Preference Programming

September 16, 2004

Ahti Salo and Raimo P. Hämäläinen
Systems Analysis Laboratory
Helsinki University of Technology
P.O. Box 1100, 02015 HUT, Finland

Abstract

Methods for dealing with incomplete preference information in hierarchical weighting models have continued to attract attention in the literature on multi-criteria decision analysis (MCDA). In this paper, we give a structured overview of several such methods which (i) accommodate incomplete preference information, (ii) offer dominance concepts and decision rules for the generation of decision recommendations and (iii) support the iterative exploration of the decision maker's preferences. By doing so, we synthesize much of the relevant literature and provide an integrative perspective on these methods which are here subsumed under the term 'preference programming'. We then demonstrate that these methods may outperform conventional decision analyses when the costs of preference elicitation are high and, moreover, provide guidelines for responsible uses of preference programming. We conclude by outlining topics for future research.

Keywords: Multi-criteria decision analysis, hierarchical weighting models, incomplete preference information, group decision and negotiation, decision support systems.

1. INTRODUCTION

Hierarchical weighting methods - such as value trees (Keeney and Raiffa, 1976) and the Analytic Hierarchy Process (Saaty, 1980)) - are widely employed in the analysis of decision problems characterized by incommensurate objectives, competing alternatives and conflicting stakeholder interests (see, e.g., Corner and Kirkwood, 1991; Keefer et. al, 2004; Hämäläinen, 2004). In these methods, the decision makers (DM) are encouraged to structure their objectives as a hierarchy of attributes, whereby the very process of developing such a representation can be helpful but challenging, too (Belton and Stewart, 2001; Keeney, 1992). In effect, the hierarchical representation provides a framework for synthesizing information about (1) how the alternatives perform on the attributes (=scores) and (2) how important the attributes are (=weights) so that an overall performance measure can be associated with each alternative.

Decision analyses are usually based on the assumption that complete information about the model parameters (scores, attribute weights) can be elicited. Yet, this assumption can be questioned on several grounds. For instance, it may be impossible to obtain complete information about the alternatives; and even if such information can be obtained, it may come at a high cost, suggesting that it is of interest to examine what tentative conclusions might be supported by the available but possibly incomplete information (Weber, 1987; Kim and Han, 1999). Also, if the decision makers (DM) are forced to provide more complete a preference specification than what they feel confident with, they may distrust the results when these are based on undependable inputs. Moreover, from the viewpoint of sensitivity analysis, it is advisable to check if the recommendations would change when the model parameters are allowed to vary within plausible limits (Rios Insua and French, 1991).

The above arguments have motivated the development of methods which accommodate incomplete preference information in hierarchical weighting models (see, e.g., White et al., 1984; Eum et al., 2001; Kim and Han, 1999, 2000; Park and Kim, 1997; Mármol et al., 1998; Puerto et al., 2000; Salo and Hämäläinen, 1992, 1995, 2001; Weber, 1987). Despite their differences, these methods exhibit similarities, most notably in (1) the use of set inclusion as a means of modelling incomplete preference information, (2) the application of dominance structures and decision rules in the generation of decision recommendations and (3) the presentation of supplementary information for guiding the DM in preference elicitation. In view of these similarities, we employ 'preference programming' – a term that was introduced by Arbel (1989) to denote incomplete preference information in the AHP – as a general term for methods which fulfil at least the two first of the above features. This terminological convention seems pertinent also because the DM is engaged in an interactive exploration of preferences while mathematical programming techniques are used in the generation of results. A further advantage of the term 'preference programming' is that it allows us to address issues that apply to several methods.

The literature on preference programming has been largely dominated by incremental theoretical contributions, although promising real-life applications have also been reported (see, e.g., Hämäläinen and Pöyhönen, 1996; Gustafsson et al., 2001). Yet, little work has been done to synthesize 'lessons learned' from applied work towards a general framework, or to examine when and how preference programming should be applied. In what follows, we address these questions by developing decision models in which (1) the costs of preference elicitation are explicitly incorporated and (2) the requirement of treating all alternatives on equal terms is operationalized. These models suggest that preference programming may outperform conventional approaches based on a complete preference specification which is characterized by point estimates for all parameters.

The remainder of this paper is structured as follows. Section 2 presents an overall framework for preference programming and discusses selected applications thereof. Section 3 examines decision models where the costs of preference elicitation and non-optimal choices are explicitly accounted for. Building on these models, Section 4 identifies fruitful contexts for preference programming and derives guidelines for preference elicitation. Section 5 concludes with topics for further theoretical and applied research.

2. ELEMENTS OF PREFERENCE PROGRAMMING

2.1. Additive Preference Representation

We assume there are n attributes at the lowest level of the hierarchical problem representation. The importance of the i-th attribute is measured by its *weight* $w_i \in [0,1]$; these weights are normalized so that $\sum_{i=1}^{n} w_i = 1$. There are m alternatives $x^1, \ldots, x^m$, and the achievement level of the j-th alternative on the i-th attribute is denoted by x_i^j. The single-attribute value associated with this achievement level is the *score* $v_i(x_i^j) = v_i^j \in [0,1]$. If the model parameters (i.e., weights $w = (w_1, \ldots, w_n) \in W = \{w | w_i \geq 0, \sum_{i=1}^{n} w_i = 1\}$ and scores $v^j = (v_1^j, \ldots, v_n^j), j = 1, \ldots, m)$ are known, the (overall) value of alternative x^j is given by the sum $V(x^j) = \sum_{i=1}^{n} w_i v_i^j$.

In what follows, 'incomplete preference information' refers to settings where the constraints implied by the DM's preference statements are satisfied by several parameter values (i.e., attribute weights and score vectors). More formally, this means that the sets of feasible weight and score vectors S_w, S_{v^j} (where $w \in S_w \subset W$ and $v^j = (v_1^j, \ldots, v_n^j) \in S_j \subset [0,1]^n, j = 1, \ldots, m)$ are non-empty, and that at least one of them contains more than a single element.

Even other terms – such as 'imprecise information' (e.g., Salo and Hämäläinen, 1991; Eum et al., 2001) or 'partial information' (e.g., Carrizosa et al., 1995) – have been employed in much the same sense. Nevertheless, 'incompleteness' seems more adequate because it stresses that while the available preference information (as captured by the sets S_w and S_j) may not lead to a full ranking of alternatives, such information can be completed through the elicitation of additional preference statements. 'Incompleteness' can also be defended by noting that the constraints on feasible weights and scores are not 'imprecise': for instance, the lower and upper bounds of interval-valued ratio statements in weight elicitation are crisp numbers with no associated uncertainties.

Although we do not review approaches based on probabilistic modelling (e.g., Saaty and Vargas, 1987), fuzzy sets (e.g., Salo, 1996a) or outranking relationships (e.g., Stewart and Losa, 2003), there are computational and other parallels between these approaches and preference programming: the consideration of incompletely specified probabilities in Bayesian updating, for instance, leads to related problem formulations (Salo, 1996b). Many of our remarks are consequently relevant to other approaches, too, particularly as regards the use of decision rules in the development of decision recommendations.

2.1 Preference Elicitation

In general, preference programming methods extend the range of approaches that can be used in the elicitation of DM's preferences. Many of them incorporate ordinal preference information and relax the requirement for exact (crisp) numerical estimates by allowing the DM to provide interval-valued estimates. For example, extensions of popular ratio-based techniques – such as the AHP (Saaty, 1980) and SMART (Edwards, 1977) – have inspired methods where the DM may provide interval-valued statements about the relative importance of attributes (Arbel, 1989; Salo, 1995; Salo and Hämäläinen, 1992, 2001; Mustajoki et al., 2001). In broad terms, Kim and Ahn (1999) distinguish between the following approaches to the elicitation of attribute weights: (1) weak ranking (i.e., $w_i \geq w_j$), (2) strict ranking ($w_i - w_j \geq \alpha_{ij}$) (3) ranking with multiples ($w_i \geq \alpha_{ij} w_j$), (4) interval form ($\alpha_i \leq w_i \leq \alpha_i + \epsilon_i$), (5) ranking of differences ($w_i - w_j \geq w_k - w_l$ for $j \neq k \neq l$,) where $\alpha_., \epsilon_. \geq 0 \, \forall \, i$. All of these statements correspond to linear constraints on attribute weights.

In a recent extension, Salo and Punkka (2003) present an approach entitled Rank Inclusion in Criteria Hierarchies (RICH) which allows the DM to provide incomplete ordinal information about the relative importance of attributes (e.g., 'cost is among the three most important attributes' or 'the most important attribute is either cost or quality'). Such statements lead to possibly non-convex sets of feasible attribute weights, but decision recommendations can still be derived through the application of dominance concepts and decision rules.

2.2 Determination of Dominance Structures

Once an incomplete preference specification has been elicited (as characterized by feasible weights S_w and scores S_j), it is of interest to know which alternatives dominate others in view of the resulting specification. Such dominance results can be based on the concept (pairwise) dominance: i.e., alternative x^k dominates x^l if the (overall) value of x^k is higher than that of x^k for all feasible model parameters or, equivalently, if

$$\sum_{i=1}^{n} w_i v_i^k \geq \sum_{i=1}^{n} w_i v_i^l, \tag{1}$$

holds for all combinations of feasible weights $w \in S_w$ and scores $v^j \in S_j$. If (1) holds, the value of x^k remains at least as high as that of x^l, even if additional preference statements are added until the sets of feasible weights and scores become singletons.

2.3 Application of Decision Rules

If there are several non-dominated alternatives in (1), it is not possible to conclude which alternative is 'best': this is because for any alternative x^k, there exists a combination of feasible weights and scores such that the overall value of some other alternative x^l is higher than that of x^k. As a result, other principles – called *decision rules* – are needed to derive a decision recommendation if possibilities for continued preference elicitation are limited. Several approaches have been proposed towards this end:

1. **Choice of representative parameters:** Based on 'representative' parameters from feasible regions, the alternative for which the corresponding overall value is highest can be recommended. For instance, one of the decision rules in PRIME uses *central weights* which are near the center of the feasible weight set (Salo and Hämäläinen, 2001). Likewise, alternative approaches to the computation of rank based weights (e.g., rank sum, rank reciprocal, rank exponent, rank order centroid; see, Barron and Barrett, 1996; Edwards and Barron, 1994; Stillwell et al., 1981) are essentially ways of converting ordinal preference information into representative weight vectors.

2. **Properties of alternatives' value ranges:** Recommendations can be based on an analysis of the ranges of values that alternatives may take, without making assumptions about the probabilities of feasible parameters. Examples include the *maximax rule* (i.e., choose the alternative with the highest possible value), *maximin* (i.e., choose the alternative with the smallest possible value) and *central*

values (i.e., choose the alternative with the highest mid-point of its value interval) (see Salo and Hämäläinen, 2001). The advantage of these rules is that they can be readily computed and communicated.

3. **Pairwise value differences between alternatives:** Decision rules can be based on measures on how well alternatives perform relative to others. One such measure is the *maximum loss of value* which indicates by how much the value of some other alternative can at most exceed that of x^i, if the DM were to choose x^i (Salo and Hämäläinen, 2001). This *minimax regret* rule is appealing in that the implications of incomplete information are linked to *ex post* decision quality, allowing the DM to take an informed decision on whether or not to continue with preference elicitation.

4. **Maximization of expected value:** If sufficient assumptions about the probabilities of feasible parameters can be made, the alternative with the highest expected overall value can be recommended (see, e.g., Weber, 1985). Despite its conceptual appeal, this decision rule may be problematic, because the elicitation of fully specified probability distributions would call for a major effort. Computational difficulties, too, may emerge.

5. **Likelihood maximization of potentially optimal alternatives:** If probability distributions on the feasible regions can be elicited, the alternative which has the highest probability of receiving the largest overall value can be offered as the decision recommendation. In effect, this approach is the SMAA method (Lahdelma at al., 1998) which recommends only potentially optimal alternatives, thus precluding non-dominated alternatives which perform reasonably well across the entire feasible set, but which are not optimal for any combination of feasible parameters. In this sense, the SMAA method represents a rather optimistic stance in the face of incomplete information.

Even though several decision rules have been proposed, the literature offers little guidance as to which decision rules should be applied in particular decision contexts. Recent simulation studies suggest that decision rules based on the use of central values tend to outperform others in terms of minimizing the expected loss of value (Salo and Hämäläinen, 2001). But because this conclusion depends on a number of context-specific assumptions (e.g., absence of correlations among alternatives), further research is needed to provide well-grounded advice for choices among alternative decision rules.

The fact that different decision rules may favor different alternatives is a source of potential concern. In group decision making, for instance, the problem of choosing among alternatives may become interpreted as the problem of making choices among competing decision rules, whereby opportunistic group members may insist on decision rules that support their favorite alternatives. It therefore seems that principles for the application of decision rules should be agreed upon in advance, to avoid situations where these principles become a source of irremediable conflict.

2.4 Management of Inconsistencies

The derivation of decision recommendations from an incompletely specified preference model presumes that the DM's preference statements are consistent so that they define a non-empty set of feasible weights and scores. Without adequate decision support, however, the DM may provide inconsistent statements so that the feasible set becomes empty. Two main approaches have been proposed to avoid this possibility:

- **Consistency restoration:** Taking the set of conflicting constraints as a point of departure, the DM can be requested to modify or withdraw earlier statements until the remaining, possibly revised constraints do not conflict with each other any more (see, e.g., White et al. 1984, Kim and Han 1999).
- **Consistency preservation:** Before the elicitation of each new preference statement, full information about the implications of earlier preference statements can be computed and presented to the DM, to ensure that the new statement is not in conflict with the earlier ones (see, e.g., by Salo and Hämäläinen 1992, 1995, 2001).

Consistency restoration may be problematic if the DM is not able or willing to revisit earlier statements. The withdrawal of earlier statements may also undermine the credibility of the analysis, because it suggests that there are 'errors' in some inputs without guaranteeing that the other inputs are less 'erroneous'. Also, although automated procedures can be used to identify the least number of constraints that should be removed to re-establish consistency, such procedures are computational interventions with little interaction on the part of the DM (see, e.g., White et al., 1984): at worst, this approach may lead to the removal of statements the DM feels most confident about. But from a computational point of view, the advantage of consistency restoration is that it can be applied in a batch mode so that possible problems with inconsistencies – if they do arise – can be addressed after the preference elicitation phase.

Consistency preservation requires that the implications of earlier preference statements are presented to the DM whenever new statements are elicited. This approach thus puts high demands on the ability to analyze the preference model in view of all those statements that the DM may wish to state next. Even so, if the DM does wish to enter a statement which violates these resulting *consistency bounds* (Salo and Hämäläinen, 1992), it may be difficult to determine what earlier statements should be removed or adjusted to produce a revised specification that is consistent with the DM's next statement. This question has also broader implications for preference elicitation: for example, should the DM be encouraged to provide relatively 'narrow' statements (which tend to support more conclusive dominance results, but are more prone to inconsistencies) or 'broad' statements (which entail a lower risk of inconsistencies, but are likely to produce less conclusive dominance results)?

Table 1 gives an overview of selected methods of preference programming, with particular attention to preference elicitation as well as conceptual, mathematical and computational advances. Overall, the general trend is from the mere incorporation of incomplete information towards the generation of supplementary information (e.g., consistency bounds, measures of incompleteness, decision rules) which helps the DM to decide whether or not the elicitation phase should be continued and, if so, how such a continuation should be organized. Another important development is the increasing availability of corresponding decision support tools which have enabled several case studies.

Insert Table 1 about here

2.5. Case Studies and Decision Support Tools

While most papers on preference programming are concerned with theoretical advances, the literature contains insightful reports on applied work, too. Anandaligam (1989), for instance, describes an application of incomplete preference information in mitigating the harmful consequences of acid raid. Hämäläinen et al. (1992) consider the use of preference programming in assisting groups of decision makers in the comparison of energy policy options. Hämäläinen and Pöyhönen (1996) report experiences from the development of policies for traffic management, highlighting the impacts of preference programming on the decision outcome and the decision support process. The use of multi-attribute decision analysis in nuclear emergency management is discussed by

Hämäläinen et al. (2000) who focus on the consideration of uncertainties. Cristiano et al. (2001) accommodate incomplete preference information in quality function deployment and demonstrate how such information help in the design of a surgical product. Gustafsson et al. (2001) apply the PRIME method to the valuation of a high technology company and illustrate how preference programming can be used in scenario-based forecasting problems.

Several software tools have been developed to support the application of preference programming. WINPRE (Workbench for INteractive PREference Programming; Helenius and Hämäläinen, 1998) provides computerized support for both PAIRS (Salo and Hämäläinen, 1992) and the use of interval valued-statements in the AHP (Salo and Hämäläinen, 1995). The RINGS system (Kim and Choi, 2001) allows the DMs to analyze range-based information in multi-attribute group decision making. PRIME Decisions (Gustafsson et al., 2001) features *elicitation tours* which assist the DM in the specification of interval-valued ratio statements. RICH Decisions (Salo et al., 2003) is an internet-based decision support tool which admits incomplete ordinal preference information and offers several decision rules (Salo and Punkka, 2003).

Experiences from case studies such as these lend support to the following observations:

- Even though the DMs are able and willing to provide incomplete preference information, the sheer number of possibilities in preference elicitation can be a major challenge. It is not necessarily apparent what kinds of elicitation questions should be posed to the DM, in order to obtain information that effectively contributes to the identification of preferred alternatives. This concern with effectiveness is warranted because the DM needs to specify two crisp numbers in order to define a single interval, for example.
- If the preference model remains relatively incomplete (so that the feasible parameter sets are large), dominance structures are unlikely to help in the determination of most preferred alternatives. In such situations, it is necessary to apply decision rules or to seek other ways of terminating the analysis (e.g., informal negotiations among the group members; see Hämäläinen and Pöyhönen, 1996). Thus, principles for choosing among decision rules and for carrying out MCDA-assisted negotiations are essential.

In effect, alternative approaches to preference elicitation, choices among decision rules and principles for terminating the decision support process are all intertwined. For example, because the recommendations of decision rules depend on the DM's earlier preference statements, these recommendations depend on just how early or late in the preference elicitation process these rules are applied. Close attention must consequently be given to what kinds of elicitation questions should be posed to the DM, and when and how intermediate results should be presented.

3. MODELS OF DECISION QUALITY

3.1. Possible Loss of Value in Preference Programming

The costs of preference elicitation have not been addressed in most applications of hierarchical weighting models, partly because the DM is usually requested to provide a complete specification and partly because in some contexts (e.g., siting of large facilities) these costs are negligible in comparison with the costs of the alternatives. Even in the preference programming literature – where elicitation costs are cited among the motivating reasons – it has not been formally explored what these costs imply for choices among different elicitation approaches and decision rules.

We now turn to the above question by developing illustrative models, with the aim of developing implications for the structuring of preference elicitation processes and the application of decision rules. The preference statements are expressed as ratio comparisons, because several preference programming methods employ such comparisons and because the resulting models are more tractable than those based on arbitrary preference statements. By construction, these models are simple representations of hypothetical decision problems; however, from the viewpoint of our purposes this is not a major restriction because similar implications would apply to more complex models as well. Nor is our focus on the maximax decision rule restrictive as related models in support of these implications can be established for other decision rules, too.

To begin with, let there be two attributes and an infinite number of alternatives such that there is an alternative for all pairs of scores (v_1, v_2) such that $v_1, v_2 \geq 0$ and $v_1^2 + v_2^2 = 1$. Now, if the DM states that the first attribute is more important than the second, the set of feasible attribute weights is $S_0 = \{(w_1, w_2)|w_1 \geq w_2,\ w_1 + w_2 = 1\})$. If the maximax decision rule is adopted – meaning that the best alternative is the one for which the maximum overall value over the set S_0 is highest – the corresponding decision recommendation is the alternative $v^1 = (v_1^1, v_2^1) = (1, 0)$ which gives the highest value to the sum $\max_{w \in S_0}[w_1 v_1 + w_2 v_2]$.

We assume that the true (unknown) weight vector $w^* = (w_1^*, w_2^*)$ is a random variable, determined by the ratio $r = w_1^*/w_2^*$ which follows the probability density function $f_r(\cdot)$

over the interval $[1, \infty[$. The optimal alternative (v_1^*, v_2^*) (which corresponds to the true weight vector w^*) is characterized by the condition $v_1^*/v_2^* = w_1^*/w_2^* w = r$, i.e., $(v_1^*, v_2^*) = \frac{1}{\sqrt{1+r^2}}(r, 1)$. The loss of value which results from choosing v^1 (as opposed to v^*) is therefore

$$\begin{aligned} \mathrm{LV}(v^1) &= (w_1^*, w_2^*)(v_1^*, v_2^*)^T + (w_1^*, w_2^*)(1, 0)^T \\ &= \frac{r}{1+r}\frac{r}{\sqrt{1+r^2}} + \frac{1}{1+r}\frac{1}{\sqrt{1+r^2}} - \frac{r}{1+r} \\ &= \frac{1}{1+r}(\sqrt{1+r^2} - r). \end{aligned}$$

Next, assume that incomplete information about the true ratio r between the weights w_1^* and w_2^* is obtained so that the DM specifies an upper bound $\overline{r} = r\epsilon$ on the ratio r (where ϵ follows the probability density function $f_\epsilon(\cdot)$ over the interval $[1, \infty[$). When this information is added to the other constraints on S_0, the application of the maximax decision rule leads to the alternative $\overline{v} = (\overline{v}_1, \overline{v}_2) = \frac{1}{\sqrt{1+(r\epsilon)^2}}(r\epsilon, 1)$. For this alternative, the loss of value (relative to the optimal alternative v^*) is

$$\begin{aligned} \mathrm{LV}(r, \epsilon) &= (w_1^*, w_2^*)(v_1^*, v_2^*)^T - (w_1^*, w_2^*)(\overline{v}_1, \overline{v}_2)^T \\ &= \frac{r}{r+1}\frac{r}{\sqrt{1+r^2}} + \frac{1}{r+1}\frac{1}{\sqrt{1+r^2}} - \frac{r}{r+1}\frac{x\epsilon}{\sqrt{1+(r\epsilon)^2}} - \frac{1}{r+1}\frac{1}{\sqrt{1+(r\epsilon)^2}} \\ &= \frac{1}{r+1}[\sqrt{1+r^2} - \frac{1+r^2\epsilon}{\sqrt{1+(r\epsilon)^2}}]. \end{aligned} \tag{2}$$

Because r and ϵ are random variables, the expected loss of value (based on the acquisition of upper bound $\overline{r}$ and the application of the maximax rule) is given by

$$\mathrm{ELV}(f_r, f_\epsilon) = \int_1^\infty [\int_1^\infty \mathrm{LV}(\rho, \chi) f_\epsilon(\chi) d\chi] f_r(\rho) d\rho. \tag{3}$$

The derivative of (2) with respect to ϵ is $[r^2/(\sqrt{1+(r\epsilon)^2})^3](\epsilon - 1) > 0$, and thus the loss of value $\mathrm{LV}(r, \epsilon)$ is minimized when ϵ attains its lower bound: in fact, it is zero if and only if $\epsilon = 1$. If ϵ follows a log-normal distribution (i.e., $\epsilon = e^\varepsilon$ where $\varepsilon \geq 0$ and $\varepsilon \propto N(0, \sigma_\varepsilon^2)$), it follows that the expected loss of value (3) approaches zero as the term σ_ϵ tends to zero.

In the above model, the upper bound $\overline{r}$ is likely to be near the true ratio r only if the variance σ_ε^2 is small. The elicitation of such a 'good' upper bound is likely to entail higher costs than that of an upper bound with a larger variance. When the costs of preference elicitation costs are accounted for by a decreasing cost function $C(\sigma_\epsilon)$, the DM is essentially faced with the problem of minimizing

$$\min_{\sigma_\epsilon}[ELV(f_r, f_\epsilon) + C(\sigma_\epsilon)] \tag{4}$$

where the first term stands for the expected loss of value and the second term denotes the costs of reaching a given level of accuracy.

If these costs are sufficiently high for small values of σ_ϵ, the optimum solution to (4) is obtained at some $\sigma_\epsilon > 0$, indicating that the use of incomplete information leads to lower total costs than attempts to elicit complete information. Another important result suggested by (4) is that – at any stage of analysis – the costs of further elicitation efforts should be balanced by equal or greater improvements in decision quality, as measured by the reduction in the expected loss of value value.

The above probability model is essentially a conceptual framework, because the specification of requisite probability distributions would entail an inordinate elicitation effort. It is simplistic in that the DM is assumed to be risk-neutral with respect possible loss of value. Nor is the elicitation of probability distributions in keeping with the 'spirit' of preference programming methods which use set inclusion in the modelling of incomplete information.

Yet, if probabilistic modelling were to be pursued despite the above remarks, the DM should exploit her beliefs about the distribution of the error term ϵ when revising her assumptions about the distribution of the model parameters (such as the weight vector w^*). For instance, if the ratio $r = w_1^*/w_2^*$ and the error term ϵ follow the distribution $f_r(\tau) = f_\epsilon(\tau) = 1/\tau^2, (\tau \geq 1)$, the DM's revised probability for the statement $w_1^* < \varrho w_2^*$, conditioned on the upper bound $\overline{r}$, becomes

$$\begin{aligned} F_r(r < \varrho | \epsilon r = \overline{r}) &= K \int_1^{\varrho} [\int_1^{\overline{r}/x} (\frac{1}{y})^2 dy](\frac{1}{x})^2 dx \\ &= K[1 - \frac{1}{\varrho} - \frac{1}{\overline{r}} \ln \varrho], \end{aligned} \tag{5}$$

where $K = 1/[1 - (1/\overline{r})(1 + \ln \varrho)]$ is the normalization constant. Here, the conditional distribution (5) has a different functional form than $F_r(\varrho) = [1 - (1/\varrho)]$. These and other computational complexities suggest that the DM cannot be expected to master Bayesian inferences without extensive support for probability judgments.

The above Baysian model has implications for the performance evaluation of preference programming methods. Often, a sample of representative weights is generated from a uniform distribution over the weight set $S_0 = \{(w_1, \ldots, w_n) | \sum_{i=1}^n w_i = 1, w_i \geq 0\}$. Using these weights, an incomplete preference specification is derived by applying a constant term $\Delta > 1$ to the 'true' ratios between the components of weight vectors: for instance, if the true ratio between the two first components of the weight vector is $r_{12} = w_1/w_2$, the lower and upper bounds in an interval statement might be defined as $w_1/w_2 \in [(1/\Delta)r_{12}, \Delta r_{12}]$ (Salo and Hämäläinen, 2001).

Even though it does support the comparison of different decision rules, the above evaluation approach can be criticized on the grounds that the 'true' ratio (and hence the underlying weight vector) can, in principle, be inferred from the interval because the *same* error term Δ is used in the computation of *both* lower and upper bounds. In fact, this symmetric distribution of lower and upper bounds may explain why decision rules based on estimates nearer the middle of these intervals (e.g., central weights) tend to outperform decision rules which choose alternatives that are optimal at some extreme point of the feasible weight set (e.g., maximax rule).

3.2 Equitable Treatment of Alternatives

Decision rules can be used at any stage of the analysis to obtain a decision recommendation based on the statements that the DM has provided up to that point of analysis. Yet, as a matter of principle, it is plausible to require that all alternatives should be treated in the same way, in the sense that no alternative is disadvantaged due to the particular sequencing of preference elicitation steps or the timing of decision rules. We next develop a decision model which helps formalize what this requirement for equitable treatment of alternatives implies for preference elicitation.

Let $v_1^1 = v_1(x_1^1)$ and $v_1^2 = v_1(x_1^2)$ be the scores of alternatives x^1, x^2 on attribute a_1. Without losing generality, we assume that these scores belong to the range $[0, 1[$. These scores are characterized by the ratios $r_1 = 1/(1 - v_1^1)$ and $r_2 = 1/(1 - v_1^2)$ which, by construction, are greater than one. Conclusions about which alternative performs better on attribute a_1 can be inferred from these ratios: for instance, x^1 is better than x^2 with regard to a_1 if and only if $r_1 > r_2$.

If there are no *a priori* reasons for concluding that either alternative is better than the other, the ratios r_1, r_2 can be regarded as random variables with the same probability density function f_r over the interval $[1, \infty[$. In this case, incomplete score information can be modelled through upper bounds r_1', r_2' on the ratios r_1, r_2, defined by $r_1' = r_1\epsilon_1$ and $r_2' = r_2\epsilon_2$ where ϵ_1, ϵ_2 are random variables over the interval $[1, \infty[$ with probability density functions f_1, f_2. Now, if a decision recommendation is derived on the basis of this information and the application of the maximax decision rule (with regard to the first attribute only), the recommended alternative would be x^1 if $r_1' > r_2'$ and x^2 if $r_2' > r_1'$.

One way of interpreting the requirement for an 'equitable treatment of alternatives' is that, in the absence of any distinguishing information, the probability of erroneously choosing x^1 when x^2 is the better alternative should be the same as the probability that of choosing x^2 when x^1 is the better alternative. These two probabilities are given by the integrals

$$P(v_1'^1 > v_1'^2 | v_1^2 > v_1^1) = \int_1^\infty \{ \int_{r_1}^\infty [\int_1^\infty (\int_{\frac{r_2 \epsilon_2}{r_1}}^\infty f_1(\epsilon_1) d\epsilon_1) f_2(\epsilon_2) d\epsilon_2] f_r(r_2) dr_2 \} f_r(r_1) dr_1 \quad (6)$$

$$P(v_1'^2 > v_1'^1 | v_1^1 > v_1^2) = \int_1^\infty \{ \int_{r_2}^\infty [\int_1^\infty (\int_{\frac{r_1 \epsilon_1}{r_2}}^\infty f_2(\epsilon_2) d\epsilon_2) f_1(\epsilon_1) d\epsilon_1] f_r(r_1) dr_1 \} f_r(r_2) dr_2 \quad (7)$$

The ratios r_1, r_2 appear symmetrically in the above integrals. Thus, by interchanging the roles r_2 and r_1 in one of the integrals, it follows that the probabilities (6) and (7) are equal if the probability density functions $f_1(\cdot), f_2(\cdot)$ and their cumulative probability functions $F_1(\cdot), F_2(\cdot)$) satisfy the following (sufficient) condition

$$\int_1^\infty [f_1(\epsilon) F_2(\alpha\epsilon) - F_1(\alpha\epsilon) f_2(\epsilon)] d\epsilon = 0, \quad (8)$$

where $\alpha = \max\{r_1/r_2, r_2/r_1\}$ is a constant.

The condition (8) is satisfied if the error terms follow the same distribution (i.e., $f_1(\cdot) = f_2(\cdot)$). If not, the condition (8) hardly ever holds: for instance, if ϵ_1 and ϵ_2 follow the log-normal distribution and ϵ_1 has the higher variance, then alternative x^1 has the higher probability of being (erroneously) regarded as the better alternative (on the basis of the maximax decision rule) when the alternatives perform equally well. There are, however, some highly exceptional cases with non-identical probability distributions where (8) may hold: this is the case for distributions $f_1(\cdot)$ and $f_2(\cdot)$ defined by $f_1(2) = 0.4$, $f_1(4) = 0.2$, $f_1(8) = 0.4$ and $f_2(2) = 0$, $f_2(4) = 1$, $f_2(8) = 0$, for example.

If both ratios are equally 'difficult' to estimate (meaning that ϵ_1 and ϵ_2 follow the same distribution), the above result implies that the DM should provide the upper bounds r_1', r_2' using the same confidence level (e.g., fractile of the cumulative probability function), for else one alternative will be disadvantaged relative to the other. From the viewpoint of preference elicitation, the key implication of this is that the same level of *thoroughness* should be pursued when addressing with all alternatives (e.g., by asking the DM to establish intervals which contain the scores at some fixed confidence level). Furthermore, decision rules *should not* be applied before score information on all alternatives has been elicited. Otherwise, there is a possibility that those alternatives that have not been considered at all (or at the same level of thoroughness) are unduly disadvantaged.

The above analysis can be readily adapted to the elicitation of attribute weights, too. Assume there are three attributes and two alternatives x^1, x^2 such that the scores of the first alternative are $v^1 = (v_1(x_1^1), v_2(x_2^1), v_3(x_3^1)) = (1, 0, 0)$ while those of the second one are $v^2 = (0, 1, 0)$. Also, assume that the least important of the three attributes is the third one, and that this attribute is used as the reference attribute in two ratio-based weight elicitation questions. In this case, the overall values of the two alternatives can be determined on the basis of estimates about the weight ratios $r_1 = w_1/w_3$ and $r_2 = w_2/w_3$. But if it is possible to obtain only upper bounds for these ratios such that $r_1' = r_1\epsilon_1$ and $r_2' = r_2\epsilon_2$, the maximum value for the first alternative, computed on the basis of this estimate, is $w_1' = r_1\epsilon_1 w_3'$ while that of the second is $w_2' = r_2\epsilon_2 w_3'$. Thus, the probability of choosing the second (but inferior) alternative x^2 is $P(r_1 x_1 > r_2\epsilon_2 | r_1 > r_2)$, which is essentially the same expression as in (6) and (7). It therefore follows that just as in score elicitation, the same level of thoroughness should be pursued in assessing attribute weights.

4. IMPLICATIONS FOR THE PRACTICE OF DECISION ANALYSIS

4.1 Preference Elicitation

The fundamental philosophy in preference programming is that (1) an effort is made to elicit as much preference information as reasonably possible and (2) if the resulting preference specification does not lead to the identification of the most important alternative, a decision recommendation is produced with the help of a suitable decision rule. In this context, the decision models in Section 3 suggest several guidelines for the practice of decision analysis:

- An attempt should be made to obtain equally 'complete' score information for all alternatives, in the sense that the DM is equally confident in that the preference specification of each alternative contains the 'true' scores. This can be encouraged by attaching a progression of confidence intervals to the preference statements, for example. One may also apply fuzzy mathematics in the aggregation of such confidence intervals (see, e.g., Salo, 1996a).
- The same level of 'thoroughness' should also be pursued when assessing the relative importance of attributes. If this requirement is violated, there is a possibility that some alternative becomes disadvantaged, only because it has its highest score (relative to the other alternatives) on an attribute about which little weight information is available (meaning that this attribute would count less in the application of the maximin rule, for instance).

- In the elicitation of attribute weights, it is advisable not to mix different types of preference elicitation questions, because this may result in a feasible region that is less symmetric than what is obtained by relying on a single question type (e.g., interval-valued ratio statements). Limiting the number of different types of elicitation questions may also put a smaller cognitive load on the DM.

If an *a priori* decision is taken to continue the analysis until the most preferred alternative is determined by dominance structures (rather than by decision rules), the above recommendations do not apply with full force. This is because the recommended alternative would surely remain the one with the highest value, even if additional preference statements were to be acquired. But because it is difficult to know *in advance* if such dominance structures can be established, it is advisable to observe the above recommendations in all situations.

The structure of the elicitation process should be driven by an assessment to what extent the expected costs of preference elicitation are outweighed by the expected benefits (in terms of reductions in expected loss of value (4)). For example, if the number of alternatives is large, the problem can be first analyzed at high confidence levels (e.g., wide intervals which contain the 'true' parameters almost surely), followed by further analyses using lower confidence levels (e.g., intervals which contain the parameters at a 50 % confidence level; see Salo and Hämäläinen, 2001). This is because the statements at high confidence levels may eliminate several dominated alternatives so that further elicitation efforts can be focused on the remaining non-dominated alternatives. Conversely, if there are few alternatives which resemble each other, statements at a high confidence level are unlikely to eliminate any alternatives on the basis of dominance structures, wherefore it may be advisable to start with narrower intervals at a lower confidence level. The choice of confidence levels is thus an important decision in its own right.

One can also argue that the choice of decision rules should be contingent on the particular confidence level that is being applied. At high confidence levels (e.g., large intervals), it seems that alternatives should be excluded only on the basis of robust principles (i.e., dominance structures): otherwise, there is a possibility that further elicitation steps at a lower confidence level would give precedence to alternatives that might have been eliminated, if less restrictive decision rules had been applied earlier on.

4.2. Applicability of Preference Programming

The recognition of preference elicitation costs makes it possible to characterize contexts where preference programming methods are likely to be helpful. To begin with, it is plausible to assume that the costs of score elicitation are roughly proportional to the number of alternatives. Then, if the specification of the DM's preferences is relatively precise and the number of alternatives is large, even relatively incomplete score information may exclude several dominated alternatives from further consideration. This suggests that – rather than being seen as a tool for supporting final comparisons among competing alternatives – preference programming can serve as a front-end for conventional decision analyses. Here, the 'added-value' of preference programming comes from the ability to *prune* the number of relevant alternatives to a manageable few, thus ensuring that further analytical efforts are focused on the most relevant alternatives.

Preference programming also holds considerable potential in settings where reliable information on the DMs' preferences or the alternatives cannot be obtained, for one reason or another (i.e., the cost term in (4) is large). Examples of such settings include (i) the multi-criteria evaluation of preventive measures in the context of nuclear emergency management (Hämäläinen et al., 2000) and (ii) the evaluation of risk mitigation strategies with highly uncertain consequences, whereby the *precautionary principle* should be invoked (see, e.g., Stirling, 1999; Salo, 2001). Instead of recommending seemingly 'optimal' alternatives, the inherent limitations of available information in such settings must be recognized and reflected in a *strategic posture* which gives precedence to robust alternatives that perform reasonably well across a large range of permissible parameter values.

4.3. Dominance Structures, Decision Rules and Rank Reversals

Rank reversal is a phenomenon where the introduction or removal of a third alternative changes the relative ranks of two other alternatives. In effect, rank reversals imply that the DM's preferences for a given alternative depend not only on the properties of the alternative itself, but also on what other alternatives are included in or excluded from the analysis. Rank reversals have been a source of considerable controversy, and are regarded by many as a major flaw in decision support methodologies such as the Analytic Hierarchy Process (see, e.g., Belton and Gear, 1984; Salo and Hämäläinen, 1997).

It is therefore pertinent to ask which decision rules exhibit rank reversals. One may readily conclude that rank reversal cannot occur with decision rules in which the alternative's performance measure depends only on the properties of this alternative alone (in terms of its scores and fixed attribute weights). Clearly, no such *unitary* performance measure will change, even if additional alternatives are introduced or existing alternatives are removed. Several decision rules in Section 2 (e.g., maximax, maximin, central values, central weights) are therefore immune to rank reversals.

However, if a decision rule is based on the comparison of performance measures that depend on two or more alternatives, rank reversals may occur. Thus, for instance, the normalization of attribute-specific scores may cause rank reversals in the AHP (Belton and Gear, 1984; Salo and Hämäläinen, 1997). Another example is the Stochastic Multiobjective Acceptability Analysis (SMAA; Lahdelma et al., 1998) where the decision recommendation is based on the weight sets for which a given alternative is optimal. That is, the set $W(x^j) = \{w \in W | \sum_{i=1}^n w_i v_i^j \geq \sum_{i=1}^n w_i v_i^k \; \forall \; x_k \neq x_j\}$ (where $W = \{(w_1, \ldots, w_n) | \sum_{i=1}^n w_i = 1, w_i \geq 0, i = 1, \ldots, n\}$) consists of weights for which the overall value of alternative x^j is higher than or equal to that of all other alternative. This set is used to derive a performance measure – called the acceptability index $AI(x^j)$ – which is defined as the ratio between the volumes of $W(x^j)$ and W.

Now, assume that two alternatives x^1, x^2 are evaluated with SMAA with regard to two attributes using the score information $v^1 = (v_1^1, v_2^1) = (0.7, 0)$ and $(v_1^2, v_2^2) = (0, 0.6)$. Because the corresponding weight sets are $W(x^1) = \{w \in W \mid w_1 \geq \frac{6}{13}\}$ and $W(x^2) = \{w \in W | w_1 \leq \frac{6}{13}\}$, alternative x^1 has the larger acceptability index and is therefore the recommended alternative.

Next, assume that a third alternative x^3 with scores $v^3 = (v_1^3, v_2^3) = (0.6, 0.225)$ is introduced. After this addition, the weight sets become $W(x^1) = \{w \in W | w_1 \geq \frac{9}{13}\}, W(x^3) = \{w \in W | \frac{5}{13} \leq w_1 \leq \frac{9}{13}\}$ and $W(x^2) = \{w \in W | w_1 \leq \frac{5}{13}\}$. Alternative x^2 now has the acceptability index so that it becomes the recommended alternative, even though it was inferior to x^1 before the introduction of x^3: a rank reversal thus has occurred.

Although reports from applied work (see, e.g., Lahdelma et al., 2002) suggest that the DMs do feel comfortable with the SMAA method, one can raise also other concerns about its validity as a decision support methodology. In its original version, SMAA does not encourage the DM to explicate her preferences: thus, while the weight set $W(x^j)$ may be larger than $W(x^k)$ for all $x^k \neq x^j$, this set may consist of weights that are not aligned with the DMs (unstated) preferences. Hence, additional checks are needed to ensure that the weights underlying the acceptability index and the DM's preferences are compatible (see also Tavares, 1999).

Rank reversals may also occur when the decision recommendation is based on the minimization of maximum loss of value (MLV). To demonstrate this, assume there are two alternatives with scores $v^1 = (v_1^1, v_2^1) = (0.4, 0.6)$ and $v^2 = (v_1^2, v_2^2) = (0.6, 0.3)$, and that no weight information is available (i.e., the set of feasible weights is $W = \{(w_1, w_2)|w_1 + w_2 = 1, w_i \geq 0, i = 1, 2\}$). In this situation, the MLV of alternative x^1 is $\max_{w \in W}[(6-4)w_1 + 0w_2] = 2$ and that of x^2 is 3, i.e, x^1 is the recommended alternative. Now, if a third alternative with scores $v^3 = (v_1^3, v_2^3) = (0.8, 0.1)$ is added, the MLVs of the three alternatives become 4,3 and 5, respectively, indicating that x^2 becomes the preferred alternative, even though previously x^1 was the alternative. In much the same way, one can show that many other measures based on the comparison of value differences between alternatives (e.g, aggregated net intensity; see Kim and Ahn, 1999; Ahn et al., 2000) may exhibit rank reversals.

If rank reversals are deemed unacceptable, the above results suggest that the MLV concept should not be employed as decision rule. This concept can still be useful because it supports informed decisions as to when the elicitation phase might be terminated. That is, while unitary decision rules can be used to obtain recommendations which do not exhibit rank reversals, the computation of MLVs still helps the DM consider if additional elicitation efforts are warranted in view of possible improvements in decision quality.

5. TOWARDS A RESEARCH AGENDA

The discussion in the two preceding sections points to several research topics:

- **Matching methods to problems:** The explicit consideration of preference elicitation costs helps assess the likely advantages and disadvantages of preference programming in different contexts. From the viewpoint of applications, it is therefore important to develop guidelines for (i) identifying the qualitative properties of such contexts and for (ii) mapping these properties into implications for decision support process. For instance, recurring decision problems with modest stakes may be best supported by a relatively precise preference specification (e.g., narrow intervals) and decision rules that are based on estimates nearer the middle of these intervals (e.g., central values). But when the stakes are high and the development a complete preference specification is hampered by major uncertainties, wider confidence intervals and more ‘precautionary’ decision rules (e.g., maximin) may be called for. Moreover, preference programming methods can also be seen as a means of conducting *ex post* sensitivity analysis with regard to all model parameters (Mustajoki et al., 2003).

- **Computational efficiency analysis:** The development of above guidelines calls for simulation studies on which elicitation approaches and decision rules perform best, subject to different assumptions about problem size and costs of preference elicitation. Apart from the work of Salo and Hämäläinen (2001), the analyses by Barron and Barrett (1996) as well as Salo and Punkka (2003) fall into this stream of research, even though these analyses are concerned with ordinal preference information and few decision rules only. Nor have these earlier studies modelled elicitation costs in explicit terms, or sought to capture the staged nature of preference elicitation.

- **Development of elicitation strategies:** Further work should be devoted to the question of how preference information should be elicited from the DMs. Some advances have been made by structuring the elicitation process into structured sub-sequences (see, e.g., Mustajoki et al., 2003). Also, the PRIME Decisions software encourages the DM to complete several elicitation tours, each of which consist of a sequence of elicitation tasks (Gustafsson et al. 2001). Such 'building blocks' help in the development of elicitation strategies and the legitimate application of decision rules. Related research should also address to what extent preference programming approaches may mitigate behavioral biases in the elicitation attribute weights – or even create new ones (see, e.g., Weber and Borcherding, 1993).

- **Concepts and tools for group decision support:** Although many authors have noted that preference programming is suitable for group decision support (see, e.g., Hämäläinen and Pöyhönen, 1996; Kim and Ahn, 1999), there is a need for suitable preference elicitation concepts. Here, *a priori* agreements on principles of democratic decision making – analogous to the requirement for the equitable treatment of alternatives in Section 3 – can be used to derive constraints on the preference model: for example, the group may agree that no member shall have to yield more than a fourth of the value that she would get, if the group were to choose her preferred choice. After such cross-cutting constraints have been introduced, the group's aggregate utility may then be maximized, in the assurance that the resulting recommendation complies with the decision making principles that the group has set for itself. Similar principles can also be used to establish new decision rules: for instance, the recommendation may be selected by minimizing the relative loss of value that any group member has to endure (e.g., $\min_k \max_{w \in S_w} [V_k(x^{*,k}) - V_k(x^i)]/[V_k(x^{*,k}) - V_k(x^{\circ,k})]$ where $x^{*,k}$ and $x^{\circ,k}$ are the most and least preferred alternatives of the k-th group member).

- **Development of dedicated software tools and reflective case studies:** While conventional decision analyses can be carried out without dedicated software tools, such tools are indispensable for preference programming where the determination of dominance structures and decision rules often lead to extensive optimization problems. Even though there already exist several user-friendly tools (e.g., WINPRE, PRIME Decisions, RICH Decisions; see Hämäläinen, 2003), high priority much be attached to further tool development. Also, because the advantages and disadvantages of preference programming techniques are realized in the context of applications, the research agenda should be geared towards reflective case studies. Apart from the decision outcome, such studies should also analyze the impacts that preference programming has on the decision support process.

6. CONCLUSION

In this paper, we have discussed the salient features of preference programming methods for the incorporation of incomplete preference information in hierarchical weighting models. By building on illustrative models and experiences from reported applications, we have outlined guidelines for ensuring that the consecutive phases of preference elicitation and synthesis are carried out so that all alternatives are treated fairly and cost-effectively.

We have also discussed the advantages and disadvantages of preference programming in different decision contexts. Preference programming seems particularly useful in problems where (i) the elicitation complete information is either impossible or involves a very high cost; or where (ii) the elicitation of incomplete information during the early phases can be motivated by an attempt to prune dominated alternatives from a large set of alternatives. Our results also suggest that if the costs of information elicitation are explicitly accounted for, preference programming methods may outperform conventional approaches in terms of the overall cost-benefit characteristics of the decision support process. We conjecture that this result – together with the increasing availability of decision support tools – may contribute to the wider use of preference programming in the practice of decision analysis.

References

1. Ahn, B.S., Park, K.S., Han, C.H. and Kim, J.K. (2000). Multi-Attribute Decision Aid under Incomplete Information and Hierarchical Structure. *European Journal of Operational Research*, vol. 125, pp. 431–439.

2. Anandaligam, G. (1989). A Multiagent Multiattribute Approach for Conflict Resolution in Acid Rain Impact Mitigation. *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 19, pp. 1142–1153.

3. Arbel, A. (1989). Approximate Articulation of Preference and Priority Derivation. *European Journal of Operational Research*, vol. 43, pp. 317–326.

4. Barron, F.H. and Barrett, B.E. (1996). Decision Quality Using Ranked Attribute Weights. *Management Science*, vol. 42/11, pp. 1515–1523.

5. Belton, V. and Gear, T. (1983). On a Short-coming of Saaty's Method of Analytic Hierarchies. *Omega*, vol. 11, 228-230.

6. Belton, V. and Stewart, T.J. (2001). *Multiple Criteria Decision Analysis: An Integrated Approach.* Kluwer Academic Publishers, Boston.

7. Carrizosa, E., Conde, E., Fernández, F.R. and Puerto, J. (1995). Multi-Criteria Analysis with Partial Information About the Weighting Coefficients. *European Journal of Operational Research*, vol. 81, pp. 291–301.

8. Corner, J.L. and Kirkwood, C.W. (1991). Decision Analysis Applications in the Operations Research Literature, 1970–1989. *Operations Research*, vol. 39/2, pp. 206–219.

9. Cristiano, J.J. and White III, C.C. (2001). Application of Multiattribute Decision Analysis to Quality Function Deployment to Target Setting. *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 31/3, pp. 366–382.

10. Edwards, E. (1977). How to Use Multiattribute Utility Measurement for Social Decisionmaking. *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 7, pp. 326–340.

11. Edwards, E. and Barron, F.H. (1994). SMARTS and SMARTER: Improved Simple Methods for Multiattribute Utility Measurement. *Organizational Behavior and and Human Decision Processes,* vol. 60, pp. 306–325.

12. Eum, Y.S., Park, K.S. and Kim, S.H. (2001). Establishing Dominance and Potential Optimality in Multi-Criteria Analysis with Imprecise Weight and Value. *Computers & Operations Research*, vol. 28, pp. 397–409.

13. Gustafsson, J., Salo, A. and Gustafsson, T. (2001). PRIME Decisions: An Interactive Tool for Value Tree Analysis. In: M. Köksalan and S. Zionts (eds.), *Multiple Criteria Decision Making in the New Millennium, Lecture Notes in Economics and Mathematical Systems 507*, Springer-Verlag, Berlin, pp. 165-176.

14. Haze, G.B. (1986). Partial Information, Dominance, and Potential Optimality in Multiattribute Utility Theory. *Operations Research*, vol. 34, pp. 296-310.

15. Helenius, J. and Hämäläinen, R.P. (1998). *WINPRE - Workbench for INteractive PREference Programming*. Helsinki University of Technology, Systems Analysis Laboratory. Downloadable software (http://www.sal.hut.fi/Downloadables/winpre.html).

16. Horsky, D. and Rao, M.R. (1984). Estimation of Attribute Weights from Preference Comparisons. *Management Science* vol. 30/7, pp. 801–822.

17. Hämäläinen, R.P. (2004). Reversing the Perspective on the Applications of Decision Analysis. *Decision Analysis* (to appear).

18. Hämäläinen, R.P. (2003). Decisionarium – Aiding Decisions, Negotiating and Collecting Opinions on the Web. *Journal of Multi-Criteria Decision Analysis*.

19. Hämäläinen, R.P. and Pöyhönen, M. (1996). On-Line Group Decision Support by Preference Programming in Traffic Planning. *Group Decision and Negotiation*, vol. 5, pp. 485–500.

20. Hämäläinen, R.P., Lindstedt, M. and Sinkko, K. (2000). Multi-Attribute Risk Analysis in Nuclear Emergency Management. *Risk Analysis*, vol. 20/4, pp. 455–468.

21. Hämäläinen, R.P., Salo A. and Pöysti, K. (1992). Observations about Consensus Seeking in a Multiple Criteria Environment. *Proceedings of the Twenty-Fifth Hawaii International Conference on System Sciences*, Vol. IV, Hawaii, January 1992, pp. 190–198.

22. Keefer, D.L., Kirkwood, C.W. and Corner, J.L. (2004). Perspective on Decision Analysis Applications, 1990–2001. *Decision Analysis* (to appear).

23. Keeney, R.L. (1992). *Value Focused-Thinking: A Path to Creative Decisionmaking*. Harvard University Press, Cambridge, MA.

24. Keeney, R.L. and Raiffa, H. (1976). *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. John Wiley & Sons, New York.

25. Kim, J.K. and Choi, S.H. (2001). A Utility Range-Based Interactive Group Support System for Multiattribute Decision Making. *Computers & Operations Research*, vol. 28, pp. 485–503.

26. Kim, S.H. and Ahn, B.S. (1999). Interactive Group Decision Making Procedure under Incomplete Information. *European Journal of Operational Research*, vol. 116, pp. 498–507.

27. Kim, S.H. and Han, C.H. (1999). An Interactive Procedure for Multi-Attribute Group Decision Making with Incomplete Information. *Computers & Operations Research*, vol. 26, pp. 755–772.

28. Kim, S.H. and Han, C.H. (2000). Establishing Dominance Between Alternatives with Incomplete Information in a Hierarchically Structured Attribute Tree. *European Journal of Operational Research*, vol. 122, pp. 79–90.

29. Lahdelma, R., Hokkanen, J. and Salminen, P. (1998). SMAA - Stochastic Multiobjective Acceptability Analysis. *European Journal of Operational Research*, vol. 106, pp. 137–143.

30. Lahdelma, R., Salminen, P. and Hokkanen, J. (2002). Locating a Waste Treatment Facility by Using Stochastic Multicriteria Acceptability Analysis with Ordinal Criteria.*European Journal of Operational Research*, vol. 142, pp. 345–356.

31. Lee, K.S., Park, K.S., Eum, Y.S. and Park, K. (2001). Extended Methods for Identifying Dominance and Potential Optimality in Multi-Criteria Analysis with Incomplete Information. *European Journal of Operational Research*, vol. 134, pp. 557–563.

32. Mármol, A. M., Puerto, J. and Fernández, F. R. (1998). The Use of Partial Information on Weights in Multicriteria Decision Problems. *Journal of Multi-Criteria Decision Analysis*, vol. 7, pp. 322–329.

33. Mustajoki, J., Hämäläinen, R.P. and Salo, A. (2003). Decision Support by Interval SMART/SWING - Methods to Incorporate Uncertainty into Multiattribute Analysis. *Systems Analysis Laboratory, Helsinki University of Technology, Manuscript.* (Downloadable at www.sal.hut.fi/Publications/pdf-files/mmus03.pdf)

34. Park, K.S. and Kim, S.H. (1997). Tools for Interactive Decision Making with Incompletely Identified Information. *European Journal of Operational Research*, vol. 98, pp. 111–123.

35. Puerto, J., Mármol, A. M., Monroy, L. and Fernández, F.R. (2000). Decision Criteria with Partial Information. *International Transactions in Operational Research*, vol. 7, pp. 51–65.

36. Rios Insua, D. and French, S. (1991). A framework for sensitivity analysis in discrete multi-objective decision making. *European Journal of Operational Research*, vol. 54, pp. 176-190.

37. Saaty, T.L. (1980). *The Analytic Hierarchy Process.* McGraw-Hill, New York.

38. Saaty, T.L. and Vargas, L.G. (1987). Uncertainty and Rank Order in the Analytic Hierarchy Process. *European Journal of Operational Research*, vol. 32, pp. 107–117.

39. Salo, A. (1995). Interactive Decision Aiding for Group Decision Support. *European Journal of Operational Research*, vol. 84, pp. 134–149.

40. Salo, A. (1996a). On Fuzzy Ratio Comparisons in Hierarchical Weighting Models. *Fuzzy Sets and Systems*, vol. 84, pp. 21–32.

41. Salo, A. (1996b). Tighter Estimates for the Posteriors of Imprecise Prior and Conditional Probabilities. *IEEE Transactions in Systems, Man, and Cybernetics*, vol. 26, no.6 , pp. 820-825.

42. Salo. A. (2001). On the Role of Decision Analytic Modelling. In: A. Stirling (ed.), *On Science and Precaution in the Management of Technological Risk, Vol. II.* Institute of Prospective Technological Studies, Joint Research Centre of the European Commission, Report EUR 19056/EN/2, November 2001, pp. 123–141. (Downloadable at ftp://ftp.jrc.es/pub/EURdoc/eur19056IIen.pdf)

43. Salo, A. and Hämäläinen, R.P. (1992). Preference Assessment by Imprecise Ratio Statements (PAIRS). *Operations Research*, vol. 40, pp. 1053–1061.

44. Salo, A. and Hämäläinen, R.P. (1995). Preference Programming Through Approximate Ratio Comparisons. *European Journal of Operational Research*, vol. 82, pp. 458–475.

45. Salo, A. and Hämäläinen, R.P. (1997). On the Measurement of Preferences in the Analytic Hierarchy Process. *Journal of Multi-Criteria Decision Analysis*, vol. 6/6, pp. 309–319.

46. Salo, A. and Hämäläinen, R.P. (2001). Preference Ratios in Multiattribute Evaluation (PRIME) - Elicitation and Decision Procedures under Incomplete Information. *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 31/6, pp. 533–545.

47. Salo, A. and Punkka, A. (2003). Rank Inclusion in Criteria Hierarchies. *European Journal of Operational Research* (to appear).

48. Salo, A., Punkka, A. and Liesiö, J. (2003). RICH Decisions v1.01. Helsinki University of Technology, Systems Analysis Laboratory (http://www.rich.hut.fi/).

49. Stewart, T.J., Losa, F.B. (2003). Towards Reconciling Outranking and Value Measurement Practice. *European Journal of Operational Research* vol. 145, pp. 645–659.

50. Stillwell, W.G., Seaver, D.A. and Edwards, E. (1981). A Comparison of Weight Approximation Techniques in Multiattribute Utility Decision Making. *Organizational Behavior and Human Performance*, vol. 28, pp. 62–77.

51. Stirling, A. (1999).*On Science and Precaution in the Management of Technological Risk, Vol. I.* Institute of Prospective Technological Studies, Joint Research Centre of the European Commission, Report EUR 19056 EN, May 1999, 77 p. (Downloadable at ftp://ftp.jrc.es/pub/EURdoc/eur19056en.pdf)

52. Tavares, L.V. (1999). A Review of Major Paradigms for the Design of Civil Engineering Systems. *European Journal of Operational Research*, vol. 119, pp. 1–13.

53. Weber, M. (1985). A Method of Multiattribute Decision Making with Incomplete Information. *Management Science*, vol. 39, pp. 431-445.

54. Weber, M. (1987). Decision Making with Incomplete Information. *European Journal of Operational Research*, vol. 23, pp. 44–57.

55. Weber, M. and Borcherding, K. (1993). Behavioral Influences on Weight Judgments in Multiattribute Decision Making. *European Journal of Operations Research* 67, pp. 1–12.

56. White III, C.C., Sage, A.P. and Dozono, S. (1984). A Model of Multiattribute Decisionmaking and Trade-Off Weight Determination under Uncertainty. *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 14/2, pp. 223-229.