pick pairs of solutions to show to the DM we also include the following straightforward algorithm to determine at least one acceptable solution: go through the solutions in a random sequence, initialise the first solution as the current best and compare each solution in turn with the current best, making the challenger the current best only if the challenger is preferred. This simple algorithm is guaranteed to identify an acceptable solution but obviously requires $N - 1$ pairwise comparisons, where N is the number of alternatives to choose from. It is denoted as H0 in the tables below.

6.2. Comparison of elicitation heuristics

For the baseline experiments reported in this section, we use Algorithm 2 as the general framework, $N = 100$ alternatives, $m = 3$ attributes and $P = \{20, 80, 150\}$ pairs of solutions being indifferent. The average number of required DM interactions until the algorithm terminates is summarized in Table 2.

Clearly, all methods are much more efficient than the above mentioned simple approach that would require 99 pairwise comparisons.

Somewhat surprisingly, most of the heuristics tested perform worse than the simple random selection of pairs of solutions (H1). This shows that the use of RORi in the framework of Algorithm 2 is responsible for most of the improvement, and that it is not easy to achieve further improvements by clever heuristics on which pairs of solutions to show to the DM. The worst are the attempt to guess the DM's most likely value function by determining the most discriminatory value function (H4) and the rather sophisticated

Table 3: Percentage of DM interactions that result in an indifference statement, for Algorithm 2, $N = 100$, $P = \{20, 80, 150\}$ indifferent pairs and $m = 3$ criteria.

P	H0	H1	H2	H3	H4	H5	H6	H7
20	0.08%	0.29%	0.24%	0.00%	0.12%	0.33%	0.55%	0.00%
80	0.46%	2.01%	0.74%	1.75%	1.93%	1.95%	3.60%	0.57%
150	0.95%	4.15%	1.47%	3.37%	2.90%	4.80%	6.07%	1.63

heuristic H7 based on look-ahead and entropy calculation, presumably because minimizing entropy is not very well aligned with our goal of minimizing the number of interactions.

Only two of the look-ahead heuristics, namely H5 and H6, are significantly better than random selection. Of the two, the method that greedily minimizes the expected number of remaining solution candidates (H6) works much better than H5 which minimizes the number of remaining solution candidates in the worst case. Although H6 is computationally more expensive to compute because it needs to estimate the probabilities for the DM’s response, this extra effort seems to pay off. H6 requires between 21.6% and 24.3% less interactions than the random selection of pairs.

Table 3 reports on the percentage of interactions with the DM that result in an indifference statement. This percentage is roughly in the ballpark of the percentage of all possible pairwise comparisons for 20, 80 and 150 indifferent pairs, which is 0.4%, 1.6%, 3.0% for 100 alternatives. H7 based on entropy and H0 create the smallest percentage of indifference responses **for** $P = 20$, **and** $P = \{80, 150\}$, **respectively**, whereas Random (H1), H5 and H6 create relatively many indifference responses.

6.3. Identifying all acceptable solutions

Repeating the experiments of the previous section but this time using the framework of Algorithm 1 rather than Algorithm 2 in order to identify *all* acceptable solutions yields the results reported in Table 4. As suspected, identifying all acceptable solutions requires slightly more preference elicitations than identifying only one acceptable solution, independent of the heuristic used. Moreover, comparing Tables 3 and 5, we observe that in most cases, switching from Algorithm 1 to Algorithm 2 the percentage of indifference statements increases, too.

Table 4: Average number and standard error of required DM interactions until the algorithm terminates with the set of all acceptable solutions, for $N = 100$, $P = \{20, 80, 150\}$ indifferent pairs and $m = 3$ criteria.

P	H1	H2	H3	H4	H5	H6	H7
20	16.58	18.06	17.36	20.72	13.26	11.46	19.66
	± 0.45	± 0.65	± 0.57	± 0.88	± 0.29	± 0.31	± 0.63
80	16.74	18.08	17.96	21.14	13.5	11.44	19.2
	± 0.46	± 0.67	± 0.69	± 0.93	± 0.33	± 0.31	± 0.67
150	16.8	18.14	18.4	21.04	13.4	11.58	19.12
	± 0.46	± 0.67	± 0.70	± 0.93	± 0.34	± 0.30	± 0.71

Table 5: Percentage of DM interactions that result in an indifference statement, for Algorithm 1, $N = 100$, $P = \{20, 80, 150\}$ indifferent pairs and $m = 3$ criteria.

P	H1	H2	H3	H4	H5	H6	H7
20	0.7%	0.2%	0.5%	0.3 %	0.5%	0.5%	0%
80	2.7%	0.9%	2.1%	1.6%	1.9%	5.2%	0%
150	5.2%	1.8%	3.9%	2.8%	3.9%	6.6%	1.8%

6.4. Influence of the number of alternatives

In this section, we examine the influence of the number of alternatives on the required number of preference elicitations. In addition to $N = 100$, we examine $N = 50$ and $N = 200$. The number of criteria is still 3, and we choose δ^T in a way that results in 1.6% indifferences, i.e., 20, 80 and 320 indifferent pairs for 50, 100 or 200 alternatives, respectively. Because many of the heuristics resulted in performances worse than random selection, we only consider H1, H5, and H6 here in addition to the straightforward benchmark H0 (see Table 6(a)).

The number of preference elicitations seems to be roughly linear. A linear regression for the H6 heuristic results in $9.29 + 0.018N$ which means that for each additional alternative, the number of expected required preference elicitations only increases by about 0.02.

6.5. Influence of the number of criteria

So far, all experiments were based on problems with three criteria. Here we examine whether the results still hold in the case of 5 criteria. For these experiments, we use the base case of 100 alternatives, 80 indifference relations

Table 6: Average number and standard error of required DM interactions until the algorithm terminates with a set of acceptable solutions,...

(a) ...for different numbers of alternatives N , indifference in 1.6% of all possible pairs and $m = 3$ criteria.					(b) ...for different numbers of criteria m and indifference in $P = 80$ pairs.				
N	H0	H1	H5	H6	m	H0	H1	H5	H6
50	49	12.32	10.92	9.86	3	99	13.96	12.28	10.56
	± 0	± 0.27	± 0.26	± 0.24		± 0	± 0.34	± 0.31	± 0.31
100	99	13.96	12.28	10.56	5	99	33.38	35.02	27.64
	± 0	± 0.34	± 0.31	± 0.31		± 0	± 0.60	± 0.73	± 0.53
200	199	15.0	13.20	11.64					
	± 0	± 0.37	± 0.34	± 0.30					

and identification of the best solution. Again, we only consider the most promising heuristics.

As can be seen from Table 6(b), the necessary number of pairwise comparisons increases sharply with the number of criteria. The reason is probably that the preference model has more parameters, and requires more information in terms of preference elicitations before it is accurate enough to be helpful in pruning solutions from the set of possibly best solutions. The relative performance, however, remains similar, and H6 is still able to reduce the number of preference elicitations by 17.2% compared to RORi with random selection (24.4% for the 3 criteria case), and by 72.1% compared to the default heuristic H0 (89.3% for the 3 criteria case).

7. Discussion and Conclusions

This paper proposes new methods to support a DM to identify his most preferred solution out of a given set of solutions, based on pairwise preference elicitations. The paper makes two main contributions. First, it extends the Robust Ordinal Regression Framework to allow for indifference statements from the DM if the utilities of the two solutions do not differ by more than a user-specific threshold. Our new RORi method learns the user threshold along with the parameters of the utility function. Second, it proposes and compares a number of heuristics to guide the pairs of solutions to be shown to the decision maker, to maximize the information gained from each elicitation.

An empirical analysis shows that RORi is able to appropriately generalize the information collected, and drastically reduce the number of necessary

comparisons compared to an algorithm that considers all alternatives in turn and always keeps the best-so-far — even when the pairs of solutions shown to the DM are picked randomly. The reduction is lower if the number of criteria is higher, because the model then has more parameters and thus more information is needed to derive meaningful parameter values. Similarly, the reduction is higher with a higher number of alternatives, because the savings are low initially but grow with the number of alternatives evaluated as the model becomes more and more precise.

Regarding the proposed heuristics to decide which pair of solutions to show to the DM, surprisingly, most heuristics, including those we took from the literature, perform worse than random sampling when evaluated based on the number of preference elicitations needed to complete the selection. The best heuristic uses a sophisticated one-step look-ahead technique, which on the problems considered further reduces the necessary number of samples by about 20% compared to RORi with random selection of pairs. Together, depending on the test problem, our method saves about 90% of the required pair-wise preference elicitations compared to a simple algorithm that considers all alternatives in turn and always keeps the best-so-far.

There are various avenues for future work. While we have used the Choquet integral as preference model in this paper, an extension to other preference models should be straightforward. More heuristics can be developed, and should be tested on a wider range of benchmark problems. Finally, one may consider showing the DM not only pairs of existing alternatives, but also fictitious alternatives in order to speed up convergence.

Acknowledgments

The second and the third authors wish to acknowledge the funding by the “FIR of the University of Catania BCAEA3, New developments in Multiple Criteria Decision Aiding (MCDA) and their application to territorial competitiveness”.

References

- [1] A.E. Abbas. Entropy methods for adaptive utility elicitation. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 34(2):169–178, 2004.

- [2] S. Angilella, S. Greco, F. Lamantia, and B. Matarazzo. Assessing non-additive utility for multicriteria decision aid. *European Journal of Operational Research*, 158(3):734–744, 2004.
- [3] S. Angilella, S. Corrente, and S. Greco. Stochastic multiobjective acceptability analysis for the Choquet integral preference model and the scale construction problem. *European Journal of Operational Research*, 240(1):172–182, 2015.
- [4] N. Argyris, A. Morton, and J.R. Figueira. CUT: A multicriteria approach for concavifiable preferences. *Operations Research*, 62(3):633–642, 2014.
- [5] F.H. Barron and B.E. Barrett. Decision quality using ranked attribute weights. *Management Science*, 42(11):1515–1523, 1996.
- [6] N. Benabbou, C. Gonzales, P. Perny, and P. Viappiani. Minimax regret approaches for preference elicitation with rank-dependent aggregators. *EURO Journal on Decision Processes*, 3(1-2):29–64, 2015.
- [7] C. Boutilier, R. Patrascu, P.I. Poupard, and D. Schuurmans. Constraint-based optimization and utility elicitation using the minimax decision criterion. *Artificial Intelligence*, 170(8):686–713, 2006.
- [8] J. Branke, S. Corrente, S. Greco, and W.J. Gutjahr. Using indifference information in robust ordinal regression. In *International Conference on Evolutionary Multi-Criterion Optimization*, pages 205–217. Springer, 2015.
- [9] A. Chateauneuf and J.Y. Jaffray. Some characterizations of lower probabilities and other monotone capacities through the use of Möbius inversion. *Mathematical Social Sciences*, 17:263–283, 1989.
- [10] G. Choquet. Theory of capacities. *Annales de l’institut Fourier*, 5(54):131–295, 1953.
- [11] K. Ciomek, M. Kadziński, and T. Tervonen. Heuristics for prioritizing pair-wise elicitation questions with additive multi-attribute value models. *Omega*, 2016, DOI: 10.1016/j.omega.2016.08.012.

- [12] S. Corrente, S. Greco, M. Kadziński, and R. Słowiński. Robust ordinal regression in preference learning and ranking. *Machine Learning*, 93 (2-3):381–422, 2013.
- [13] S. Corrente, S. Greco, M. Kadziński, and R. Słowiński. Robust ordinal regression. *Wiley Encyclopedia of Operations Research and Management Science*, pages 1–10, 2014.
- [14] S. Corrente, S. Greco, and A. Ishizaka. Combining analytical hierarchy process and Choquet integral within non-additive robust ordinal regression. *Omega*, 61:2–18, 2016.
- [15] M. Ehrgott, J.R. Figueira, and S. Greco (eds.). *Trends in Multiple Criteria Decision Analysis*. Springer, New York, 2010.
- [16] P.C. Fishburn. Interval representations for interval orders and semiorders. *Journal of Mathematical Psychology*, 10(1):91–105, 1973.
- [17] P.C. Fishburn. *Interval orders and interval graphs: A study of partially ordered sets*. John Wiley & Sons, 1985.
- [18] K. Gajos and D.S. Weld. Preference elicitation for interface optimization. In *Proceedings of the 18th annual ACM symposium on User interface software and technology*, pages 173–182. ACM, 2005.
- [19] A. Giarlotta and S. Greco. Necessary and possible preference structures. *Journal of Mathematical Economics*, 49(2):163 – 172, 2013.
- [20] I. Gilboa and D. Schmeidler. Additive representations of non-additive measures and the Choquet integral. *Annals of Operations Research*, 52: 43–65, 1994.
- [21] M. Grabisch. The application of fuzzy integrals in multicriteria decision making. *European Journal of Operational Research*, 89:445–456, 1996.
- [22] M. Grabisch. k -order additive discrete fuzzy measures and their representation. *Fuzzy Sets and Systems*, 92:167–189, 1997.
- [23] A. Greco, J.R. Figueira, and M. Ehrgott. *Multiple Criteria Decision Analysis: State of the Art Surveys*. Springer, Berlin, 2016.

- [24] S. Greco, V. Mousseau, and R. Słowiński. Ordinal regression revisited: multiple criteria ranking using a set of additive value functions. *European Journal of Operational Research*, 191(2):416–436, 2008.
- [25] S. Greco, R. Słowiński, J.R. Figueira, and V. Mousseau. Robust ordinal regression. In M. Ehrgott, J.R. Figueira, and S. Greco, editors, *Trends in Multiple Criteria Decision Analysis*, pages 241–283. Springer, 2010.
- [26] V.S. Iyengar, J. Lee, and M. Campbell. Evaluating multiple attribute items using queries. In *Proceedings of the 3rd ACM conference on Electronic Commerce*, pages 144–153. ACM, 2001.
- [27] E. Jacquet-Lagrèze and Y. Siskos. Assessing a set of additive utility functions for multicriteria decision-making, the UTA method. *European Journal of Operational Research*, 10(2):151–164, 1982.
- [28] M. Kadziński and M. Michalski. Scoring procedures for multiple criteria decision aiding with robust and stochastic ordinal regression. *Computers & Operations Research*, 71:54–70, 2016.
- [29] M. Kadziński and T. Tervonen. Robust multi-criteria ranking with additive value models and holistic pair-wise preference statements. *European Journal of Operational Research*, 228(1):169–180, 2013.
- [30] R.L. Keeney and H. Raiffa. *Decisions with multiple objectives: Preferences and value tradeoffs*. J. Wiley, New York, 1976.
- [31] R. Lahdelma and P. Salminen. SMAA-2: Stochastic multicriteria acceptability analysis for group decision making. *Operations Research*, 49(3):444–454, 2001.
- [32] R. Lahdelma, J. Hokkanen, and P. Salminen. SMAA-Stochastic Multi-objective Acceptability Analysis. *European Journal of Operational Research*, 106(1):137–143, 1998.
- [33] R.D. Luce. Semiorders and a theory of utility discrimination. *Econometrica*, 24(2):178–191, 1956.
- [34] J.L. Marichal and M. Roubens. Determination of weights of interacting criteria from a reference set. *European Journal of Operational Research*, 124(3):641–650, 2000.

- [35] S. Phelps and M. Köksalan. An interactive evolutionary metaheuristic for multiobjective combinatorial optimization. *Management Science*, 49(12):1726–1738, 2003.
- [36] M. Pirlot and P. Vincke. *Semiororders: properties, representations, applications*. Springer Science & Business Media, 2013.
- [37] G.C. Rota. On the foundations of combinatorial theory I. Theory of Möbius functions. *Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 2:340–368, 1964.
- [38] B. Roy. *Multicriteria Methodology for Decision Aiding*. Kluwer Academic, Dordrecht, 1996.
- [39] B. Roy. Paradigm and Challenges. In S. Greco, J.R. Figueira, and M. Ehrgott, editors, *Multiple Criteria Decision Analysis: State of the Art Surveys*, pages 19–39. Springer, Berlin, 2016.
- [40] B. Roy, J.R. Figueira, and J. Almeida-Dias. Discriminating thresholds as a tool to cope with imperfect knowledge in multiple criteria decision aiding: Theoretical results and practical issues. *Omega*, 43:9–20, 2014.
- [41] A.A. Salo and R.P. Hämmäläinen. Preference ratios in multiattribute evaluation (prime)-elicitation and decision procedures under incomplete information. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 31(6):533–545, 2001.
- [42] P. Sarabando and L.C. Dias. Simple procedures of choice in multicriteria problems without precise information about the alternatives’ values. *Computers & Operations Research*, 37(12):2239–2247, 2010.
- [43] G. Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, 1976.
- [44] Y. Siskos, E. Grigoroudis, and N.F. Matsatsinis. UTA methods. In S. Greco, J.R. Figueira, and M. Ehrgott, editors, *Multiple Criteria Decision Analysis: State of the Art Surveys*, pages 315–362. Springer, Berlin, 2016.
- [45] R.L. Smith. Efficient monte carlo procedures for generating points uniformly distributed over bounded regions. *Operations Research*, 32(6):1296–1308, 1984.

- [46] T. Tervonen, G. Van Valkenhoef, N. Bastürk, and D. Postmus. Hit-and-run enables efficient weight generation for simulation-based multiple criteria decision analysis. *European Journal of Operational Research*, 224:552–559, 2013.
- [47] O. Toubia, J.R. Hauser, and D.I. Simester. Polyhedral methods for adaptive choice-based conjoint analysis. *Journal of Marketing Research*, 41(1):116–131, 2004.
- [48] G. van Valkenhoef and T. Tervonen. Entropy-optimal weight constraint elicitation with additive multi-attribute utility models. *Omega*, 64:1–12, 2016.
- [49] P.P. Wakker. *Additive representations of preferences: A new foundation of decision analysis*. Springer, 1989.
- [50] T. Wang and C. Boutilier. Incremental utility elicitation with the min-max regret decision criterion. In *IJCAI*, pages 309–318, 2003.