Journal of the Operational Research Society (2009) 60, 268–275 0160-5682/09

www.palgrave-journals.com/jors/

Behavioural decision theory for multi-criteria decision analysis: a guided tour☆

A Morton* and B Fasolo

London School of Economics and Political Science, London, UK

Multi-criteria decision analysis (MCDA) involves asking decision makers difficult questions, and can leave them thinking that their judgements are not as coherent as they might have thought. This experience can be distressing and may even lead to rejection of the analysis. The psychology of preference sheds light both on how people naturally make choices without decision analytic assistance, and on how people think about the MCDA elicitation questions. As such, it can help the analyst to respond helpfully to difficulties which decision makers may face. In this paper, we review research from Behavioural Decision Theory relevant to MCDA. Our review follows the MCDA process, discussing research relevant to the structuring, value elicitation, and weighting phases of the analysis, outlining relevant and important findings, and open questions for research and practice.

Journal of the Operational Research Society (2009) **60,** 268–275. doi:10.1057/palgrave.jors.2602550
Published online 9 January 2008

Keywords: decision analysis; multi-criteria decision analysis; behavioural decision theory; psychology; decision biases; experiments; problem structuring

Introduction

Operational researchers often believe that part of the contribution they can make to decision making is by helping clients to define what they are trying to achieve. This is especially true of practitioners of multi-criteria decision analysis (MCDA), whose discipline is predicated on the notion that clients can benefit from systematic preference modelling. As many writers have recognized, this activity raises numerous questions about how decision makers make value judgements, psychologically speaking. The sub-discipline of psychology in which this is studied is known as Behavioural Decision Theory (BDT).

We believe that understanding BDT can be helpful when undertaking MCDA interventions, for two reasons. Firstly, by deepening our understanding of how people might approach decisions without assistance, BDT helps us explain what MCDA models can do. Empirical examination of the decision rules which people actually use to make decisions reveals that, much of the time, people unconsciously select from a set of rather simple heuristics, such as lexicographic rules, which avoid the need to make trade-offs between conflicting objectives (Payne *et al*, 1988, 1993).

Secondly, BDT also sheds light on biases which may colour clients' responses to elicitation questions and interpretation of the models. Consistency checking and validation of model form and parameters are critical components of any MCDA intervention. Although the BDT literature tends not concern itself directly with consistency checking and validation as such, a typical finding is that responses to the same underlying questions can vary dramatically as the questions are re-contextualized (eg Slovic, 1995). As such, the literature provides a resource for analysts for ways to feedback elicited data and to challenge decision makers to reflect more deeply.

The BDT literature divides, roughly and with considerable overlap, into a literature on choice under uncertainty and choice under conflicting objectives. This distinction maps to a parallel distinction in the decision analytic literature, between probabilistic (eg decision trees, influence diagrams) and multi-criteria modelling tools (eg Multi-Attribute Value Theory (MAVT)). The aim of the current paper is to present a view of BDT research relevant to MCDA, focussing particularly on the MAVT approach.

MAVT is a widely used and practical MCDA approach. It is based on the supposition that preferences can be represented by a value function taking an additive form: the value of an option a can be represented as $v(a)=\sum_j w_j v_j(a)$ where $v_j(\cdot)$ is a cardinal 'partial value function', yielding the value which is obtained from a with respect to its performance on criterion j and w_j is a criterion weight. The assumption of an additive form is a crucial part of the craft of modelling with MAVT, as it disallows interactions in value between criteria. Failure to develop criteria which are genuinely independent can be a serious trap for the unwary analyst. However, conditions for validation of the appropriatness of the additive form are well

*Correspondence: A Morton, Operational Research, London School of Economics and Political Science, Houghton St, London WC2A 2AE, UK. E-mail: a.morton@lse.ac.uk

☆This paper is a revised version of a Keynote Paper given in the Decision Analysis stream at Young OR15 in Bath, March 2007.

known from the theoretic literature (Fishburn, 1970; Keeney and Raiffa, 1976; von Winterfeldt and Edwards, 1986; Belton and Stewart, 2002) and we do not review them here.

The purpose of this paper is to present a ‘guided tour’ of BDT relevant for MCDA. As the classic work of von Winterfeldt and Edwards (1986) shows, the two fields of BDT and decision analysis (DA) share common origins. Over recent years they have however increasingly grown apart, and our aim is to highlight and recast the connections between BDT and DA, in particular MCDA. Our coverage is avowedly partial: we have chosen topics which we find interesting and insightful, and where we see a clear relevance to application and practice. To aid our exposition, we develop a running example and explain the phenomena we discuss in some detail. However, we do not attempt in this paper an elementary presentation of MCDA: for this, see a standard text (Belton, 1990; Clemen, 1996; Belton and Stewart, 2002; Goodwin and Wright, 2004). We also do not offer an overview of BDT; for this see a range of excellent BDT textbooks (Baron, 2000; Hastie and Dawes, 2001; Beach and Connolly, 2005; Newell *et al*, 2007). Unlike these books, this ‘guided tour’ explicitly focusses on the challenges facing the practising analyst. Excellent review papers exist, in particular those of Payne *et al* (1999) and von Winterfeldt (1999). Our exclusive focus on MCDA and coverage of recent developments in BDT differentiate us from these, and in any case, our perspective differs so our coverage only partially overlaps with that of these works.

In this paper, we consider three phases of the MCDA process, structuring, assessing values, and weighting criteria. We omit sensitivity analysis not because it is either uninteresting or unimportant, but because there is no space to do it justice here. For each phase, we discuss some of the relevant and important findings, as well as open questions for research and practice. However, in a paper of this length it is impossible to be anything other than selective and we encourage the interested reader to use our inevitably idiosyncratic selection as leads for further reading.

Structuring

Systematic behavioural research on the structuring of MCDA models is rather thin on the ground, despite general agreement that structuring is a crucial part of analysis and evidence of considerable diversity in practice among working analysts (Buede, 1986; Belton, 1990; French *et al*, 1998). In 1988, Borcherding and von Winterfeldt remarked that ‘value structuring is still largely an art, left to the skill of the individual analyst’ (p 154) and the situation has not materially changed since that time.

The main tool which is used in the structuring of MCDAs is the value tree, a hierarchical representation of the decision makers’ objectives. An illustrative value tree to assist in flat choice (based on the authors’ composite experience of looking for a place to stay in London) is shown in Figure 1. Criteria are partitioned into two sets, *Costs* and *Benefits*, the former dealing with monetary outlay, and the latter with qualitative attributes of the flat as a place to live, such as spaciousness, character, and closeness to the authors’ and their partners’ places of work.

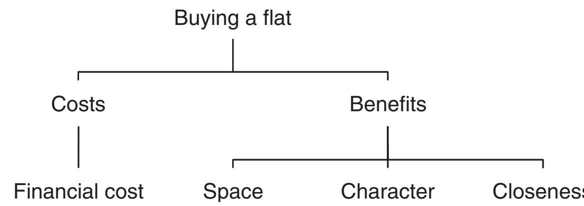

Figure 1 Value tree for buying a flat.

Developing a value tree implicitly assumes that decision makers can generate all the relevant objectives. Recent experimental work (Bond *et al*, 2007) casts doubt on this, suggesting that when people are asked for their objectives in some substantive decision problem, the number which they generate is much smaller than (typically, about half) the number of objectives which they can recognize as relevant to them from a list. This provides some support for the practice, fairly common among working analysts, of using prompts based on some generic checklist to stimulate thought on the part of decision makers.

For large problems, value trees can be quite complex, and classically, there are two main ways to structure the elicitation of criteria: ‘top-down’ (start with high level objectives, and work downwards asking, eg ‘how would you achieve that?’) and ‘bottom-up’ (start with alternatives, ask for attributes which differentiate them in ways which the decision maker cares about, and work upwards from there). Some methodologies, such as Keeney’s Value Focussed Thinking (Keeney, 1992), have a strong top-down orientation, whereas others, such as the view which informed the software MAUD (Humphreys and McFadden, 1980), have a much stronger bottom-up flavour.

Research suggests that there may be some differences in the shape of the value trees elicited by top-down and bottom-up approaches (Adelman *et al*, 1986). While interesting, this seems to us only one of the questions which one might want to ask about the difference between top-down and bottom-up structuring. For example, top-down structuring may tend to produce sets of criteria which are fairly general, but which may be hard to relate to particular option sets, whereas bottom-up structuring seems geared to produce sets of criteria which are very specifically relevant to the problem at hand, but which are not ‘portable’ to new decision situations. There is scope for more research on this sort of question.

An interesting development, particularly in the UK decision analytic community, has been the use of SODA (Strategic Options Development and Analysis), a problem structuring method, in tandem with value tree structuring (Ackermann and Belton, 1994; Bana e Costa *et al*, 1999; Montibeller and

Belton, 2006). Loosely, this involves developing a representation of the thinking of a particular decision maker or decision-making group as a qualitative network of ideas (a 'cognitive map' or 'group map'), and then extracting out of it a value tree suitable for performing value measurement. It may be that this represents a third approach to be set alongside top-down and bottom-up ('outside-in'?), and once the methodology for doing the value tree extraction has stabilized, working out how to do the three-way comparison may be an interesting topic for study.

Assessing values

Any MCDA intervention involves some sort of quantitative value assessment, that is, numbers are assigned to different options (or more generally, attribute levels of options) which represent their valuation on some particular criterion. It is conventional to write this valuation as a function $v_j(\cdot)$ over an option or attribute space where j is a criterion index. This scale may be anchored at 0 and at 100 by the least and most preferred options for criterion j (a 'local scale'); or it may be anchored by some naturally or psychologically significant levels of performance (eg an aspiration level, or a level corresponding to an absolute lower acceptable level on performance). An example of some criterion scores on a local scale for the flat choice problem is shown in Table 1.

Historically, a central issue in value assessment has been whether the values to be assessed should be von Neumann–Morgenstern utilities (established using questions about equilibrating gambles) or measurable value functions (established by asking directly about strength of preferences). Although value/utility assessment is still an active area of research, we focus in this paper on the assessment of measurable value functions, which, as Goodwin and Wright (2004) note, are widely used in practice.

One family of biases in valuation, which has been much studied in relation to economic willingness to pay studies, relates to an apparent scope insensitivity (Kahneman and Knetsch, 1992; Schkade and Payne, 1994; Frederick and Fischoff, 1998). The classic studies in this domain show that people's responses to willingness to pay questions do not seem to vary as much as one would expect with the quantity of what is being paid for. People express a readiness to pay quite similar amounts to save, for example, 2 000, 20 000, or 200 000 wildfowl in the context of setting pollution controls.

Table 1 Criterion scores for the flat choice problem

	Criteria			
Flat	*Cost*	*Closeness*	*Character*	*Space*
1	0	100	100	65
2	50	0	100	65
3	95	0	0	100
4	100	0	0	0
5	50	70	0	65

We, in common with other analysts, for example, Jones *et al* (2005) and Kleinmunz (2007), have found similar effects in work in judgemental evaluations of capital investment projects. The valuation of such projects often takes place on a 'semi-global' scale, with a fixed lower end-point or 0 at a baseline level of activity (we thank one of our referees for suggesting the term 'semi-global' to describe this). In this context, value scores reflect an idea of contributing added value on top of that baseline on some given criterion. The 100 level is set at the value of a standard, typically, the most valuable project on that criterion. Valuation of the remaining projects proceeds by direct ratio estimation (von Winterfeldt and Edwards, 1986), with decision makers being asked how value accrues from each project as a fraction of the value of the standard. A favoured method for prioritizing projects is to compute benefit/cost ratios and rank projects in value for money order (Kleinmunz, 2007; Phillips and Bana e Costa, 2007). As the benefit/cost ratio has a marginal benefit interpretation, there is a normative rationale for this procedure (Weinstein and Zeckhauser, 1973).

A common pattern of response is that when projects are evaluated, the benefit/cost ratios of small (ie low cost) projects are relatively high, and those of large (ie high cost) projects are low. This, by itself, would not be exceptional since there may be plausible reasons why small projects tend to be relatively efficient in a benefit/cost sense: for example, if a portfolio contains projects at varying stages in development, one would expect projects relatively far through their lifecycle to have relatively high values, as otherwise they would have been stopped earlier, but also relatively low costs remaining, as much of the cost will have already been incurred (we are grateful to Larry Phillips for making this point to us). However, decision makers often reject the benefit–cost ratio ordering which is implied by these judgements on the grounds that it seems disproportionately to favour small projects. This clash between holistic and decomposed judgement suggests that the original valuations may fail to take scale into account satisfactorily, in a manner similar to the scope insensitivity bias exhibited in willingness to pay studies.

We show in Figure 2 two sets of data from judgemental evaluations of scientific projects conducted by the R&D division of a large high technology company. The data shown in Figure 2(a) relate to valuations on a financial returns criterion ('NPV'), and the data shown in Figure 2(b) relate to valuations on a more qualitative criterion ('Strategic Fit'). In both cases the value scores have been divided by cost and the resulting benefit/cost ratios plotted against cost. This analysis shows that scope insensitivity may be present in both data sets, although more obvious in the latter, as the superimposed regression line shows.

This difference in the pattern of responses for the more and less readily quantifiable criteria may illustrate something important about the psychological mechanism which

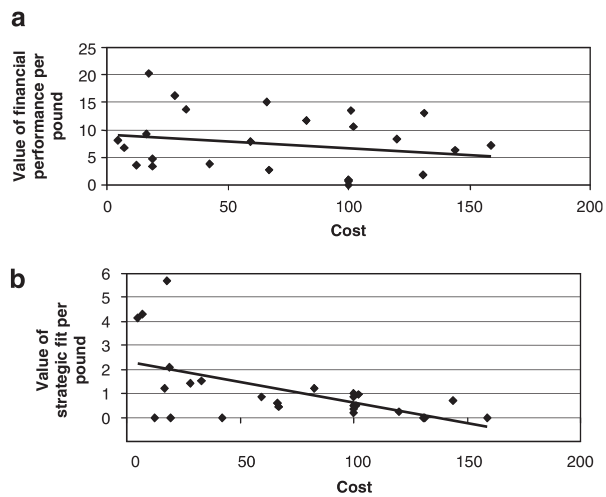

Figure 2 (a) Value of financial performance per pound against cost and, (b) Value of strategic fit per pound against cost. Examples of scope insensitivity in the evaluation of scientific projects.

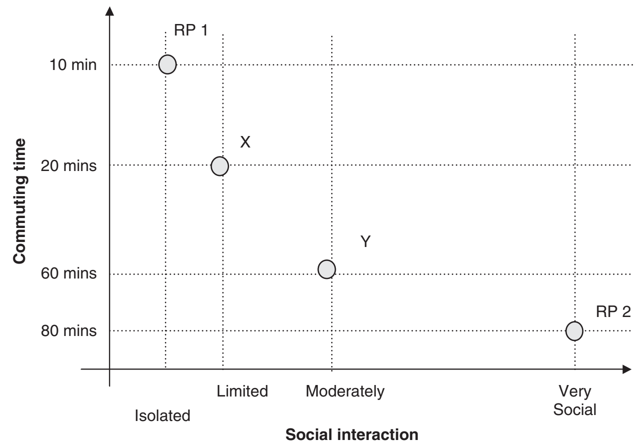

Figure 3 Illustration of Tversky and Kahneman (1991)'s reference dependence experiment.

underlies scope insensitivity. Hsee and Rottenstreich (2004) have suggested that the scope insensitivity biases arise when people rely heavily on affective response (ie feeling) in evaluation tasks. Of course, feeling is an absolutely crucial part of valuation: indeed, things are only good or bad insofar as someone somewhere feels them to be so. But feeling does not scale, and most people do not, internally, feel one hundred times as bad at the loss of 200 000 birds as at the loss of 2 000. As a result, for coherent valuation, feeling has to be supplemented with explicit quantitative thinking which Hsee and Rottenstreich refer to as 'calculation'. Based on these insights, we have recently been performing a series of experiments to see whether and what kind of alternative modes of elicitation may result in different patterns of scope insensitivity bias. Results to date are reported in two MSc projects (Morgan, 2006; El-Bassunie, 2006).

Another bias which has been widely documented is the so-called *status quo* or reference point effect (Kahneman and Tversky, 1979; Kahneman *et al*, 1991). Typically, people seem to dislike losses more strongly than they like comparable gains. A well-known experiment which relates particularly easily to the MCDA context is reported in Tversky and Kahneman (1991), which we illustrate in Figure 3. Subjects were presented with a choice between two jobs, *X* and *Y*, which were distinguished on two dimensions, the amount of social interaction and the commuting time. In one version of the experiment, subjects were told that they were currently in a job which performed well on commuting time and badly on social interaction (RP1); in the other version, the *status quo* job had the opposite characteristics (RP2). This simple manipulation flipped the majority preference: of those presented with the version 1 choice, 70% chose job *X*, whereas in the version 2 choice, only 33% did so, illustrating that the difference in social interaction between jobs *X* and *Y* is less (more) valued relative to the difference in commuting time when the former (latter) is perceived as a gain relative to the *status quo* and the latter (former) is perceived as a loss.

The existence of such effects are clearly relevant to MCDA, but drawing operational conclusions about practice is difficult. For example, in multi-criteria resource allocation (Kleinmunz, 2007; Phillips and Bana e Costa, 2007), we observe that some analysts prefer to define a zero-based baseline level of activity ('do nothing') and then add back activity on top of that baseline, whereas others define the baseline as the *status quo* and represent some options as investments and others as disinvestments. This difference in practice raises psychological questions. Setting a do-nothing level of activity plainly de-emphasizes the reference point, but does it perhaps destroy key intuitions which the decision maker may have about the problem? Defining projects as gains and losses should make any bias explicit and open to critique, but does it also make the reference point more salient in the mind of the decision maker?

A third bias which has been observed in applications where people are asked to make multi-criteria assessments is an excessive intercorrelation of criterion scores (positive when the criteria are characterized by the same direction of preference; negative otherwise; Alhakami and Slovic, 1994; Pidgeon *et al*, 2005). This has been attributed to the operation of an 'affect heuristic' (Slovic *et al*, 2007) whereby people consider first whether the object of assessment is 'good' or 'bad', and then evaluate the supposedly disaggregate criterion scores on the basis of this overall assessment. One would like to think that MCDA procedures, by encouraging decision makers to think through problems in a more explicitly disaggregate way, can 'debias' for this effect, but we know of no empirical evidence with a bearing on this hypothesis.

Finally, we note a respect in which behavioural research has been lagging decision analytic practice. Contemporary DA tools often incorporate elicitation aids which allow the decision makers to enter partial information about preferences (eg interval values or ordinal statements about value differences). Although such information may not lead immediately to a decision, it may allow dominated alternatives to be screened out, and so can help the decision maker focus on the key elements of the problem (Rios Insua and French, 1991; Bana e Costa and Vansnick, 1995; Salo and Hämäläinen, 1995).

For example, the evaluation technique MACBETH (Bana e Costa and Vansnick, 1995, 1997; Bana e Costa *et al*, 1999, 2004) requires decision makers to express preferences in terms of pairwise verbal comparisons rather than numbers (eg 'do you perceive the difference on Criterion A between Option X and Y to be non-existent, very weak, weak, moderate, strong, very strong, or extreme?'). This mode of elicitation is similar to that of the Analytic Hierarchy Process. However, MACBETH is an MAVT-based method in that it uses the elicited information to narrow the space of value functions which could characterize decision makers' preferences, which in turn enables it to propose plausible cardinal value scales, screen out options as dominated under varying assumptions about the level of measurement, and provide bounds for sensitivity analysis.

We feel that choice of evaluation technique should exploit the natural capabilities of decision makers and in ongoing research, one of us (Fasolo and Bana e Costa, in preparation) is investigating how preference for verbal scoring (as encouraged by MACBETH) and numerical scoring interact with the decision maker's level of verbal fluency and 'numeracy', the ability to use appropriate numerical principles (Lipkus *et al*, 2001), where numeracy is operationalized as the number of correct responses to a test we devised to measure comprehension and use of numerical information. The experimental findings suggest that numerical scoring is found easier by relatively more numerate decision makers, whereas the verbal scoring is found more manageable by relatively more verbally fluent and less numerate decision makers.

Weighting criteria

Once criterion scores have been established, the next stage of the MCDA process is to weight the criteria. Weighting presupposes that it is possible to represent the decision maker's preferences as a weighted additive sum of the individual criterion scores. The qualitative characteristics of the decision maker's preferences which make this representation plausible have been extensively explored in the normative literature over the years (Fishburn, 1970; Keeney and Raiffa, 1976; von Winterfeldt and Edwards, 1986). In MAVT, weights represent scaling factors or 'exchange rates' which bring the individual criterion value scores to a common value scale. Normalised weights are shown for the flat choice problem in Table 2, together with computed overall evaluations.

For most decision makers, weighting criteria is the most cognitively demanding part of the MCDA process. A feature of the weighted additive model is that it is explicitly compensatory, meaning that deficiencies in one criterion can be offset by strengths in others (compare this with a lexicographic approach where options which do not enjoy the highest level of performance on some designated 'most important' criterion are screened out). Although people only express trade-offs reluctantly (Hogarth, 1987), experimental research suggests that people can use compensatory approaches even without analytic support (eg Bettman *et al*, 1993), but this is cognitively demanding, and only possible when there is ample time and the trade-offs do not stir strong emotional reactions (Luce *et al*, 1999).

Table 2 Weights and overall values for the flat choice problem

	Criteria				
	Cost	*Closeness*	*Character*	*Space*	
Flat swing weight	0.38	0.32	0.11	0.19	*Overall value*
1	0	100	100	65	56
2	50	70	100	65	65
3	95	0	0	100	55
4	100	0	0	0	38
5	50	70	0	65	54

Even under relatively favourable conditions where analytic support is available, weight judgements exhibit predictable biases. For example, weight judgements seem to depend on the structure of the value tree. If a criterion is 'split' into subcomponents (eg if the criterion 'character' was subdivided into two subcriteria relating to the attractiveness of the neighbourhood and the aesthetic appeal of the physical building), the resulting (aggregate) weight assigned to that criterion tends to be increased (Weber *et al*, 1988; Weber and Borcherding, 1993; Pöyhönen and Hämäläinen, 1998; Pöyhönen *et al*, 2001). This bias arises from a general tendency for assessments to be 'partition dependent' (Fox *et al*, 2005): people assign quantities (eg weights, values, probabilities) relatively evenly across a 'partition' presented to them.

Of crucial importance is the weighting question which is asked. The simplest (and most naive) approach is to ask for direct judgements of relative importance: how important is 'space' relative to 'closeness'? However, alternative elicitations modes exist, such as 'swing weighting' (Goodwin and Wright, 2004) in which the differences in value between the levels of a most and least preferred option on two given criteria ('swings') are made salient and respondents are explicitly requested to consider the relative value of the swings. A second alternative is 'tradeoff weighting' (von Winterfeldt and Edwards, 1986, p 289), where respondents are requested to find an upper level of some underlying attribute which would equilibrate swings on two designated criteria (eg an upper level of 'square footage' which equilibrates the swing on 'space' with the swing on 'character').

The research literature shows that there are systematic and persistent differences in the numbers which decision makers assign when their weights are elicited with the different methods (Weber and Borcherding, 1993). But also, and more interestingly, it shows that weights elicited by different methods vary in 'range sensitivity'. Range sensitivity is a normative property which exploits the dependence of criterion weights to criteria ranges: if one particular criterion has a small range and does not discriminate between options (eg in London, all flats tend to perform quite badly in terms of spaciousness) then the difference in value between the most and least preferred options, and hence the criterion weight, should be rather small, even if the criterion is felt to be (in some global, absolute sense) 'very important'. Generally, direct importance judgements are highly range *in*sensitive, and judgements elicited through questions like swing and tradeoff weighting are rather less so, although still not as range sensitive as the normative model would predict (von Nitsch and Weber, 1993; Fischer, 1995). The conclusion for DA would seem to be that while no method is perfect, the more sophisticated methods are to be preferred to direct importance weighting.

People's ability to understand and respond to direct importance questions baffles us. Both of us regard such questions as meaningless and refuse to answer them. It could be that our intuitions have been destroyed by our atypically extensive exposure to DA. Perhaps direct importance questions *do* in fact tap into something psychologically real and meaningful, although not a criterion weight as one would think of it in conventional MCDA terms. Intriguingly, some analysts (eg Bana e Costa, personal communication) claim that it is possible to define global scales anchored by psychologically primitive notions of 'good' and 'neutral' and that if scales are defined in this way, swing or trade-off weights correspond closely to direct importance weights. We know of no direct tests of this claim, but for some interesting reflections along these lines, the reader is referred to the work of William Goldstein and his colleagues (Goldstein and Beattie, 1991; Goldstein and Mitzel, 1992; Goldstein *et al*, 2001).

Conclusion

In this paper, we have presented a selective review of behavioural findings relevant to the practice of MCDA. We have grouped these under three headings: structuring, assessing values, and weighting criteria. Under *structuring*, we have discussed decision makers' ability to generate criteria, and constrasted three approaches to structuring: bottom-up, top-down, and 'outside-in'. Under *assessing values*, we have discussed the scope insensitivity and *status quo* biases, described the operation of the affect heuristic, and outlined what we see as the challenges posed by a new generation of assessment tools, such as MACBETH. Finally, under *weighting criteria*, we have discussed a partition dependence bias (known in the context of weight assessment as the 'splitting bias') and the behavioural properties of different weighting questions.

One of the characteristics of MCDA is that it can reveal to decision makers that their thinking is not as coherent as they might have thought. While we see this as one of the key elements of the value proposition of DA (and indeed, of OR in general), it can be distressing for decision makers and can lead them to reject the method of analysis. A greater understanding of BDT can help analysts to engage with decision makers in a sympathetic and constructive fashion, reflecting back the logical implications of judgements to test robustness. For example, how does a decision maker's specific criteria for a given decision compare with a generic list of criteria for that broad decision type? How does a set of value scores appear when divided by costs and transformed to value for money? Does a given set of weights match up with or clash with pre-analytic judgements of the relative importance of a set of criteria? A psychologically aware MCDA practitioner is in an excellent position not only to identify these as worthwhile questions to ask, but also to isolate and explain clashes between the model and intuitive judgement.

Recent work has seen MCDA used in very high profile applications, such as the study by the UK's Committee on Radioactive Waste Management's analysis of disposal options for the UK's radioactive waste (CoRWM, 2006; Phillips, 2006; Morton, Airoldi and Phillips, in preparation). With the increasing demand for transparent approaches to the management of complex societal problems, we expect the scope for the deployment of decision analytic approaches to grow. Yet many of the stakeholders in these high-profile applications distrust formal modelling, which they see as technocratic and alienating, and instead cling to a belief in the capacity of the unaided human intellect to integrate complex information which is unsupported by anything we know of in the psychological literature. Understanding the psychology of preference is, in our view, essential if both Multi-Criteria Decision Analysts and Operational Researchers are to persuade a sceptical audience that their techniques are helpful.

Acknowledgements—This paper has grown out of our teaching together on courses *OR304 Decision Analysis in Theory and Practice*, *OR425 Decision Analysis in Social Context*, and Barbara Fasolo's teaching on *OR405 Advanced Topics in Behavioural Decision Theory*, and we are indebted to our co-teachers, Larry Phillips, Carlos Bana e Costa, Val Belton, Daniel Read, and Greg Davies for numerous stimulating discussions. We thank Larry Phillips for supplying us with the data on which Figure 2 is based, and Gilberto Montibeller for helpful comments. We are grateful to Ana Barcus for inviting us to present a keynote at Young OR and to the YOR audience for questions, comments, and appreciation. We thank the reviewers for helpful suggestions that improved the clarity of this paper.

References

Ackermann F and Belton V (1994). Managing corporate knowledge experiences with SODA and VISA. *Br J Mngt* **5**: 163–176.

Adelman L, Sticha PJ and Donnell ML (1986). An experimental investigation of the relative effectiveness of two techniques for stucturing multiattributed hierarchies. *Org Behav Hum Decis Process* **37**: 188–196.

Alhakami AS and Slovic P (1994). A psychological-study of the inverse relationship between perceived risk and perceived benefit. *Risk Anal* **14**(6): 1085–1096.

Bana e Costa C and Vansnick J-C (1995). General overview of the Macbeth approach. In: Pardalos PM, Siskos Y and Zopounidis C (eds). *Advances in Multicriteria Analysis*. Kluwer: Dordrecht.

Bana e Costa C and Vansnick J-C (1997). Applications of the MACBETH approach in the framework of an additive aggregation model. *J Multi-Crit Decis Anal* **6**: 107–114.

Bana e Costa CA, De Corte JM and Vansnick JC (2004). On the mathematical foundations of MACBETH. In: Figueira J, Greco S and Ehrgott M *et al*. (eds). *Multiple Criteria Decision Analysis: The State of the Art Surveys*. Kluwer: Dordrecht.

Bana e Costa CA, Ensslin L, Correa EC and Vansnick J-C (1999). Decision support systems in action: Integrated application in a multicriteria decision aid process. *Eur J Opl Res* **113**(2): 315–335.

Baron J (2000). *Thinking and Deciding*. Cambridge University Press: Cambridge.

Beach LR and Connolly T (2005). *The Psychology of Decision Making*. Sage: Thousand Oaks, CA.

Belton V (1990). Multiple criteria decision analysis: Practically the only way to choose. In: Hendry L and Eglese R (eds). *Operational Research Tutorial Papers*. Operational Research Society: Birmingham.p. 53–101.

Belton V and Stewart TJ (2002). *Multiple Criteria Decision Analysis: An Integrated Approach*. Kluwer: Boston.

Bettman JR, Johnson EJ, Luce MF and Payne JW (1993). Correlation, conflict, and choice. *J Exp Psychol-Learn Memory Cog* **19**(4): 931–951.

Bond S, Kurt AC and Keeney R (2007). Generating objectives: Can decision makers articulate what they want? *Mngt Sci* (in press).

Borcherding K and von Winterfeldt D (1988). The effect of varying value trees on multiattribute evaluations. *Acta Psychol* **68**: 153–170.

Buede DM (1986). Structuring value attributes. *Interfaces* **16**(2): 52–62.

Clemen RT (1996). *Making Hard Decisions: An Introduction to Decision Analysis*. Duxbury Press: Belmont, CA.

CoRWM (2006). Managing our radioactive waste safely: CoRWM's recommendations to Government, CoRWM, http://www.corwm.org.uk/pdf/fullreport.pdf.

El-Bassunie N (2006). *Testing the effectiveness of priming as an instrument to reduce the occurrence of scope insensitivity*. Unpublished MSc thesis, Department of Operational Research, London School of Economics and Political Science.

Fischer GW (1995). Range sensitivity of attribute weights in multiattribute value models. *Org Behav Hum Decis Process* **62**(3): 252–266.

Fishburn PC (1970). *Utility Theory for Decision Making*. Wiley: Chichester.

Fox CR, Bardolet D and Lieb D (2005). Partition dependence in decision analysis, resource allocation and consumer choice. In: Zwick R and Rapoport A (eds). *Experimental Business Research*, Volume III. Kluwer: Doredrecht.

Frederick S and Fischoff B (1998). Scope (in)sensitivity in elicited valuations. *Risk Decis Pol* **3**(2): 109–123.

French S, Simpson L, Atherton E, Belton V, Dawes R, Edwards W, Hämäläinen RP, Larichev O, Lootsma F, Pearman A and Vlek C (1998). Problem formulation for multi-criteria decision analysis: Report of a workshop. *J Multi-Crit Decis Anal* **7**: 242–262.

Goldstein WM and Beattie J (1991). Judgements of relative importance in decision making: importance of interpretation and the interpretation of importance. In: Brown DR and Smith JEK (eds). *Frontiers in Mathematical Psychology*. Springer: New York and Berlin.

Goldstein WM and Mitzel HC (1992). The relative importance of relative importance: Inferring other people's preferences from relative importance ratings and previous decisions. *Org Behav Hum Decis Proces* **51**: 382–415.

Goldstein WM, Barlas S and Beattie J (2001). Talk about tradeoffs: judgments of relative importance and contingent decision behavior. In: Goldstein WM, Barlas S and Beattie J (eds). *Conflicts and Tradeoffs in Decision Making: Essays in Honor of Jane Beattie*. Cambridge University Press: Cambridge.

Goodwin P and Wright G (2004). *Decision Analysis for Management Judgement*, 3rd edn. Wiley: Chichester.

Hastie R and Dawes RM (2001). *Rational Choice in an Uncertain World*. Sage: Thousand Oaks.

Hogarth RM (1987). *Judgment and Choice: The Psychology of Decision*. Wiley: Chichester.

Hsee CK and Rottenstreich Y (2004). Music, pandas and muggers: On the affective psychology of value. *J Exp Psychol General* **133**: 23–30.

Humphreys P and McFadden W (1980). Experiences with MAUD: Aiding decision structuring versus bootstrapping the decision maker. *Acta Psychol* **45**: 51–69.

Jones A, Walker G and Young C (2005). Tailoring decision conferencing to suit a Department of Trade and Industry application. *Presented at OR47*, Chester (unpublished). Operational Research Society: Birmingham, UK.

Kahneman D and Knetsch JL (1992). Valuing public goods: the purchase of moral satisfaction. *J Environ Econ Mngt* **27**: 57–70.

Kahneman D and Tversky A (1979). Prospect theory: An analysis of decision under risk. *Econometrica* **47**(2): 263–292.

Kahneman D, Knetsch JL and Thaler RH (1991). Anomalies: The endowment effect, loss aversion, and the status quo bias. *J Econ Persp* **5**(1): 193–206.

Keeney RL (1992). *Value focused thinking: A Path to Creative Decision Making*. Harvard University Press: Cambridge.

Keeney RL and Raiffa H (1976). *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. Wiley: Chichester.

Kleinmunz DN (2007). Resource allocation decisions. In: Edwards W, Miles RF and von Winterfeldt D (eds). *Advances in Decision Analysis*. Cambridge University Press: Cambridge.

Lipkus IM, Samsa G and Rimer BK (2001). General performance on a numeracy scale among highly educated samples. *Med Decis Making* **21**: 37–44.

Luce MF, Payne JW and Bettman JR (1999). Emotional trade-off difficulty and choice. *J Market Res* **36**(2): 143–159.

Montibeller G and Belton V (2006). Causal maps and the evaluation of decision options: A review. *J Opl Res Soc* **57**: 779–791.

Morgan M (2006). *Testing scope invariance in MCDA resource allocation*. Unpublished MSc thesis, Department of Operational Research, London School of Economics and Political Science.

Newell BR, Lagnado DA and Shanks DR (2007). *Straight Choices: The Psychology of Decision Making*. Psychology Press: Hove.

Payne JW, Bettman JR and Johnson EJ (1988). Adaptive strategy selection in decision making. *J Exp Psychol Learn, Memory Cog* **14**: 534–552.

Payne JW, Bettman JR and Johnson EJ (1993). *The Adaptive Decision Maker*. Cambridge University Press: Cambridge.

Payne JW, Bettman JR and Schkade DA (1999). Measuring constructed preferences: towards a building code. *J Risk Uncertain* **19**(1–3): 243–270.

Phillips L (2006). When Values Conflict: How citizens, stakeholders and experts contributed to formulating policy for managing the UK's radioactive waste. *Operational Research Group Public Lecture*, London School of Economics (unpublished).

Phillips LD and Bana e Costa CA (2007). Transparent prioritisation, budgeting and resource allocation with multi-criteria decision analysis and decision conferencing. *Ann Opns Res* **154**: 51–68.

Pidgeon NF, Poortinga W, Rowe G, Horlick-Jones T, Walls J and O'Riordan T (2005). Using surveys in public participation processes for risk decision making: The case of the 2003 British GM Nation? public debate. *Risk Anal* **25**(2): 467–480.

Pöyhönen M and Hämäläinen R (1998). Notes on the weighting biases in value trees. *J Behav Decis Making* **11**: 139–150.

Pöyhönen M, Vrolijk H and Hämäläinen R (2001). Behavioral and procedural consequences of structural variation in value trees. *Eur J Opl Res* **134**: 216–227.

Rios Insua D and French S (1991). A framework for sensitivity analysis in discrete multi-objective decision-making. *Eur J Opl Res* **54**: 176–190.

Salo AA and Hämäläinen RP (1995). Preference programming through approximate ratio comparisons. *Eur J Opl Res* **82**: 458–475.

Schkade DA and Payne JW (1994). How people respond to contingent valuation questions: A verbal protocol analysis of willingness to pay for an environmental regulation. *J Environ Econ Mngt* **26**(1): 88–109.

Slovic P (1995). The construction of preference. *Am Psychol* **50**: 364–371.

Slovic P, Finucane ML, Peters E and MacGregor DG (2007). The affect heuristic. *Eur J Opl Res* **177**(3): 1333–1352.

Tversky A and Kahneman D (1991). Loss aversion in riskless choice: A reference-dependent model. *Quart J Econ* **106**(4): 1039–1061.

von Nitsch R and Weber M (1993). The effect of attribute ranges on weights in multiattribute utility measurement. *Mngt Sci* **39**(8): 937–943.

von Winterfeldt D (1999). On the relevance of behavioral research for decision analysis. In: Shanteau J, Mellers B and Schum D (eds). *Decision Science and Technology*. Kluwer: Boston.

von Winterfeldt D and Edwards W (1986). *Decision Analysis and Behavioral Research*. Cambridge University Press: Cambridge.

Weber M and Borcherding K (1993). Behavioral influences on weight judgments in multiattribute decision making. *Eur J Opl Res* **67**: 1–12.

Weber M, Eisenfuhr F and von Winterfeldt D (1988). The effects of splitting attributes on weights in multiattribute utility measurement. *Mngt Sci* **34**(4): 431–445.

Weinstein M and Zeckhauser R (1973). Critical ratios and efficient allocation. *J Pub Econ* **2**: 147–157.

Received June 2007;
accepted October 2007 after one revision