Chapter 14
Eliciting Multi-Criteria Preferences: ELECTRE Models

Luis C. Dias and Vincent Mousseau

Abstract Outranking methods are a specific type of Multi-Criteria Decision Aiding methods. They are based on the construction of binary relations validating or invalidating, for any pair of alternatives (a, b), the assertion "*a outranks b*". This comparison is grounded on the evaluation vectors of both alternatives, and on additional information concerning the decision maker's preferences, typically accounting for two conditions: concordance and non-discordance. In decision processes using these methods, the analyst should interact with the decision maker in order to elicit values for the parameters that define a preference model. This can be done either directly or through a disaggregation procedure that infers parameter values from holistic judgements provided by the decision maker. In this chapter we discuss the elicitation of an outranking-based preference model, focusing on the valued outranking relation used in the ELECTRE III and ELECTRE TRI methods.

14.1 Introduction

As described in Chap. 12 in this book (Morton 2018), a common approach in the field of Multiple Criteria Decision Aiding (MCDA) is to aggregate the performances of an alternative being assessed on multiple criteria into a single number synthesizing its overall value (see also Keeney and Raiffa 1993). However, a different type of methods has been developed in parallel, which obtain a binary relation on the set of alternatives without aggregating multiple performances into a synthesis value. These methods are usually referred as outranking methods in the MCDA literature and have been, by and large, associated with the so-called European school of MCDA (see Roy and Vanderpooten 1996). This allows for decision aiding approaches able to model not only situations of preference or indifference between alternatives, but

L.C. Dias (✉)
Faculty of Economics, CeBER and INESC Coimbra, University of Coimbra,
Av. Dias da Silva 165, 3004-512, Coimbra, Portugal
e-mail: lmcdias@fe.uc.pt

V. Mousseau
Laboratoire Génie Industriel, CentraleSupélec, Châtenay-Malabry, France
e-mail: vincent.mousseau@centralesupelec.fr

L.C. Dias et al. (eds.), *Elicitation*, International Series in Operations Research
& Management Science 261, https://doi.org/10.1007/978-3-319-65052-4_14

also situations in which alternatives are deemed to be incomparable in the light of the preference information elicited. Incomparability typically occurs when the strengths and weaknesses of two alternatives are so different that one cannot conclude that one is better than the other, but it is equally unwarranted to conclude that they are indifferent (i.e., similarly preferred).

Outranking methods ground the recommendations to the Decision Maker (DM) on the construction of one (or several) binary relation(s) representing the preference among pairs of alternatives (see Roy 1991; Roy and Bouyssou 1993). A simple binary relation is dominance: an alternative dominates another one if it is better on some criteria and it is not worse in any other criterion. It does not require any subjective parameters such as criteria weights, but the relation is usually poor (i.e., it applies to few pairs of alternatives). Outranking methods use additional inputs to enrich this relation. Examples of outranking methods include ELECTRE methods (Figueira et al. 2013), PROMETHEE methods (Brans and Vincke 1985; Majid Behzadian et al. 2010), RUBIS (Bisdorff et al. 2007), NAIADE (Munda 1995), and qualitative approaches (Martel and Matarazzo 2005). This chapter will focus on preference elicitation for ELECTRE methods, but analogous procedures can be applied for other outranking-based approaches.

Let us consider a decision situation involving a finite set of alternatives $A = \{a_1, a_2, \ldots, a_l\}$ evaluated on n criteria $g_1, g_2, \ldots, g_n$, ($F = \{1, 2, \ldots, n\}$ denotes the set of criteria indices).

The construction of an outranking relation S amounts at validating or invalidating, for any pair of alternatives $(a, b) \in A^2$, an assertion aSb, whose meaning is "a is at least as good as b" or, in other words, "a is not worse than b". This comparison is grounded on the evaluation vectors of both alternatives a and b, *i.e.*, $(g_1(a), g_2(a), \ldots, g_n(a))$ and $(g_1(b), g_2(b), \ldots, g_n(b))$, and on additional information concerning the DM's preferences. To validate a statement aSb, two basic conditions should be verified: concordance and non-discordance (or non-veto).

A criterion g_k is said to be concordant with the assertion aSb if a is at least as good as b with respect to criterion g_k. The concordance condition is fulfilled for the assertion aSb when the subset of criteria concordant with aSb is a "sufficient majority". A criterion g_k is said to veto the assertion aSb if a is so much worse than b on this criterion that the difference of evaluation $|g_k(b) - g_k(a)|$ becomes incompatible with the assertion aSb, whatever the evaluation on the other criteria. The non-discordance condition is fulfilled when no criterion opposes a veto to the assertion aSb.

Constructing an outranking relation S involves the elicitation of values for preference-related parameters, such as weights, majority thresholds and veto thresholds. The next section provides details about these parameters and how they shape a model of the DM's preferences. Sections 14.3 and 14.4 in this chapter discuss how to elicit parameter values. The elicitation of preference-related parameters can be done either in a direct way centered on parameters (discussed in Sect. 14.3) or indirectly through a disaggregation procedure centered on examples, that infers the parameters values from holistic preferences provided by the DM (see Jacquet-Lagrèze and Siskos 2001) (discussed in Sect. 14.4). Inference is usually performed through an optimization program that accounts for the aggregation model and minimizes an “error function”. This disaggregation approach has been largely used in additive models (e.g. see Jacquet-Lagrèze and Siskos 2001 and Chap. 13 in this book Matsatsinis et al. 2018). Section 14.5 discusses the elicitation process, namely focusing on the order parameters are elicited and how precise should the elicitation be. Section 14.6 closes the chapter summarizing the main takeaways and highlighting the research challenges that still lie ahead.

14.2 Preference Models with ELECTRE

This section briefly presents the ELECTRE preference model, namely describing how a valued outranking relation on the set of alternatives is built in methods such as ELECTRE III (see Roy 1978) and ELECTRE TRI (see Yu 1992a,b; Roy and Bouyssou 1993).

14.2.1 Outranking Relations for a Single Criterion

ELECTRE builds, for each criterion g_j, a valued outranking relation S_j modelling the comparison of alternatives on that single criterion. For any ordered pair $(a, b) \in A^2$, $S_j(a, b)$ is defined by (14.2) on the basis of $g_j(a)$, $g_j(b)$ and two thresholds: indifference q_j and preference p_j ($0 \leq q_j \leq p_j$). We consider the thresholds p_j and q_j as constant, although it is possible to consider them as affine functions (for such cases see Roy et al. 2014). For a more compact notation, we will write:

$$\Delta_j(b, a) = g_j(b) - g_j(a), \tag{14.1}$$

which for each pair $(a, b) \in A^2$ represents the advantage of b over a on the jth criterion. This assumes, without loss of generality, that the evaluations are coded in such a way that the higher the value, the better it is (if this is not the case, one simply considers that $\Delta_j(b, a) = g_j(a) - g_j(b)$).

$S_j(a, b)$ represents the degree to which alternative a outranks (is at least as good as) b, defined as (Fig. 14.1):

$$S_j(a, b) = \begin{cases} 0, & \text{if } \Delta_j(b, a) > p_j \\ \frac{p_j - \Delta_j(b,a)}{p_j - q_j}, & \text{if } q_j < \Delta_j(b, a) \leq p_j \\ 1, & \text{if } \Delta_j(b, a) \leq q_j \end{cases} \tag{14.2}$$

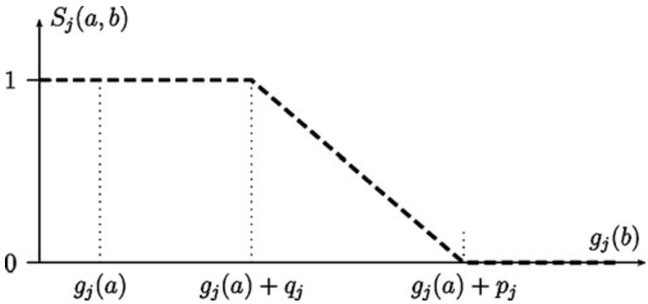

Fig. 14.1 Partial valued outranking relation

14.2.2 Concordance Relation

The valued concordance relation $C(a,b)$ aggregates the relations S_j $(j \in F)$, and it represents the level of majority among the criteria in favor of the assertion "*a* is at least as good as *b*". When computing this majority level, each criterion g_j has a weight $w_j \geq 0$ representing its voting power. Without any loss of generality, we will consider $\sum_{j\in F} w_j = 1$. Therefore, $C(a,b)$ can be written as follows:

$$C(a,b) = \sum_{j\in F} w_j S_j(a,b) \tag{14.3}$$

14.2.3 Discordance Relations

ELECTRE also builds, for each criterion g_j, a valued discordance relation d_j restricted to that criterion. This relation $d_j(a,b)$ is traditionally defined by (14.4) on the basis of $g_j(a)$, $g_j(b)$, a veto threshold v_j and a preference threshold p_j $(p_j < v_j$; note we consider $p_j < v_j$, although ELECTRE also allows $p_j = v_j)$ (see Fig. 14.2). We consider the thresholds v_j as constants (as we already did for p_j and q_j), although it is possible to consider them as affine functions.

$$d_j(a,b) = \begin{cases} 1, \quad \text{if } \Delta_j(b,a) \geq v_j \\ \frac{\Delta_j(b,a)-p_j}{v_j-p_j}, \quad \text{if } p_j < \Delta_j(b,a) < v_j \\ 0, \quad \text{if } \Delta_j(b,a) \leq p_j \end{cases} \tag{14.4}$$

The overall valued non-discordance relation $ND(a,b)$ as originally defined (Roy 1978) is grounded on $C(a,b)$ and on the relations d_j, $j \in F$; it represents the degree to which the minority criteria collectively oppose a veto to the assertion "*a* is at least as good as *b*". The classical way of defining $ND(a,b)$ is given in (14.5). $ND(a,b) = 0$ corresponds to a situation where some minority criteria are totally opposed to aSb

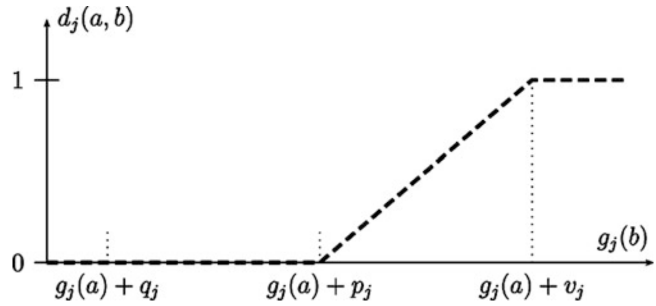

Fig. 14.2 Partial valued discordance relation

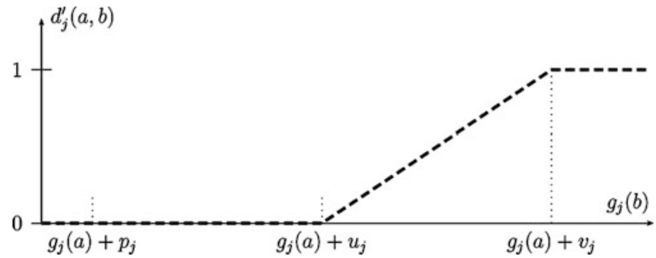

Fig. 14.3 Partial discordance relation $d'_j(a,b)$

whereas $ND(a,b) = 1$ means that none of the criteria oppose a veto to aSb.

$$ND(a,b) = \prod_{j \in \overline{F}} \frac{1 - d_j(a,b)}{1 - C(a,b)} \quad \text{where } \overline{F} = \{j \in F \text{ such that } d_j(a,b) > C(a,b)\} \tag{14.5}$$

This expression is equivalent to (14.6):

$$ND(a,b) = \prod_{j \in F} ND_j(a,b), \tag{14.6}$$

where:

$$ND_j(a,b) = Min \left\{ 1, \frac{1 - d_j(a,b)}{1 - C(a,b)} \right\}. \tag{14.7}$$

Mousseau and Dias (2004) have proposed an alternative valued non-discordance relation defined by (14.8)–(14.9), where $u_j \in [p_j, v_j[$ is a new parameter (discordance threshold) for the j-th criterion (Fig. 14.3):

$$ND'(a,b) = \prod_{j \in F} ND'_j(a,b) = \prod_{j \in F} \left(1 - d'_j(a,b)\right) \tag{14.8}$$

$$d'_j(a,b) = \begin{cases} 1 \quad \text{if } \Delta_j(b,a) \geq v_j \\ \frac{\Delta_j(b,a)-u_j}{v_j-u_j} \quad \text{if } u_j < \Delta_j(b,a) < v_j \\ 0 \quad \text{if } \Delta_j(b,a) \leq u_j \end{cases} \tag{14.9}$$

A second alternative to define a valued non-discordance relation is the following (see Mousseau and Dias 2004), which is simpler but only takes the highest discordance into account:

$$ND''(a,b) = Min_{j \in F} ND'_j(a,b) \tag{14.10}$$

Both definitions (14.8) and (14.10) follow ELECTRE's intention of allowing one minority criterion to veto the conclusion sustained by the majority of the criteria, if the performance difference is too large (and worse). These two definitions are mainly relevant when used in indirect elicitation (regression) processes, since they allow easier mathematical programming models to infer parameter values, especially variant $ND''(a,b)$.

14.2.4 Valued Outranking Relations

ELECTRE's valued outranking relation combines the concordance and non-discordance relations:

$$S(a,b) = C(a,b)\ ND(a,b), \tag{14.11}$$

or, according to the two alternative definitions of the discordance concept,

$$S'(a,b) = C(a,b)\ ND'(a,b) \tag{14.12}$$

$$S''(a,b) = C(a,b)\ ND''(a,b) \tag{14.13}$$

From a valued outranking relation such as $S(a,b)$, $S'(a,b)$ or $S''(a,b)$ it is possible to define a family of nested "crisp" outranking relations S_λ. These crisp relations correspond to λ-cuts of $S(a,b)$, where the cutting level $\lambda \in [0.5, 1]$ represents the minimum value for $S(a,b)$ so that $aS_\lambda b$ holds.

14.2.5 Exploitation of the Outranking Relation

Depending on the type of decision problem different ELECTRE methods can be applied. Roy (1996) identifies three main "problématiques" depending on the type of result sought:

- Selection (or choice): to identify the best alternative (or a predefined number of best alternatives) among a set of possibilities. Example: to select the best project among a set of possible variants.
- Ranking: to obtain a preference order among the alternatives, from best to worst. Example: a prioritization of projects defining the order by which they should be implemented.
- Sorting problems aim at assigning alternatives to categories, which are typically defined a priori and ordered. Example: sorting projects among the categories "not urgent", "urgent" and "very urgent".

Methods ELECTRE I and IS (Roy and Bouyssou 1993; Roy and Skalka 1984) have been proposed to deal with selection problems. Since the outranking relation is usually not transitive and not complete, often these methods are unable to identify a single winner. Their purpose is more modestly to identify a subset, named kernel, of candidates to be the most preferred alternative. The methods try to make this kernel as small as possible by excluding alternatives that are outranked. Alternatives in the kernel are incomparable, which typically means they are too different to be compared with the information requested by ELECTRE.

Methods ELECTRE II, III, and IV (Roy and Bertier 1973; Roy and Bouyssou 1993; Vallée and Zielniewicz 1994) have been proposed to deal with ranking problems. As in the case of choice, the lack of transitivity and incompleteness of the outranking relation hinder obtaining a clear-cut result. These methods yield a partial preorder as an output, i.e., an incompletely defined ranking allowing ties and in which some of the alternatives are incomparable.

Although ELECTRE methods for ranking and sorting have been used in many applications (Govindan and Jepsen 2016), the inconclusiveness of its results may disappoint some DMs and analysts. On the other hand, this inconclusiveness may be seen as a strength in that ELECTRE I-IV do not force the result to be more conclusive than warranted by the data and the preferences elicited. Another concern that has been much debated (e.g., Figueira and Roy 2009) is the fact that adding, removing, or modifying a possibly irrelevant alternative can change the relative position of the remaining alternatives. New ranking methods overcoming the latter issue have been proposed more recently (Rolland 2013).

Finally, ELECTRE TRI (Yu 1992b) and its variants are devoted to sorting problems. Since they do not compare the alternatives being evaluated against each other, adding, removing, or modifying an alternative has no effects of the results concerning the other alternatives. In ELECTRE TRI the alternatives are sorted based on how they compare to the profiles that define the available categories. These profiles are multidimensional preference vectors (each profile indicates one performance value for each criterion), which constitute new preference-related parameters to be elicited.

The original version of ELECTRE TRI (Yu 1992b) defined category profiles as bounds delimiting the categories: a profile b^1 separates category C^1 from category C^2 (b^1 can be considered a lower bound for C^2 and an upper bound for C^1); a profile

b^2 separates category C^2 from category C^3, and so on. This version is sometimes referred to as ELECTRE TRI-B.

A subsequent version of ELECTRE TRI is ELECTRE TRI-C (Almeida-Dias et al. 2010), which proposes to define profiles as central elements of the categories: a profile b^1 it the typical (characteristic) element of category C^1, a profile b^2 it the typical element of category C^2, and so on. Later, the extension ELECTRE TRI-nC was proposed to allow each category to be defined by more than one characteristic element (Almeida-Dias et al. 2012). By analogy, it is also possible to create an ELECTRE TRI-nB version (Fernández et al. 2017).

14.3 Direct Elicitation

14.3.1 Single-Criterion Concordance Parameters

It makes sense to start the process of eliciting an ELECTRE model by the single-criteria concordance parameters, since the parameters to be elicited afterwards are used in computations that refer to the relations S_j. Furthermore, the discussion about these parameters is not as cognitively demanding as for other parameters, and allows introducing the cornerstone concept of concordance in ELECTRE.

Given a pair of alternatives (a, b), $S_j(a, b)$ assesses the degree to which a outranks (is at least as good as) b according to the criterion g_j. According to (14.2), this depends on the advantage of a over b, denoted $\Delta_j(a, b) = -\Delta_j(b, a)$ and two parameters to be elicited: the indifference threshold q_j and the preference threshold p_j.

In the oldest ELECTRE methods (Roy 1968, 1971) the single-criterion concordance would be an absolute yes, i.e. $S_j(a, b) = 1$, if $\Delta_j(a, b) \geq 0$, or it would be an absolute no, i.e. $S_j(a, b) = 0$ otherwise. If a was worse than b on criterion g_j then there was no concordance at all, however small this difference might be. There are however some reasons why this model might be inadequate for some criteria:

- A small difference might be considered insignificant in relative terms concerning orders of magnitude. For instance, a difference of \$1 between two projects involving over \$1 million is not likely to be valued by any DM.
- Performances may be assessed in an imprecise way using measuring instruments or statistics. If the performance of a is 100 ± 5 and the performance of b is 101 ± 5, many DMs will be indifferent between one or the other because the performance difference is much lower than the acknowledged imprecision.
- Performance assessed may be just an imperfect indicator (or even a proxy) of real-world performance. For instance, the advertised fuel consumption of a car corresponds to its behavior in an idealized circuit (e.g., the New European Driving Cycle). If the performance of a is $5.0\,\text{l}/100\,\text{km}$ and the performance of b is $4.9\,\text{l}/100\,\text{km}$, many DMs will be indifferent between them because they know that none of these values correspond to real-life performance. Likewise, when

recruiting a college graduate, the DM knows that the grade point average (GPA) of their degrees is just an imperfect indicator: a student with a GPA of 3.5 is not guaranteed to be more knowledgeable than another one with a GPA of 3.4.

The ELECTRE methods introduced later acknowledge these situations by allowing the DM to set an indifference threshold q_j, which is the largest difference such that the DM does not distinguish between two alternatives in terms of preference. A question addressing the need for such a threshold can be the following:

"If the difference between two alternatives on criterion j is not equal to zero then one of them must be preferred on that criterion, or can this difference be so small that you would not distinguish them in terms of preference?"

In the latter case, it is possible to ask for a limit to this indifference situation:

"How large can a performance difference be until you start hesitating about the indifference between two alternatives?"

The DM can reply in absolute terms, e.g., 2.0, or in relative terms, e.g., 2%. Although most ELECTRE software implementations allow modelling $q_j(g_j(a))$ as an affine function $\alpha_j + \beta_j g_j(a)$ for some parameters α_j and β_j, typically this option is not used and this function is either a constant value ($\beta_j = 0$) or a proportion of the performance ($\alpha_j = 0$). For simplicity, in the remainder of this text we assume it is a constant value. When thresholds are modelled as functions of the performance levels special care must be taken to ensure their consistency (Roy et al. 2014).

It is possible to ask verification questions and adjust the parameter by trial and error:

"If $g_j(a)$ has value x_a and $g_j(b)$ has value x_b (for some relatively close values x_a and x_b) would you say that on criterion j the two alternatives are indifferent, or would you have a clear preference?"

It is not uncommon that up to a difference δ_1 the DM feels clearly indifferent, for a difference larger than δ_2 ($\delta_2 > \delta_1$) the DM has a clear preference, and for differences in-between δ_1 and δ_2 the DM exhibits some hesitation in answering such a question. This allows setting $q_j = \delta_1$ and $p_j = \delta_2$, since the preference threshold p_j corresponds to the minimum difference such that the DM has a clear preference for one of the alternatives.

The elicitation of q_j and p_j can therefore result in one of these typical situations:

- $p_j > q_j > 0$, meaning that some differences are too small to warrant preference, and that up to a certain point there is a clear indifference, then some hesitation, and finally a clear preference as the difference in performance increases.
- $p_j > q_j = 0$, meaning that above a given threshold there is a clear preference, and below this threshold the DM hesitates if the alternatives are indifferent or one is better than the other.
- $p_j = q_j$, and possibly both are null, in some cases concerning a discrete scale (e.g., number of rooms in a house, or number of stars of an hotel).

14.3.2 Weights and Cutting Level

As presented in Sect. 14.2.2, criteria weights are used to aggregate the concordance of the different criteria concerning an outranking relation. Although they are used in a weighted sum of concordance indices, they should not be interpreted as trade-off weights. Unlike a typical additive aggregation model (e.g., Keeney and Raiffa 1993) weights are not scaling coefficients such that the ratio of two weights indicates the conversion rate between units of value (or utility) on two different criteria.

An adequate analogy for eliciting weights in ELECTRE is that of voting. Suppose for the moment that all indifference and preference thresholds are null, i.e., the single-criterion concordances are either 0 or 1. Suppose also that there is no discordance (veto thresholds are not set or are extremely high), so that $S(a,b) = C(a,b)$ for any pair of alternatives (a,b). Then, a outranks b if the weights of the coalition of criteria that add up to $C(a,b)$ reaches at least cutting level λ. Then, the cutting level λ can be interpreted as representing the required majority for establishing an outranking relation. Typical values for this parameter are 0.50 or 0.51 (a simple majority), 0.67 (requiring a 2/3 majority), etc., up to 1 (requiring unanimity). A direct elicitation question could be:

"How strong must the majority of the criteria that agree that a is at least as good as b be, in order to establish this conclusion, taking criteria weights into account? (in the absence of strong discordance)"

In a trial-and-error process tentative symbolic majority levels can be suggested, such as 1/2, 2/3 or 3/4. Otherwise, communicating in terms of percentages is preferable to decimal numbers (i.e., 60% communicates better the sense of a required majority than 0.60). The higher the majority level required, the less will the number of outranking relations be but the stronger is their justification. Often, a compromise is sought between the richness of the relation (number of pairs for which outranking holds) and the strength of the justification, by observing the effects of varying this parameter.

In the particular case of sorting problems with ELECTRE TRI (ELECTRE TRI-B) it may be more appropriate to inquire about the cutting level in a way that matches more directly its effects on the results:

"How strong must the majority of the criteria that agree that a is at least as good as the lower profile of a category be, taking criteria weights into account, to warrant that an alternative can be sorted on that category, if not better? (in the absence of strong discordance)"

Indeed, to be sorted in a given category (if not better) an alternative must outrank the category's lower profile. This parameter can be interpreted as denoting how much demanding the decision maker is. A high cutting level makes it more difficult for the alternatives to be classified in the best categories. Again, symbolic majority levels can be tentatively suggested.

Having the voting majority analogy in mind, then criteria weights simply reflect how much they count in the formation of such majorities. This means that weights

Table 14.1 Example of criteria weights

Criteria:	g_1	g_2	g_3	g_4	g_5
Weights (w_k):	0.15	0.20	0.15	0.15	0.35

match the analogy of weights in the physical world. A direct elicitation question could be:

"Considering that the support of all criteria for an outranking relation amounts to a 100% majority (unanimity), how much weight (or voting power) would you assign to criterion g_j alone?"

Confirmation questions can be asked concerning the elicited weights. Consider for instance the weights in Table 14.1. Since $w_1 < w_2$ one should confirm that having the support of the first criterion for an outranking relation is less important than having the support of the second criterion. Since $w_1 + w_2 = w_5$, one should confirm that the last criterion counts as much as the other two criteria. These are just two examples among many possible. Further confirmatory questions may interrelate the elicited weights and the cutting level. For instance, if $\lambda = 0.55$, one should confirm that:

- No criterion alone is strong enough to warrant an outranking relation.
- The only coalition of two criteria strong enough to warrant an outranking relation is g_2 together with g_5 (since $w_2 + w_5 = 0.55 = \lambda$)
- No coalition of three criteria is a sufficient majority unless g_5 is in it.
- Any coalition of four criteria is a sufficient majority (at the minimum, $w_1 + w_2 + w_3 + w_4 = 0.65$, which is larger than λ).

If indifference and preference thresholds are not null, the single-criterion concordances can be any value between 0 or 1, but this does not change the logic of the elicitation process. One simply has to reason that if, for instance, the performance of alternatives a and b is such that $S_j(a, b) = 0.50$, then criterion g_j contributes with half of its weight to the coalition supporting that a outranks b.

An alternative to directly asking for numerical criteria weights has been proposed by Simos (1990) and later revised by Figueira and Roy (2002). DMs can use cards with criteria names to indicate how they would rank the criteria by order of importance. Two or more cards can be placed together to indicate the respective criteria should have the same weight. In addition, DMs can place blank cards to indicate a higher difference in weights between ranks. For instance, DMs could indicate the following ranking: g_1 and g_2, g_3, (blank), g_4, (blank), (blank), g_5. This indicates that g_1 and g_2 are the two criteria with higher weight, followed by g_3, then g_4 and finally g_5. The blank cards in this example entail that one should have $w_4 - w_5 = 3 * (w_2 - w_3)$ and that $w_3 - w_4 = 2 * (w_2 - w_3)$. Since there are many weight vectors fulfilling these inequalities the revised Simos technique requires DMs to set a ratio between the first and the last ranked weights. The authors also propose a rounding technique if the resulting weights are required to have a predefined maximum number of decimal digits (for details see Roy and Figueira 2002).

14.3.3 Discordance Parameters

14.3.3.1 Parameters Defining $d_j(a,b)$

Being a noncompensatory preference model, ELECTRE allows specifying that a large disadvantage on one criterion may not be compensated by advantages on other criteria. Let us recall the way the non-discordance condition is implemented through $ND(a,b)$ as in Eq. (14.5). If $g_j(b) - g_j(a)$ exceeds v_j for at least one criterion then aSb is invalidated, *i.e.*, $\exists j \in F : d_j(a,b) = 1 \Rightarrow S(a,b) = 0$. This may happen even when the total concordance $C(a,b)$ is higher than the cutting level λ.

Traditionally, $ND(a,b)$ accounts both for the values of $d_j(a,b)$ and $C(a,b)$: the way $ND(a,b)$ accounts for $d_j(a,b)$ is amplified when $C(a,b)$ is low. This reflects the idea that a veto situation should be accentuated when the concordance relation is not firmly established. On the other hand, if $C(a,b)$ is high, then low values of $d_j(a,b)$ are not taken into account: the overall non-discordance relation defined in (14.5) considers the $d_j(a,b)$ only for criteria such that $d_j(a,b) > C(a,b)$.

The interplay between $d_j(a,b)$ and $C(a,b)$ in measuring discordance makes the process of eliciting veto thresholds v_j prone to misunderstandings. The typical question asked is often:

"What would be a performance difference in criterion j so large that it cannot be compensated, i.e., that would make this criterion oppose a veto to any concordant majority of other criteria?"

Suppose for instance that the previous steps of the elicitation process had let to set $p_j = 10$ and $k_j = 0.20$, for some $j \in F$. Suppose also that the answer to the previous question had led to set $v_j = 50$, possibly by "trial and error". The DM was found to have the opinion that if the performance difference is equal to 50 units or more, then there would be a veto, but if the difference was less than 50 then an outranking would be allowed. However, in this case any difference higher than 45 would necessarily veto an outranking relation:

From (14.4), $p_j = 10$, $v_j = 50$, and $\Delta_j(b,a) > 45$ imply $d_j(a,b) > 7/8$.

Even assuming that there is no other discordance and $C(a,b) = 1 - k_j = 0.8$, Eq. (14.7) together with $d_j(a,b) > 7/8$ yield $ND_j(a,b) < 0.625$.

Finally, Eqs. (14.6) and (14.11) yield $S(a,b) < 0.5$.

Since λ, the required majority, is at least 0.5, a cannot outrank b.

This means that the traditional question for eliciting a difference large enough to warrant a veto situation leads to an overestimation of this difference. A more rigorous way to question about the veto threshold, provided that criteria weights have been elicited, is the following (assuming the parameter values of this example):

"Suppose that j is the only discordant criterion, meaning that a coalition of 80% of the criteria weights agrees that aSb. What would be a performance difference in criterion j so large that it cannot be compensated, i.e., that would make this criterion oppose a veto to that coalition, even if λ was as low as 0.5?"

It the DM provided the same answer, 50, then to obtain the desired behavior it would be necessary to set:

$$v_j = p_j + \frac{C(a,b)(\Delta_j(b,a) - p_j)}{C(a,b) - 0.5(1 - C(a,b))} = 10 + \frac{0.8(50 - 10)}{0.8 - 0.5(1 - 0.8)} = 55.71429. \tag{14.14}$$

14.3.3.2 Parameters Defining $d'_j(a,b)$ for Relation $S'(a,b)$ or $S''(a,b)$

The indices $d'_j(a,b)$ are defined by (14.9) on the basis of $g_j(a)$, $g_j(b)$, a veto threshold v_j and an additional threshold u_j which we call *discordance threshold*. u_j represents the difference of evaluation $g_j(b) - g_j(a)$ above which the discordance condition starts to weaken concordance $C(a,b)$ in the definition of $S'(a,b)$. This discordance threshold u_j can be considered either:

- as an additional preferential parameter to be elicited through an interaction with the DM, or
- as a technical parameter (rather than a preference-related one), an option that should be used only when the DM does not wish to use the added flexibility offered by u_j, preferring to work with the thresholds v_j only. In such cases, a reasonable "rule-of-thumb" is to set $u_j = p_j + 0.75(v_j - p_j)$ (see Mousseau and Dias 2004).

In case the discordance threshold u_j is to be elicited, then the main difference in the use of relation $S'(a,b)$ rather than $S(a,b)$ is that criteria that intervene in the product are not restricted to those for which $d'_j(a,b) > C(a,b)$, *i.e.*, small values of $d'_j(a,b)$ will impact $ND'(a,b)$. Moreover, the concordance relation $C(a,b)$ does not intervene in the non-discordance implementation.

In model $S'(a,b)$, the discordance $d'_j(a,b)$ corresponds to a correction factor to the concordance of all other criteria taken together. One possibility is to ask two questions defining the performance differences that correspond to two distinct $d'_j(a,b)$ values, e.g., a 10% correction (decrease) and 25% correction. For the former case the question would be (the question pertaining the latter is similar):

"Suppose that j is the only discordant criterion, meaning that all other criteria agree that aSb. What would be a performance difference in criterion j that would warrant decreasing the weight of all concordant criteria by 10%?" (Note that unlike relation S there is no need to refer to the exact weight of the criteria).

If, for instance, the DM would state that $\Delta_j(b,a) = 40$ warrants decreasing the weight of all concordant criteria by 10% and $\Delta_j(b,a) = 50$ warrants decreasing the weight of all concordant criteria by 25% then, based on Eq. (14.9), solving the system

$$\begin{cases} \frac{40-u_j}{v_j-u_j} = 0.10 \\ \frac{50-u_j}{v_j-u_j} = 0.25 \end{cases} \tag{14.15}$$

leads to the solution $u_j = 100/3$ and $v_j = 100$.

It is also possible to ask only one of the above questions, and use a different question to elicit u_j:

"Suppose that j is the only discordant criterion, meaning that all other criteria agree that aSb. At what point (performance difference) would a veto effect start to occur, in that the weight of all concordant criteria would start to be decreased?"

If, for instance, the DM would reply that a veto effect gradually begins at a difference of 40, and that $\Delta_j(b,a) = 50$ warrants decreasing the weight of all concordant criteria by 25% then, based on Eq. (14.9),

$$\frac{50-40}{v_j-40} = 0.25 \text{ yields } v_j = 95. \tag{14.16}$$

14.3.4 Profiles in Sorting Problems

The elicitation of profiles in sorting problems in the framework of ELECTRE models must take into account their distinct nature in different variants of ELECTRE TRI: in the original version (ELECTRE TRI-B) the profiles are limits separating the consecutive categories, whereas in ELECTRE TRI-C the profiles are central elements of the categories.

Let us first address the original version (ELECTRE TRI-B). Here, a profile b^k separates category C^k from category C^{k+1} (it can be considered a lower bound for C^{k+1}). If there are n_{cat} categories, then $n_{cat}-1$ profiles need to be elicited. A lower bound for the first (worst) category, b^0, needs not be elicited by assuming that aSb^0 is true for every conceivable alternative a. Similarly, An upper bound for the last (best) category, $b^{n_{cat}}$, needs not be elicited by assuming that $aSb^{n_{cat}}$ is false and $b^{n_{cat}}Sa$ is true for every conceivable alternative a.

Considering the convention that C^1 is the worst category and $C^{n_{cat}}$ is the best category the following conditions should be ensured:

- Each profile dominates the profiles of lower categories: if $k' > k$ then $g_j(b^{k'}) \geq g_j(b^k)$ for criteria g_j to be maximized and $g_j(b^{k'}) \leq g_j(b^k)$ for criteria g_j to be maximized, with at least one of these inequalities being strict.
- Profiles should not be so close to each other that an alternative might be indifferent to both: for two different profiles $b^{k'}$ and b^k there is no alternative a such that $aSb^{k'}$ and $b^{k'}Sa$ and at the same time aSb^k and b^kSa.

The sorting of alternatives in ELECTRE TRI can be performed according to a pessimistic (pseudo-conjunctive) perspective or an optimistic (pseudo-disjunctive) perspective. Whenever the alternative to be sorted is incomparable to some profiles, the pessimistic perspective places it a lower category than the optimistic perspective; otherwise, both perspectives sort it in the same category. In this chapter we will consider the pessimistic perspective, according to which an alternative is sorted in a category C^k if it is good enough to outrank its lower bound but not good enough to outrank its upper bound:

$$a_i \text{ is sorted in } C^k \Leftrightarrow a_i S b^{k-1} \wedge \neg a_i S b^k \tag{14.17}$$

Elicitation of the profiles can be conducted by considering one criterion at a time. For each criterion g_j, the profiles for the successive categories can be asked in ascending order (starting from the worst one) or in descending order (starting from the best one). In descending order the performance value for $g_j(b^{n_{cat}-1})$ can be asked as follows:

"On criterion g_j, what level of performance is required for this criterion to vote in favor of sorting an alternative in the best category, $C^{n_{cat}}$?"

Then, the performance value for $g_j(b^{n_{cat}-2})$ can be asked as follows:

"On criterion g_j, what level of performance is required for this criterion to vote in favor of sorting an alternative in category $C^{n_{cat}-1}$? (if not better)"

Then, performances $g_j(b^{n_{cat}-3})$,…,$g_j(b^1)$ would be elicited in the same way, before moving on to a different criterion. Focusing on one criterion at a time makes the task easier for decision makers, who are in this way invited to consider how each criterion would sort the alternatives, if there was not any other criterion.

As an alternative, the elicitation can focus on one category at a time, considering all the criteria, but often this task is harder. Decision makers would have to provide multi-criteria performances for a profile $b^{n_{cat}-1}$ such that all alternatives outranking it would be placed in the best category. Then, they would need to provide multi-criteria performances for a profile $b^{n_{cat}-2}$ such that all alternatives outranking it (but not outranking $b^{n_{cat}-1}$) would be placed in the second best category, and so on.

Let us now address the central profiles version ELECTRE TRI-C. Here, a profile b^k is the most representative (also called characteristic) element of category C^k. If there are n_{cat} categories, then n_{cat} profiles need to be elicited. For these profiles to be consistent, a profile for one category, say b^k, cannot be better than the profile b^{k+1} from a better category. At the minimum, $S(b^k, b^{k+1}) < 1$, but more stringent conditions such as $S(b^k, b^{k+1}) < 0.5$ or $S(b^k, b^{k+1}) < 0$ can be placed (Almeida-Dias et al. 2010). The basic idea of this method is to sort each alternative to the category such that the alternative outranks and is at the same time outranked by the profile as much as possible, i.e., with the largest $min\{S(a_i, b^k), S(b^k, a_i)\}$ (for details, see Almeida-Dias et al. 2010).

As in the case of ELECTRE TRI-B, elicitation of the profiles can be conducted by considering one criterion at a time. For each criterion g_j, the profiles for the different categories can be asked, in any order. The performance value for $g_j(b^k)$ can be asked as follows:

"On criterion g_j, what level of performance best characterizes an alternative in category, C^k?"

As an alternative, the elicitation can focus on one category at a time, considering all the criteria. In this case, a profile can be regarded as an ideal example characterizing the sort of performances the decision maker associates with each category. Method ELECTRE TRI-nC (Almeida-Dias et al. 2012), which extends ELECTRE TRI-C, even allows the decision maker to provide different examples of profiles to characterize each category.

14.4 Indirect Elicitation (Regression)

Assigning values to the parameters involved in the definition of an ELECTRE model might be a difficult task for the DM. The disaggregation approach (see Jacquet-Lagrèze and Siskos 2001) allows to infer parameter values from holistic preferences (i.e., global preferences rather than a criterion-by-criterion analysis). Holistic statements might be a ranking of a set of alternatives, comparisons of alternatives, or, in the case of sorting problems, the proposal of classification examples. The alternatives that are compared in a holistic manner might be a small subset of a much larger set of alternatives to be evaluated, or alternatives considered in past decision processes (possibly knowing how well they performed previously), or even examples constructed in a way that facilitates comparisons.

The disaggregation approach is usually performed using mathematical programs. Such inference programs can either be partial if only a subset of parameters is being inferred (the values of the other parameters being fixed), or global if all parameters are to be inferred. The inputs for the mathematical program are the holistic preference statements and the values of the parameters that are not being elicited. The decision variables are the parameters to be inferred. The objective function is to minimize an "error function" measuring how well the holistic preferences are reproduced by the inferred model. The constraints reflect the holistic preferences and also constraints that the method imposes on the model (e.g., weights are nonnegative and they add up to 1).

As described in Sect. 14.2.5, in ELECTRE methods the final choice set, or ranking, or sorting result is derived from the outranking relation. For ELECTRE TRI's pessimistic (or pseudo-conjunctive) variant, a statement in the form of a sorting example can be translated in two statements concerning outranking relations (Mousseau and Dias 2004). For instance, a statement "alternative a should be classified at least in the second category and at most in the third category" is translated into two outranking statements: "a outranks the lower profile of the second category" and "a does not outrank the lower profile of the fourth category". Unfortunately, statements based on ELECTRE methods devoted to choice or ranking problems, such as ELECTRE I-IV, do not have an easy direct translation into outranking statements. Therefore, the literature has concentrated on the cases of sorting problems or inferring parameters from outranking statements.

In order to elicit values for preference-related parameters (*i.e.*, w_j, $v_j(g_j)$, $p_j(g_j)$, $q_j(g_j)$, and limits of categories in ELECTRE TRI) it is possible to proceed using a disaggregation procedure that infers the parameters values from holistic preferences provided by the DM. Hence, it is necessary to formalize $S(a,b)$ through an optimization program that minimizes an "error function" that measures how much the values of the inferred parameters contradict the stated holistic preferences. However, $S(a,b)$ is rather "optimization unfriendly". Difficulties arise mainly from the way the non-discordance condition is implemented, *i.e.*, the way $ND(a,b)$ is defined.

More precisely, two features of the non-discordance relation are concerned. First, the subset of criteria $\overline{F}$ (see (14.5)) is difficult to integrate into an optimization program. Second, the fact that $C(a,b)$ intervenes in the definition of $ND(a,b)$ implies that the optimization program will necessarily be non-linear, even when all the parameters are fixed except the weights.

The problem of inferring the parameters of an ELECTRE method (ELECTRE TRI) based sorting examples translated into outranking statements was initially studied by Mousseau and Slowinski (1998). The resulting mathematical programming was nonlinear and would require global optimization techniques to find a solution. A simpler formulation was proposed to infer only the weights and the cutting level in situations without veto thresholds, in which $S(a,b) = C(a,b)$. In this case, an easy to solve linear programming formulation could be devised.

If veto thresholds are allowed, then the problem can no longer be solved by linear programming, even if the only parameters to be inferred are weights and the cutting level. Indeed, $S(a,b)$ is a non-differentiable and quasi-concave nonlinear function of the weights in the domain where it is strictly positive and therefore a constraint like $S(a,b) < \lambda$ (which reflects a holistic statement of the form $\neg aSb$) does not define a convex set (Dias and Climaco 1999). For this reason, Mousseau and Dias (2004) proposed variants for the outranking relation, S' and S'' (presented in Sect. 14.2.4) that allow using linear programming in such cases.

To provide an example of the mathematical programming approach to inference, the following section briefly recalls the inference of weights and cutting level for relation S'. The ensuing section overviews the literature on eliciting other subsets of parameters.

14.4.1 Inferring Weights and Cutting Level from S' Outranking Statements

Let us suppose that the DM is not able (or not willing) to assign directly values to the preference-related parameters involved in the outranking relation, but can state crisp statements about this relation for some specific pairs of alternatives (a,b), *i.e.*, either aSb (a outranks b) or $\neg aSb$ (a does not outrank b). Our purpose is to find criteria weights and a cutting level that restore the DM's statements.

Let A denote a set of alternatives. Let $S^+ = \{(a,b) \in A^2$ such that the DM stated $aSb\}$ and $S^- = \{(a,b) \in A^2$ such that the DM stated $\neg aSb\}$. Then, a combination of parameter values is able to restore the DM's request iff $S(a,b) \geq \lambda,\ \forall (a,b) \in S^+$ and $S(a,b) < \lambda,\ \forall (a,b) \in S^-$, which may be written as $S(a,b) - \lambda \geq 0,\ \forall (a,b) \in S^+$ and $\lambda - S(a,b) + \varepsilon \geq 0,\ \forall (a,b) \in S^-$ (ε being a small positive value).

The mathematical program given below (14.18)–(14.23) maximizes a common slack α for all these constraints, to obtain a relatively "central" combination of parameter values. Whenever the optimum value of α is negative, there is no combination of parameter values complying to all the constraints, *i.e.*, the DM provided inconsistent information (a procedure to deal with such inconsistencies is proposed in Mousseau et al. 2003). Alternative objective functions can be considered (see Beuthe and Scannella 2001 and Mousseau and Slowinski 1998).

$$Max \;\; \alpha \tag{14.18}$$

$$s.t. \; \alpha \leq S(a,b) - \lambda, \quad \forall (a,b) \in S^{+} \tag{14.19}$$

$$\alpha \leq \lambda - S(a,b) + \varepsilon, \quad \forall (a,b) \in S^{-} \tag{14.20}$$

$$\lambda \in [0.5, 1] \tag{14.21}$$

$$v_j(g_j) > p_j(g_j) > q_j(g_j) \geq 0, \quad \forall j \in F \tag{14.22}$$

$$\sum\nolimits_{j=1}^{n} w_j = 1; \quad w_j \geq \varepsilon, \; \forall j \in F. \tag{14.23}$$

Some additional constraints can be added to this program, in order to integrate explicit statements of the DM concerning the values of some parameters. From (14.5) and (14.11), it is obvious that this is a difficult nonlinear program if all the parameters were considered as variables. A solution to circumvent this difficulty is to formulate partial inference programs, where only a subset of the parameters are considered as variables, while the remaining ones are elicited by other means. Among the partial inference problems, previous research on related problems has focused mainly on inferring the weights and the cutting level (see Mousseau et al. 2000; Dias et al. 2002; Miettinen and Salminen 1999). This is an important partial inference problem because the weights and the cutting level are the only parameters involving inter-criteria judgements (the remaining parameters do not interrelate the criteria).

Let us consider the case where S' is used. In this case each product $\prod_{j \in F}(1 - d'_j(a,b)) = ND'(a,b)$ is a fixed constant $\forall (a,b)$. The following constraints concerning outranking statements are hence linear, since $C(a,b)$ is an affine function of the weights.

$$\alpha \leq C(a,b) \prod_{j \in F} \left(1 - d'_j(a,b)\right) - \lambda, \quad \forall (a,b) \in S^{+} \tag{14.24}$$

$$\alpha \leq \lambda - C(a,b) \prod_{j \in F} \left(1 - d'_j(a,b)\right) + \varepsilon, \quad \forall (a,b) \in S^{-} \tag{14.25}$$

Considering $S'(a,b)$ instead of $S(a,b)$, the weights and the cutting level can be inferred by solving a linear program whose variables are α, $w_1, \ldots, w_n$, and λ, where (14.24) and (14.25) appear as (14.27) and (14.28):

$$Max \;\; \alpha \tag{14.26}$$

$$s.t. \; \alpha \leq \sum\nolimits_{j=1}^{n} w_j \, S_j(a,b) \, ND'(a,b) - \lambda, \quad \forall (a,b) \in S^{+} \tag{14.27}$$

$$\boldsymbol{\alpha} \leq \lambda - \sum_{j=1}^{n} w_j \, S_j(a,b) \, ND'(a,b) + \varepsilon, \quad \forall (a,b) \in S^- \tag{14.28}$$

$$\lambda \in [0.5, 1], \tag{14.29}$$

$$\sum_{j=1}^{n} w_j = 1 \quad w_j \geq \varepsilon, \ \forall j \in F \tag{14.30}$$

If the maximum value of α is positive, then the values of $w_1, \ldots, w_n$, and λ at the optimum are able to restore all the statements defining S^+ and S^-. Otherwise, the inferred values provide suggestions for changing those examples. The DM should ponder whether they want to change the sets S^+ and S^-, or to analyze the values of $ND'(a,b)$. Indeed, some of the differences among the current model and the DM's requests may stem from inadequate values for the veto and discordance thresholds. Considering $S''(a,b)$ instead of $S'(a,b)$ leads to a similar linear program.

As a particular case, the pessimistic procedure of ELECTRE TRI assigns alternative a to category C_h (b_{h-1} and b_h being the lower and upper profiles of C_h, respectively) iff $S(a, b_{h-1}) \geq \lambda$ and $S(a, b_h) < \lambda$ ($\lambda \in [0.5, 1]$ is the chosen cutting level).

Suppose the DM has specified a set of assignment examples, *i.e.*, a subset of $A^* \subset A$ such that each $a_k \in A^*$ is associated with $C^M(a_k)$ ($C^m(a_k)$, respectively) the maximum (minimum, respectively) category to which a should be assigned according to his/her holistic preferences. Hence $[C^m(a_k), C^M(a_k)]$ defines an interval of possible categories to which a_k can be assigned to. $C^m(a_k) = C^M(a_k) = C_{h_k}$ means that the DM wants a_k to be assigned to C_{h_k} precisely (we will note $a_k \rightarrow_{DM} C_{h_k}$ such statement), while $C^m(a_k) < C^M(a_k)$ corresponds to an imprecise statement ($a_k \rightarrow_{DM} [C^m(a_k), C^M(a_k)]$).

Inferring all ELECTRE TRI parameters is a difficult nonlinear program (Mousseau and Slowinski 1998). But if we consider $S'(a,b)$ instead of $S(a,b)$, the weights and the cutting level can be inferred by solving a linear program (all other parameters being given as inputs). The linear program for this partial inference problem is equal to (14.26)–(14.30) if we define:

$$S^+ = \left\{(a_k, b_{C^m(a_k)-1}) \in A^* \times B \ : \ a_k \rightarrow_{DM} [C^m(a_k), C^M(a_k)]\right\} \tag{14.31}$$

$$S^- = \left\{(a_k, b_{C^M(a_k)}) \in A^* \times B \ : \ a_k \rightarrow_{DM} [C^m(a_k), C^M(a_k)]\right\} \tag{14.32}$$

Considering $S''(a,b)$ instead of $S'(a,b)$ leads to a similar linear program.

14.4.2 Inferring Different Parameters for Sorting Problems

In recent years, several papers dealt with the learning of ELECTRE TRI parameters.

As mentioned previously, the first paper devoted to the learning of ELECTRE TRI parameters has been proposed by Mousseau and Slowinski (1998). The learning algorithm takes as input a set of assignment examples and their associated vector of performances with respect to the problem criteria. The paper shows the difficulties to learn the parameters of ELECTRE TRI without veto. The main difficulty is the non-linearity of the partial concordance indices. Indeed, it makes the concordance index not differentiable which prevents the use of gradient optimization algorithms. In order to tackle this difficulty, Mousseau and Slowinski (1998) propose to approximate the partial concordance indices by sigmoid functions.

Learning all the parameters of an ELECTRE TRI model involves the determination of a lot of parameters. It requires a lot of cognitive effort from the user. Mousseau et al. (2001) consider the subproblem of finding the weights and the cutting threshold of an ELECTRE TRI model with fixed profiles and indifference and preference thresholds. In the paper, a linear program is proposed and some experiments are conducted. It shows that learning only a subpart of the ELECTRE TRI model simplifies the problem. Fewer assignment examples are required to obtain good results.

Ngo The and Mousseau (2002) proposed a mixed integer program in order to infer the profiles of an ELECTRE TRI model with fixed weights and thresholds. The mixed integer program presented in the paper finds the partial concordance indices in a first step. The second step consists in deducing the values of the profiles from the partial concordance indices. They propose to use this mixed integer program in combination with the linear program of Mousseau et al. (2001) in order to determine the whole set of parameters of an ELECTRE TRI model.

Mousseau and Slowinski (1998), Mousseau et al. (2001) and Ngo The and Mousseau (2002) consider only ELECTRE TRI models without veto. Dias and Mousseau (2006) present a manner to learn vetoes of an ELECTRE TRI model with fixed profiles, thresholds and weights. In the paper, two subproblems are treated. The first one considers the inference of veto parameters for a single criterion. The second considers the inference of all veto parameters for multiple criteria at the same time.

Doumpos et al. (2009) proposed a metaheuristic in order to learn all the parameters of an ELECTRE TRI model, including the veto thresholds. They developed a genetic algorithm in order to learn all the parameters of the model at the same time. The interest of this approach is that it allows to deal with larger data sets than mixed integer program based algorithms.

However ELECTRE TRI integrates a large number of preference parameters that are to be determined. MR-SORT is a simplified version of ELECTRE TRI which keeps the philosophy of ELECTRE TRI with the advantage of using less parameters (no veto thresholds and no discrimination thresholds are considered). Leroy et al. (2011) propose a mixed integer program in order to learn the parameters of such a model based on assignment examples. The experimental results presented in the paper show that the mixed integer program is able to find MR-SORT models which perform well in generalization. However, the experiments show the limitation of such an algorithm in terms of computing time. For a small problem involving five categories and three criteria, more than 100 s are required to restore all the parameters of a MR-SORT model on the basis of 100 assignment examples.

Damart et al. (2007) are the first to consider the problem of learning the parameters of an ELECTRE TRI model in the context of multiple decision makers. They propose an approach that aims at determining a set of fictitious alternatives that contain enough information to obtain a model that is satisfactory for all the DMs. The procedure is applied to an illustrative example.

Later, Cailloux et al. (2012) developed two mixed integer programs in order to learn the parameters of a MR-SORT model in the context of multiple DMs. The first mixed integer program aims at finding a set of profiles that is common to all the decision makers. The second mixed integer program learns a set of weights compatible with the preferences of each DM. The paper presents experimental results on real and fictitious applications.

Recently, Sobrie et al. (2013), Sobrie (2016) proposed an heuristic to efficiently infer MR-SORT parameters (weights and profiles) from large sets of assignment examples (over several thousands).

14.5 Elicitation Process

After reviewing elicitation techniques, we now focus on elicitation as a process that evolves in time, involving at least one DM and an analyst conducting the process. Two issues are discussed: elicitation sequence and numerical precision.

14.5.1 Elicitation Sequence

The elicitation sequence defines which parameters are elicited, in which order (or simultaneously), and using which technique.

All the parameters of an ELECTRE model should be discussed with the DMs, but not necessarily elicited from them. Indeed, there are at least three situations in which some parameters are not elicited:

- Indifference and preference thresholds, unlike preference-based parameters such as weights, may be considered technical parameters (Rogers and Bruen 1998; Roy et al. 2014) that can be set by the analyst, possibly with the help of experts on the domain that a criterion refers to. For instance, an analyst may set both thresholds equal to zero if a scale is ordinal, or an expert may set these thresholds based on considerations about the method that measures the performance of the alternatives on a cardinal scale, or a scientist may inform which differences in, say, noise levels, are negligible because a human cannot perceive them (Rogers and Bruen 1998).
- Veto thresholds may not be necessary, at least for all the criteria. The DMs may deem that no veto power is granted to some criteria, meaning that the discordance from those criteria is always null.

- The DMs may feel uncomfortable about setting criteria weights. In such cases, they may resort to ELECTRE IV, a method that does not ask for weights (Roy and Hugonnard 1982), or they may consider some freedom in setting the weights, as discussed in the following section. Many DMs may simply ask that all criteria have the same weight, but such a conclusion should result from (or be confirmed by) elicitation questions (Sect. 14.3.2).

There is no mandatory order by which parameters should be elicited. A possible sequence is the one followed by Sect. 14.3. Indifference and preference thresholds are clearly related and thus should be elicited simultaneously, one criterion at a time. Then, since the concordance part of the outranking relation is being addressed, the elicitation of weights may ensue. If the cutting level λ is communicated as a required majority level, then this parameter can be discussed simultaneously with weights, as described in Sect. 14.3.2. Finally, the possibility of veto is discussed, eliciting veto and non-discordance thresholds.

A different strategy is to initially focus on one criterion at a time and elicit indifference, preference, discordance and veto thresholds for each criterion. Then, criteria weights and the cutting level, which interrelate multiple criteria, would be elicited.

When an indirect elicitation (regression approach) is followed, multiple types of parameters can be inferred simultaneously, although that is a difficult optimization problem. Inferring only a subset of the parameters at a time allows overcoming this difficulty, and has an additional advantage. Since the DMs interactively revise the information they provide and observe the results of the mathematical program, partial inference problems allow them to focus their attention on a subset of parameters at a time and to better understand the consequences of modifying the examples they provide. We believe that inference programs should not be considered as a problem to be solved once, but rather as problems to be solved many times throughout an interactive learning process. Furthermore, it is possible to mix direct and indirect elicitation questions for different sets of parameters, and even for the same parameters (for confirmation purposes). Finally, the notion that parameters are elicited in a sequence does not mean that the elicitation process is linear. Often, the analyst may find out that the discussion concerning a subset of parameters puts into question the values elicited previously for another subset of parameters.

14.5.2 Numerical Precision

The issue of precision (and accuracy) arises in both direct and indirect elicitation. By precision we mean the freedom of variation one accepts for a parameter. For instance, setting the weight of the first criterion as $w_1 = 0.288$ is more precise than setting $w_1 \in [0.28, 0.29]$, which is more precise than setting or $w_1 \in [0.25, 0.30]$. The elicitation process is developed during a finite time window in which the DMs are available (and attentive!). Therefore, one has to accept the elicitation results might not be "accurate" in the sense that they include the exact parameter values that would result from a much longer process. In a direct elicitation process, a DM would hardly state that $w_1 = 0.288$. Probably he or she would state 0.29 or 0.3 which are "rounder" numbers. Typically these inputs are accepted even knowing they might be slightly inaccurate: no analyst would ask if it should really be a value of 0.299 or 0.301 instead of 0.3. Analysts know that rounder numbers are more comfortable for the DMs and reckon it would not be worthwhile to trouble a DM for a degree of precision that might be irrelevant to the results of the analysis. These concerns can be addressed at the end by means of a sensitivity or a robustness analysis (Roy 1998).

In indirect elicitation processes the mathematical programs might admit many different solutions able to reproduce the examples provided by DMs. For instance, experimental studies have been developed (Mousseau et al. 2001) showing that to infer relevant values for w_j and λ, the cardinality of S^+ and S^- should be "sufficiently" large. On the other hand, accepting less precision leads to higher confidence that the elicitation results (a subset of the parameter space) contains the parameter vector that would result from an ideally long elicitation process.

There are two possible outcomes of an indirect elicitation process: a set of constraints defining a partial information set (a subset T of the parameter space) or a (precise) vector of parameter values $t^* \in T$ (the best fit found by a mathematical program). For instance, the IRIS implementation (Dias and Mousseau 2003) of an indirect elicitation process for ELECTRE TRI (Dias et al. 2002) infers a suggested parameter vector and displays the resulting sorting of the alternatives, but it always displays all other sorting possibilities that are compatible with examples and other constraints provided by DMs.

Often, precision is not required for a model to be requisite (as defined by Phillips 1984). The analyst can follow a strategy of progressively reducing the variation for the parameters by means of new questions depending on the observation of results that are robust relatively to information provided before Dias (2007). The process stops when the DMs feel the precision in the results is requisite for their purposes. As an example we can mention an application for sorting plots of land according to their suitability for photovoltaic plants (Sánchez-Lozano et al. 2014). A subset of 20 plots was considered as potential sorting examples. At the outset, an interval of weights was considered based on the maximum and minimum values indicated by a panel of stakeholders. Then, a DM observed the range of categories in which each plot could be sorted given their characteristics and the weight intervals considered. The DM then sorted a few of these plots according to his experience-based opinion, one at a time, and observed how the range of possible categories for each plot was reduced as a result of the new constraints associated with the example. After sorting the seventh plot the number of constraints collected defined a region in the parameter space that was sufficiently precise to be able to sort each one of the remaining 13 plots into a single category. The model was considered to be requisite, concluding the elicitation process.

Setting a precise figure for each parameter value may also be an elusive goal when seeking the agreement of multiple DMs, due to differences in their preferences. It is easier for them to agree that $w_1 \in [0.25, 0.30]$ than to agree that $w_1 = 0.288$, and often conclusions are robust to vector variations within a subset of the parameter space. DMs may agree on a result although they would not be able to agree on precise values for the input parameters (Dias and Clímaco 2000). In such cases, DMs can start with little information and progressively constrain the subset of the parameter space they consider.

Avoiding eliciting precise figures is also a possibility to cope with situations in which DMs do not wish to set criteria weights, particularly in sensitive situations (e.g., impacts on the environment and on human health, or social impacts). Such DMs wish to treat criteria in a value-neutral way. An alterative to considering all criteria have the same weight is to consider that all criteria share a common interval of weights (for an example, see Domingues et al. 2015). This makes no distinction between the criteria importance, but does not entail they have the same weight. In this case, DMs would discuss the acceptable interval of weights for the criteria, discussing for instance that no criterion should weight more than all other criteria ($k_j < 0.5$), or defining a maximum acceptable ratio between any two weights (e.g., a criterion's weight cannot be more than α times greater than any other criterion's weight, Domingues et al. 2015).

14.6 Concluding Remarks

A large literature exists concerning the way by which ELECTRE methods can be implemented in practice and in particular with respect to the integration of the DM judgement in the preference model. Preference elicitation for ELECTRE methods have been largely developed and this chapter provides a synthesis of the corresponding literature.

However, there are still many challenges to be faced. An important one concerns the indirect elicitation of ELECTRE models for ranking problems: as ELECTRE methods are not invariant with respect to third a alternative, i.e, a DM can provide a statement "*a* is preferred to *b*", the inferred model will reproduce this comparison, but when applied to rank a larger set of alternatives, *b* can be better ranked than *a*.

Another challenge related to inference of ELECTRE model is related to the multiplicity of preference parameters. When eliciting indirectly these preference parameters, we usually can obtain a rather limited amount of preference statements (e.g. pairwise comparisons, or assignment examples). The contrast of the great flexibility of the preference models with the limited preference information makes it difficult to set the values of the preference parameters without some form of arbitrariness. In some applications, it might be relevant to consider some simplification of the original ELECTRE methods (avoiding some of the parameters). Another path is to collect la large amount of preference information, but this implies computational challenges related to the inference of ELECTRE models with large sets of preference statements.

A third challenge is to elicit and integrate “soft” requests, such as “I would like that criteria weights are not too different”, or “I would like that more important criteria have greater veto power than the remaining ones” in direct and especially in indirect elicitation processes.

Finally, group decision making places many different challenges. A strategy to deal with lack of agreement is working with less precise information, as suggested in the previous section. But if the DMs wish to somehow aggregate their opinions assigning different weights for the DMs’s requests (e.g. reflecting their expertise or past performance), then there is lack of research on how to take this into account in eliciting ELECTRE’s parameter values.

References

Almeida-Dias J, Figueira JR, Roy B (2010) ELECTRE TRI-C: a multiple criteria sorting method based on characteristic reference actions. Eur J Oper Res 204(3):565–580

Almeida-Dias J, Figueira JR, Roy B (2012) A multiple criteria sorting method where each category is characterized by several reference actions: the Electre Tri-nC method. Eur J Oper Res 217(3):567–579

Behzadian M, Kazemzadeh RB, Albadvi A, Aghdasi M (2010) PROMETHEE: a comprehensive literature review on methodologies and applications. Eur J Oper Res 200(1):198–215

Beuthe M, Scannella G (2001) Comparative analysis of UTA multicriteria methods. Eur J Oper Res 130(2):246–262

Bisdorff R, Meyer P, Roubens M RUBIS: a bipolar-valued outranking method for the choice problem. 4OR 6(2):143–165 (2007)

Brans JP, Vincke P (1985) A preference ranking organization method. Manag Sci 31(6):647–656

Cailloux O, Meyer P, Mousseau V (2012) Eliciting ELECTRE TRI category limits for a group of decision makers. Eur J Oper Res 223(1):133–140

Damart S, Dias LC, Mousseau V (2007) Supporting groups in sorting decisions: Methodology and use of a multi-criteria aggregation/disaggregation DSS. Decis Support Syst 43(4):1464–1475

Dias LC (2007) A note on the role of robustness analysis in decision-aiding processes. In: Roy B, Ali Aloulou M, Kalaï R (eds) Robustness in OR-DA, Annales du LAMSADE, No. 7. Université-Paris Dauphine, Paris, pp 53–70

Dias LC, Climaco JN (1999) On computing ELECTRE’s credibility indices under partial information. J Multi-Criteria Decis Anal 8(2):74–92

Dias LC, Clímaco JN (2000) ELECTRE TRI for groups with imprecise information on parameter values. Group Decis Negot 9(5):355–377

Dias LC, Mousseau V (2003) IRIS: a DSS for multiple criteria sorting problems. J Multi-Criteria Decis Anal 12(4-5):285–298

Dias L, Mousseau V (2006) Inferring ELECTRE’s veto-related parameters from outranking examples. Eur J Oper Res 170(1):172–191

Dias LC, Mousseau V, Figueira J, Clímaco JN (2002) An aggregation/disaggregation approach to obtain robust conclusions with ELECTRE TRI. Eur J Oper Res 138(2):332–348

Domingues AR, Marques P, Garcia R, Freire F, Dias LC (2015) Applying multi-criteria decision analysis to the life-cycle assessment of vehicles. J Clean Prod 107:749–759

Doumpos M, Marinakis Y, Marinaki M, Zopounidis C (2009) An evolutionary approach to construction of outranking models for multicriteria classification: the case of the ELECTRE TRI method. Eur J Oper Res 199(2):496–505

Fernández E, Figueira JR, Navarro J, Roy B (2017) ELECTRE TRI-nB: a new multiple criteria ordinal classification method. Eur J Oper Res 263(1):214–224. DOI 10.1016/j.ejor.2017.04.048.

Figueira JR, Roy B (2008) A note on the paper, "Ranking irregularities when evaluating alternatives by using some ELECTRE methods", by Wang and Triantaphyllou. Omega 37(3):731–733

Figueira J, Greco S, Roy B, Slowinski R (2013) An overview of ELECTRE methods and their recent extensions. J Multi-Criteria Decis Anal 20:61–85

Govindan K, Jepsen MB (2016) ELECTRE: a comprehensive literature review on methodologies and applications. Eur J Oper Res 250(1):1–29

Jacquet-Lagrèze E, Siskos Y (2001) Preference disaggregation: 20 years of MCDA experience. Eur J Oper Res 130(2):233–245

Keeney RL, Raiffa H (1993) Decisions with multiple objectives-preferences and value tradeoffs. Cambridge University Press, Cambridge, New York

Leroy A, Mousseau V, Pirlot M (2011) Learning the parameters of a multiple criteria sorting method. In: Brafman R, Roberts F, Tsoukiàs A (eds) Algorithmic decision theory. Lecture notes in artificial intelligence, vol 6992. Springer, Berlin, pp 219–233

Martel J-M, Matarazzo B (2005) Other outranking approaches. In: Figueira J, Greco S, Ehrgott M (eds) Multiple criteria decision analysis: state of the art surveys. Springer, Berlin, pp 198–219

Matsatsinis NF, Grigoroudis E, Siskos E (2018) Disaggregation approach to value elicitation. In: Dias LC, Morton A, Quigley J (eds) Elicitation: the science and art of structuring judgement. Springer, New York

Miettinen K, Salminen P (1999) Decision-aid for discrete multiple criteria decision making problems with imprecise data. Eur J Oper Res 119(1):50–60

Morton A (2018) Multiattribute value elicitation. In: Dias LC, Morton A, Quigley J (eds) Elicitation: the science and art of structuring judgement. Springer, New York

Mousseau V, Dias LC (2004) Valued outranking relations in ELECTRE providing manageable disaggregation procedures. Eur J Oper Res 156(2):467–482

Mousseau V, Slowinski R (1998) Inferring an ELECTRE TRI model from assignment examples. J Glob Optim 12(2):157–174

Mousseau V, Slowinski R, Zielniewicz P (2000) A user-oriented implementation of the ELECTRE TRI method integrating preference elicitation support. Comput Oper Res 27(7-8):757–777

Mousseau V, Figueira J, Naux JP (2001) Using assignment examples to infer weights for ELECTRE TRI method: Some experimental results. Eur J Oper Res 130(2):263–275

Mousseau V, Dias LC, Figueira J, Gomes C, Clímaco JN (2003) Resolving inconsistencies among constraints on the parameters of an MCDA model. Eur J Oper Res 147(1):72–93

Munda G (1995) Multicriteria evaluation in a fuzzy environment. Physica Verlag, Heidelberg

Ngo The A, Mousseau V (2002) Using assignment examples to infer category limits for the ELECTRE TRI method. J Multi-Criteria Decis Anal 11(1):29–43

Phillips LD (1984) A theory of requisite decision models. Acta Psychol 56:29–48

Rogers M, Bruen M (1998) Choosing realistic values of indifference, preference and veto thresholds for use with environmental criteria within ELECTRE. Eur J Oper Res 107(3):542–551

Rolland A (2013) Reference-based preferences aggregation procedures in multi-criteria decision making. Eur J Oper Res 225(3):479–486

Roy B (1968) Classement et choix en présence de points de vue multiples: Le méthode ELECTRE. Revue Francaise d'Informatique et de Recherche Opérationnelle. 8:57–75

Roy B (1971) La méthode ELECTRE II. Technical report. METRA, Direction Scientifique, Note de Travail n. 142

Roy B (1978) ELECTRE III: Un algorithme de classements fondé sur une représentation floue des préférences en présence de critères multiples. Cahiers du CERO 20(1):3–24

Roy B (1991) The outranking approach and the foundations of ELECTRE methods. Theory Decis 31(1):49–73

Roy B (1996) Multicriteria methodology for decision aiding. Kluwer Academic, Dordrecht

Roy B (1998) A missing link in OR-DA, Robustness analysis. Found Control Eng 23(3):141–160

Roy B, Bertier P (1973) La méthode ELECTRE II - une application au média-planning. In: Ross M (ed) OR'72. North-Holland Publishing Company, Amsterdam, pp 291–302

Roy B, Bouyssou D (1993) Aide Multicritère à la Décision: Méthodes et Cas. Economica, Paris

Roy B, Figueira J (2002) Determining the weights of criteria in the ELECTRE type methods with a revised Simos' procedure. Eur J Oper Res 139:317–326

Roy B, Hugonnard JC (1982) Ranking of suburban line extension projects of the Paris metro system by a multicriteria method. Transp Res 16A:301–322

Roy B, Skalka JM (1984) ELECTRE IS : Aspects méthodologiques et guide d'utilisation. Technical report, Document du LAMSADE N^o 30, Université Paris-Dauphine, Paris

Roy B, Vanderpooten D (1996) The European school of MCDA: emergence, basic features and current works. J Multi-Criteria Decis Anal 5(1):22–37

Roy B, Figueira JR, Almeida-Dias J (2014) Discriminating thresholds as a tool to cope with imperfect knowledge in multiple criteria decision aiding: theoretical results and practical issues. Omega 43:9–20

Sánchez-Lozano JM, Antunes CH, García-Cascales MS, Dias LC (2014) GIS-based photovoltaic solar farms site selection using ELECTRE-TRI: Evaluating the case for Torre Pacheco, Murcia, Southeast of Spain. Renew Energy 66:478–494

Simos J (1990) Evaluer limpact sur lenvironnement: Une approche originale par l'analyse multicritère et la négociation. Presses Polytechniques et Universitaires Romandes, Lausanne

Sobrie O (2016) Learning preferences with multiple-criteria models. PhD thesis, UMONS and Université Paris Saclay

Sobrie O, Mousseau V, Pirlot M (2013) Learning a majority rule model from large sets of assignment examples. In: Perny P, Pirlot M, Tsoukiás A (eds) Algorithmic decision theory. Lecture notes in artificial intelligence, vol 8176. Springer, Brussels, pp 336–350

Vallée D, Zielniewicz P (1994) ELECTRE III-IV, version 3.x, Aspects Méthodologiques (tome 1), Guide d'utilisation (tome 2). Document du LAMSADE no. 85 et 85 bis, Université de Paris Dauphine, France

Yu W (1992a) Aide multicritère à la décision dans le cadre de la problématique du tri: méthodes et applications. PhD thesis, LAMSADE, Université Paris Dauphine, Paris

Yu W (1992b) ELECTRE TRI : Aspects méthodologiques et manuel d'utilisation. Document du lamsade no 74, Université Paris-Dauphine