**LECTIO PRAECURSORIA** **11.1.2013**

**ANTTI PUNKKA**

Doctor Custos, Doctor Opponent, Ladies and Gentlemen

People make quick, unstructured decisions all the time and do just fine. Yet, some of us have spent a considerable amount of time on some decisions. These decisions can have had several objectives and constraints, long-lasting consequences, they might have been to some extent irreversible, or may have involved several stakeholders. Examples of such decisions can include purchasing a new apartment or a house, or even selecting how to spend a vacation. Apartments for example can differ in several attributes – such as price, location, size, condition, expected future renovation need and so on, but not all tens of thousands of apartments that are for sale fulfill your constraints concerning price, location, size, and so forth. From another perspective, we might be interested in the reasons behind a political decision. For this to be possible, this decision must be justifiable in some way.

Without any knowledge of formal decision analytic methods, many people have approached such decisions by listing the pros and cons of a few decision alternatives, or created a table where each of the alternatives is evaluated with regard to the same criteria with some suitable, verbal or numerical, scale.

I, for example, have found the latter way of decomposing the decision problem useful in many settings. For example, when choosing accommodation for a holiday trip, for me there are often too many constraints and objectives, such as minimization of price, minimization of the distance to the nearest ski lift or the nearest beach, maximization of the number of bedrooms, and so on to make the decision without systematically keeping track of the alternatives' properties. Finally, after familiarizing myself with a number of options, I might even be in a position to make trade-offs: Am I or are we willing to pay 300 euros to stay two hundred meters closer to the closest ski lift, when all other attributes considered are approximately equal? This list, including the attributes considered, can then serve as a starting point for the actual selection for me and my travel companions.

Many key principles of multi-criteria decision analysis are met with above-like simple procedures: recognize your objectives – or values – and constraints, construct a set of decision alternatives and choose attributes and suitable evaluation scales to measure the achievement of the objectives. Finally, evaluate the alternatives, elicit preferences and identify differences between the alternatives in a systematic way.

Methods, models and frameworks for multi-criteria decision analysis aim to make such procedures transparent, defensible and systematic, and to provide decision recommendations which make use of all relevant data and are in line with the decision maker's preferences.

Decision analysis is indeed prescriptive instead of descriptive; that is, it does not try to describe how people make decisions, but rather gives insights on how decision problems should be approached and solved. Many methods for decision analysis, however, have made use of the results of descriptive decision theory. For example, research on descriptive decision theory has revealed several common systematic errors and cognitive biases related to expressing preferences. These have later been acknowledged in the design of some preference elicitation methods.

The main results of a decision analytic process are that the alternatives become equitably and transparently evaluated, differences between them can be communicated in a justified way, and, furthermore and hopefully, the people involved in the process have common, well-defined definitions for the concepts used, including objectives, criteria, attributes and measurement scales.

Although building an evaluation framework is perhaps the most time-consuming and the most important phase of the decision support process, it is not sufficient for producing decision recommendations. In decision analysis, comparison of alternatives is fundamentally based on certain axioms of rationality. These axioms seem very obvious. They assume for example that preferences are transitive, meaning that if I prefer a hundred euros to 99 euros, and 99 euros to 98 euros, I should not prefer 98 euros to 100 euros. Another central assumption is independence of irrelevant alternatives: If I prefer hotel room A to hotel room B, but become aware of an inferior hotel room, which is very small and expensive, located far away from the ski lifts, and so on, I should not change my opinion between A and B due to my increased awareness concerning *other* alternatives. Though these conditions sound like something we could assume everyone to agree upon, there are many popular multi-criteria decision making methods that do not obey these axioms.

While the axioms of rationality are the foundation of decision analysis, modeling the decision maker’s preferences with a multi-attribute value function requires closer examination of the preferences. The definitions of the conditions that should be examined are not necessarily the easiest ones to understand. Most students who take the course exam in decision analysis here at Aalto University fail to define even the simplest of these, namely the condition of preferential independence. Furthermore, thorough examination of these conditions is quite laborious. Fortunately, an experienced decision analyst can be capable of recognizing possible violations of these conditions by looking at the model’s attributes and asking a few questions. Small changes in the attributes can then help fulfill these conditions.

Even so, value functions, especially additive value functions which this thesis focuses on, are widely used to evaluate and compare alternatives by people who have no clue about these underlying preference assumptions. Many, if not all of us, are familiar with multi-criteria product comparisons that can be found in various kinds of magazines. In these, the compared products are often evaluated with regard to the attributes on a numerical scale; from 0 to 10, for example. For each product, these numerical evaluations are then aggregated by first multiplying them by attribute-specific weighting factors called *attribute weights* and then summing these terms up to get the product’s overall value. The products are then rank-ordered based on these values. As another example, recent news about problems in competitive bidding of public infrastructure maintenance have strengthened my understanding that additive value functions are very widely used across the public administration. The users of these models just might call these models scoring models, or something else.

One possible reason for the popularity of additive models among practitioners is that the models themselves are mathematically very simple, but yet account for multiple objectives. It is indeed a very positive feature that these models can be readily applied, because the application of the mathematical model always requires the definition of objectives, attributes and alternatives. Applying the model is then just a small part of the whole decision support process.

However, preference elicitation can be seen as a part of applying the model and it can cause difficulties for practitioners who are not at all familiar with the underlying theory behind the models. One common, big mistake is to define the attribute weights without properly linking them to the attributes' measurement scales. For example, although it is common to say that the economy is more important than environmental factors, such a statement has no clear mathematical interpretation – at least in value theory. In fact, with a suitable choice of measurement scales, this can be interpreted to mean that 10 euros are more valuable than saving hundreds of endangered species from extinction.

The fundamental idea of value functions is indeed that they are a mathematical representation of the decision-maker's preferences. Therefore, it is essential that the decision maker's preferences can be captured by responses to readily understandable elicitation questions. In other words, the decision maker should fully understand the meaning of his preference statements.

Conventionally, the preferences are captured by specifying equally preferred changes over one attribute at a time, and equally preferred alternatives that often differ only on some attributes. I might for example be asked to state a price for a hotel room that is located immediately next to a ski lift so that I would be indifferent between this room and a room which costs 150 euros per night and is located 500 meters away from the nearest ski lift. My answers to such tradeoff questions together with preference statements on single attributes would eventually reveal my multi-attribute value function.

However, it was noted already in the early 1980s that giving such tradeoff statements can be time-consuming and difficult, and the obtained value functions can depend on the questions that are presented. For example, I might think that I would be willing to pay some 5-15 euros for staying next to the ski lift. Thus my answer to the aforementioned elicitation question would be between 155 and 165 euros. But as I have to describe my preferences with one number, I might choose to answer 160 euros. This would leave me unaware of how answering 164, 162 or 158 euros would have affected the results of the analysis.

Such considerations have motivated the development of methods that allow the decision maker to give incomplete information about his preferences. For example, I could respond to the above question by an interval: the price of such hotel room is between 155 and 165 euros. Or, equivalently, I could state that I would pay 155 euros for the hotel room that is next to the lifts, but I would not pay more than 165 euros for that room.

Theoretically, such statements lead to a setting, in which the decision maker's preferences cannot be represented by only one value function – or, to be precise, by one value function and its positive affine transformations which all represent the same preferences. Instead, there are infinitely many value functions that describe *different* preferences, but yet fulfill the decision maker's preference statements. The decision recommendations should then be based on examining all these value functions.

This thesis proposes novel kinds of preference statements to be used in preference elicitation by introducing the notion of incomplete ordinal information. This form of information allows the decision maker to give rank-based preference statements about the relative importance of the attributes, and about the alternatives' properties with regard to the attributes. For example, the decision maker can state that hotels D and E are the two most preferred ones with regard to the distance to the nearest ski lift, if information to make more complete statements is lacking. Or he can give preference statements between attributes by stating that reducing the hotel room price from 165 to 150 euros and decreasing the distance to the ski lifts from 500 meters to 0 are at least as valuable as some other, specified improvements.

These modeling possibilities can be particularly beneficial for a group of persons that seek to construct a joint value model, as they leave room for different opinions and do not necessarily require the application of preference statements that involve numerical evaluations.

Mathematically, incomplete preference information can be represented by a set of parameters so that these parameters correspond to the value functions that fulfill the preference statements. This parameter set truly is a playground: it is bounded, and often convex and closed by linear constraints. Furthermore, an additive value function is linear in these parameters.

The development of computers has made linear programming and Monte Carlo simulation standard tools in decision analysis, too. This has been very valuable from the perspective of decision support system development. The decision maker can start with loose preference statements and characterizations of alternatives and see how entering more restrictive statements changes the decision recommendations. The information provided by such interaction can help the decision maker to better understand the problem at hand, and can save the decision maker's efforts in obtaining more complete data about the alternatives than is needed to reach a decision recommendation.

However, adaption of these techniques has possibly shifted the focus of methodological research to a slightly wrong direction. In the 90's, the starting point of many scientific articles was no longer on a set of well-grounded preference statements obtained from comparing hypothetical alternatives. Instead, the starting point was that there are various forms of constraints between the parameters that represent the decision maker's preferences. No questions were asked concerning the origin of such constraints. Specifically, no one seemed to care what kinds of preference statements would lead to these constraints. For example, intervals of attribute weights have received much attention in the scientific literature, although the preference statements behind such constraints require – at least to my knowledge – comparisons of alternatives that differ in all attributes.

Linear programming and Monte Carlo simulation have found their use in producing decision recommendations, and carrying out sensitivity analyses, too. Different measures and indexes have been developed based on the examination of the parameter set that represents the decision maker's preferences. These indexes measure for example distances between preferences, or the share of preferences which recommend one alternative over another, and provide decision recommendations for example in a maximin solution sense; that is, choose the alternative for which the worst outcome, measured in terms of overall value, is the best.

In such methodological development, one fundamental property of value functions has been left unnoticed. The applied parameter set is merely one of the infinitely many ways to represent these preferences and these alternatives. This set is based on defining two, often hypothetical, alternatives, whose values are fixed to be constant throughout the set. The values of the other alternatives can then vary across the parameter set and they are always computed compared to the value difference between the two hypothetical alternatives.

This procedure, which is often referred to as normalization, is indeed required to get *numerical* overall values for the alternatives and to be able to define the attribute weights. However, it is not required for rank-ordering the alternatives or their value differences. These rank-orderings do not depend on the choice of the normalization. Nor do the *dominance relations* which are based on examining the mutual superiority of two alternatives with all value functions that fulfill the decision maker’s preference statements.

This thesis shows that many of the measures, indexes and recommendations that have been proposed in the scientific literature can be manipulated by changing the normalization of the value function. Ultimately, following the recommendations of these indexes is prone to rank reversals. That is, the recommendations can favor alternative A over B with some normalization, but favor B over A with some other normalization, although these normalizations represent exactly the same model, with exactly the same preferences.

As a partial solution to these shortcomings, this thesis proposes the computation of the alternatives’ ranking intervals. That is, if the preference information is incomplete and we consider all value functions that fulfill the preference statements, what rankings can a particular alternative have?

Although these intervals do not serve as decision rules, they provide easy-to-understand means to communicate how the alternatives perform over *all* the value functions that fulfill the preference statements. In addition, they are a well-grounded attempt to perform sensitivity analysis on the whole ranking of the alternatives: they can be used for example to study sensitivity of university ranking lists, which are often based on a naïve weighting scheme where most attributes are given equal weights. Because the ranking intervals communicate the best and worst possible rankings of the alternatives, they can also be used to support the selection of not one but many alternatives.

---

**I ask you Professor Robert Clemen, as the opponent appointed by the School of Science at the Aalto University to make any observations on the thesis which you consider appropriate.**