European Journal of Operational Research 242 (2015) 951–959

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Decision Support

The effect of bi-criteria conflict on matching-elicited preferences

Stéphane Deparis [a], Vincent Mousseau [b,*], Meltem Öztürk [c], Caroline Huron [d]

[a] IBM Research, Damastown Industrial Estate, Mulhuddart, Dublin 15, Ireland
[b] Laboratoire Génie Industriel, Ecole Centrale Paris, Châtenay-Malabry 92290, France
[c] LAMSADE, Université Paris Dauphine, Paris 75016, France
[d] INSERM Cognitive Neuroimaging Unit, CEA Neurospin, Gif-sur-Yvette F-91191, France

ARTICLE INFO

Article history:
Received 25 November 2013
Accepted 2 November 2014
Available online 20 November 2014

Keywords:
Multicriteria decision
Behavioural decision analysis
Preference elicitation
Multicriteria conflict
Matching

ABSTRACT

We focus on multicriteria preference elicitation by matching. In this widely employed task, the decision maker (DM) is presented with two multicriteria options, *a* and *b*, and must assess the performance value on one criterion for *b*, left blank, so that she is indifferent between the two options. A reverse matching, which is normatively equivalent, can be created by integrating the answer to the description of *b* and letting the DM adjust a performance value on the previously totally specified option *a*. Such a procedure is called a bi-matching. Consistency requires that isopreferences resulting from the forward and backward matchings be identical, but they empirically differ in a systematic direction. In a matching task, multicriteria conflict refers to the magnitude of the advantage or disadvantage to be compensated. We investigate the effect of the multicriteria conflict, or trade-off size, on the difference of judgement between forward and backward matchings. We observed that the difference of judgement is increased both by multicriteria conflict and by asking deteriorating rather than improving judgements at both steps of the bi-matching. We derive some implications for the practice of preference elicitation.

1. Introduction

Supporting a decision maker (DM) involved in decision processes often requires using preference models, which incorporate the value systems and judgements of the DM. Hence the concept of preference is crucial in decision making, and accounts for the way the DM evaluates and compares alternatives. The field of preference modelling has grown steadily over the past fifty years and numerous models have been proposed, such as utility-based models (Keeney & Raiffa, 1976) or outranking-based models (Roy, 1996). Preference models define rationality from a normative point of view, that is, norms to which the decision maker should conform.

Implementing preference models to support a DM involved in decision processes requires incorporating her judgements in the model. This calls for acquiring preferences through an interaction with the DM, and integrating these preferences in the model: this is called preference elicitation. Numerous preference elicitation methodologies, most of which are interactive methods (Steuer, 1986), have been proposed in the literature (Bana e Costa & Vansnick, 1994; Edwards & Barron, 1994; Jacquet-Lagrèze & Siskos, 1982). They allow the analyst to capture the DM's point of view and integrate it into a specific preference model. These preference models can be based on utility, but also on outranking relations (Mousseau & Slowinski, 1998) or decision rules (Greco, Matarazzo, & Slowinski, 2001). It is clear that the preference elicitation process is crucial for the preference model to represent faithfully the opinion of the DM, and lead to recommendations that can be viewed as reasonable and helpful by the DM.

However, preference elicitation methodologies implicitly hypothesize rationality principles to which the DM does not always conform. This can lead to elicited preference models that do not faithfully represent the DM's judgements. It is therefore important to understand decision behaviour fully through empirical studies in order to avoid misinterpretation of the DM preference statements and to avoid decision biases. For instance, preference elicitation methods usually assume "*procedure invariance which requires strategically equivalent methods of elicitation to yield the same preference order*" (Tversky, Slovic, & Kahneman, 1990).

Violations of this procedure invariance principle have been observed when using matching vs. pairwise comparison tasks (e.g., Fischer, Carmon, Ariely, & Zauberman, 1999; Fischer & Hawkins, 1993; Tversky, Sattath, & Slovic, 1988; Willemsen & Keren, 2002). In pairwise comparisons, the DM compares two alternatives and says whether she prefers one to the other or is indifferent between them. In matching tasks, two alternatives are presented to the DM, one with assigned evaluations on all criteria and the other missing a value for one

* Corresponding author.
E-mail addresses: stephane.deparis@ie.ibm.com (S. Deparis), vincent.mousseau@ecp.fr (V. Mousseau), ozturk@lamsade.dauphine.fr (M. Öztürk), caroline.huron.unicog@gmail.com (C. Huron).

http://dx.doi.org/10.1016/j.ejor.2014.11.001
0377-2217/

criterion. She is then asked to provide the missing value that makes the two alternatives indifferent. Experiments have shown that preferences elicited by matching and choice systematically differ since the more prominent dimension "*looms larger in choice than in matching*" (Tversky et al., 1988).

A positive feature of matching compared to other procedures is that it provides rich information on trade-offs from a limited number of questions (see Carmon & Simonson, 1998, for a discussion on this topic). For compensatory and attribute-based strategies, matching appears to be a natural elicitation procedure (e.g., Payne, Bettman, & Luce, 1998) and is widely used for decision-making under certainty or uncertainty (see for instance, Keeney & Raiffa, 1976 for utility assessment methods under uncertainty and Bana e Costa & Vansnick, 1994; Edwards & Barron, 1994 for utility assessment without uncertainty). Hence, matching is considered a suitable tool for preference elicitation because it offers a good compromise between the effort required from the DM and the obtained preference information.

Nevertheless, it is important to analyse empirically how DMs actually respond to matching questions, and whether their answers conform to rationality principles required by the preference models. For instance, suppose a DM assesses a matching between two alternatives a and b, providing a value on one criterion for alternative b (*forward matching*); a reverse task (*backward matching*), which is normatively equivalent, can be created by integrating the answer into the description of b and letting the DM adjust a performance value on the previously totally specified option a. This ordered pair of two matchings is called a *bi-matching* and the answers to this bi-matching should be consistent, i.e., isopreferences resulting from the *forward* and *backward* matching should be identical, or should not vary significantly. Previous experimental studies (Delquié, 1997; Willemsen & Keren, 2002, 2003) showed that DMs significantly deviate from such consistency in a bi-matching. In this paper, we extend these previous results and study the effect of the multicriteria conflict, or trade-off size, on the asymmetry (difference) of judgement between forward and backward matching questions.

In pairwise comparisons, multicriteria conflict arises when each alternative outperforms the other one on some criterion, which imposes a trade-off. Comparing two alternatives can be difficult if their respective advantages are of high magnitude (Deparis, Mousseau, Öztürk, Pallier, & Huron, 2012). In a matching task, the multicriteria conflict is related to how much the performances of the two options differ on the non-adjusted criterion. We propose an experimental setting to investigate the following questions: Is the observed difference in adjustment increased by the multicriteria conflict in a bi-matching task? Do the observed effects depend on the direction of the matching (evaluating an advantage or a disadvantage in the matching). As some elicitation procedures use matching questions with pairs of alternatives involving a strong conflict between criteria (e.g., Bana e Costa & Vansnick, 1994; Edwards & Barron, 1994; Keeney & Raiffa, 1976), we analyse the consequences of using such alternatives in matching questions when eliciting the preferences of a DM.

The paper is organized as follows. The next section is devoted to a description of the matching procedure. In Section 3, we formulate hypotheses about the effect of multicriteria conflict on the DM's answers. The experimental design is described in Section 4; the results are presented in Section 5 and discussed in Section 6. In the last section, we draw conclusions and propose further research avenues.

2. Eliciting preferences using matching

2.1. Matching procedure

Consider two alternatives a and b. We assume the decision problem only involves two criteria, called Criterion 1 and Criterion 2, that take real values (e.g. a salary in Euros). We write a_i to refer to the evaluation of alternative a on Criterion i. We indifferently write a or (a_1, a_2) to refer to alternative a in the following. When an evaluation is unspecified on any criterion, it is noted ?. For instance, $(b_1, ?)$ refers to an alternative b, whose evaluation on Criterion 2 is not yet defined. In a bi-criterion setting, a matching between two alternatives $a = (a_1, a_2)$ and $b = (b_1, ?)$, where the DM must specify a performance for b on Criterion 2, can be denoted as: $(a_1, a_2)\, I\, (b_1, ?)$, where I stands for "indifferent to". Matching can also be defined in the general case of n criteria by considering two alternatives differing on two of the criteria and equivalent in all the others, or considering two alternatives differing on more than two criteria. We restrict ourselves to the bicriteria case in this paper, because it allows to focus on the simplest, irreducible form of a trade-off. For the remainder of Section 2, we assume that preference increases with the value on both criteria. All the results hold when preference decreases with marginal value, as is the case in the experiment presented in Section 4.

In the above example, a is called the *stimulus alternative*, b the *adjusted alternative*. The criterion for which an evaluation is missing is called the *adjusted criterion* (Criterion 2 in the above example). Non dominance between the two alternatives implies that some values should not be possible for b_2: for instance, if $a_1 > b_1$, the DM will choose b_2 greater than a_2.

Matching questions can be distinguished according to the *direction* of matching. As an example, let us consider the following elicitation: *Consider a decision between jobs that vary according to the annual salary (expressed in Euros) and the number of vacation days per year. Suppose you currently hold a job with 20 days vacation, and a 30,000 Euros salary.* In the following question, an unfavourable matching is asked: *A new position is available with 25 days vacation, but an inferior salary. What loss in salary would make that new job equally satisfying to your current job?* A favourable matching can alternatively be asked: *A new position is available with only 15 days vacation, but a better salary. Which salary compensation would make that new job equally satisfying to your current job?*

The *direction* of matching is defined as such (see Fig. 1):

- *Favourable matching*: The DM estimates an advantage with respect to the adjusted criterion. When $a_1 > b_1$ (as in Fig. 1a), adjusting on Criterion 2 comes down to evaluate the proper advantage of b over a on Criterion 2 ($b_2 > a_2$), in order to compensate for the inferior evaluation on Criterion 1.
- *Unfavourable matching*: The DM estimates a disadvantage with respect to the adjusted criterion. When $a_1 < b_1$, (as in Fig. 1b) adjusting on Criterion 2 comes down to evaluate the proper disadvantage of b compared to a on Criterion 2 ($b_2 < a_2$), in order to compensate for the superior evaluation on Criterion 1.

Another important characteristic of a matching question is related to the intensity of the multicriteria conflict involved. We define the *conflict* of the matching $(a_1, a_2)\, I\, (b_1, ?)$ as the difference of evaluation between the two matched alternatives on the non adjusted criterion, i.e. $|a_1 - b_1|$. Note that, in order to compare the conflict associated to two matching questions, the adjusted criterion should be the same for the two questions.

2.2. Bi-matching

As we explained in the Introduction, a bi-matching is composed by a sequence of two consecutive matchings that we call *forward matching* and *backward matching*.

In the following paragraph, a performance value on a criterion is underlined when it was elicited from the DM, and an alternative is underlined when any of its performance was elicited from the DM. Table 1 presents an example of bi-matching where $A = (a_1, a_2)$ is the stimulus alternative in the forward matching, $\underline{B} = (b_1, \underline{b_2})$ is the adjusted alternative in the forward matching but the stimulus alternative in the backward matching, $\underline{A} = (a_1, \underline{a_2})$ is the adjusted alternative in the backward matching. Notice that a bi-matching can involve any