Multi-criteria decision analysis

Michael Wood
14 October 2009

Many decisions involve multiple criteria that may conflict, so the decision maker has the task of balancing these criteria. For example, you may be faced with deciding between a job which pays well but which may be lonely and stressful, and another job which does not pay as well but is likely to be less stressful and comes with a better social life. On a larger scale, a decision about building a new airport would involve balancing criteria such as reducing journey times for travellers, reducing noise levels for the local population and reducing carbon emissions. What is the best way to proceed with decisions like these?

This problem arises in many contexts – for example:

- Deciding on a suitable location for an office (Goodwin and Wright, 1998, chapter 2),
- Deciding whether to go ahead with a project like launching a new product,
- Choosing accommodation options in higher education (Bowhill, 1994),
- Evaluating alternative energy technologies (Clemen and Reilly, 2001, chapter 15),
- Choosing a car (Clemen and Reilly, 2001, chapter 15),
- Deciding how to invest money taking account of risk, return and ethical considerations,
- Decisions about buying something (either by an individual or an organisation),
- Choosing a university

I will use a personal decision – my decision about which mode of transport to use to get to work – as an example to illustrate the method described below, but the problem, and approaches to tackling it, are of far wider concern.

Many approaches to multi-criteria decision problems have been suggested, among them SMARTS (Edwards and Baron, 1994; also described in Goodwin and Wright, 1998, Clemen, 1996, and Clemen and Reilly, 2001, chapter 5 – the last two books do not use the acronym SMARTS but the basic concepts are all there), AHP (the analytic hierarchy process – originally due to T. Saaty and described in Goodwin and Wright, 1998, chapter 15) and Even swaps (Hammond et al, 1998 and 1999)[1].

The method I suggest here involves the steps listed below.

Don't forget that the aim of your model is to help the decision maker make a decision. As well as the objective information about each alternative option, this will also depend on the values and preferences of the decision maker[2]. *This means that different decision makers are likely to come to different conclusions.*

You should treat multi-criteria decision models with caution! The information on which they are based is often imprecise, uncertain or wrong, and the assumptions on which they depend are often very shaky. For each step of the method, ask yourself if it makes sense. If the answer is no, either adjust the method, or stop and go for something simpler. Do what seems right for your decision - but you must be prepared to explain why. Part of the purpose of decision modelling is to justify your decision to others (and yourself).

Preliminary steps

Check that you know who or what the decision maker (DM) is. It may be a particular individual, or any individual finding himself or herself in a particular position. Alternatively it may be an

organisation, or a committee or individual or group within an organisation. There are obvious problems with decisions made by groups – it is then particularly important to have a structured method of analysis to clarify the reasons for the decisions made.

You should also check that the decision problem is defined in a sensible way.

I am obviously the DM for my decision about transport to work (see the spreadsheet at http://userweb.port.ac.uk/~woodm/dra/TransportDecisionDra.xls). I am modelling *my* decision; the analysis for someone else would probably be very different.

Step 1: Clarify alternatives (options)

Decisions involve a choice between *alternatives* (otherwise known as *options*, or *courses of action* or *strategies*). The first step is to list these alternatives. Sometimes this is obvious (to bid for a contract or not to bid), but sometimes it isn't. There may be a large number of alternatives in a decision about where to locate a business, or about a career for an individual. Obviously, it is a good idea to do some research to try get as many promising alternatives as possible in the list, perhaps by using some creative thinking techniques (see below)[3].

Step 2: Clarify criteria (for evaluating alternatives)

This is an important step. It is often helpful to use a tree structure so that you end up with a *critieria tree (often called a value tree).* In the Transport decision the top level criteria are costs, journey quality, timing issues, health costs/benefits and environmental. These are then broken down into finer detail – for example the journey quality is broken down into five sub-criteria.

To produce this tree, or a simple list of criteria, you need to find out what criteria are relevant to the decision. There are various ways of doing this. For example:

- Develop a wish list. What does the DM want? What does the DM value?
- Take the alternatives in pairs. Ask the DM what the pros and cons of each are. These pros and cons should then suggest the criteria that are important.

This process is often difficult and deserves care. You must check that the tree (or list of criteria) is complete (nothing left out) and without overlap (otherwise you may double count some things). You should also check that it makes sense to assess each criterion separately, and that your list or tree is as simple as possible.

If you have several monetary objectives, try to use a measure like cost per year to combine them into a single figure (as I have done in the Transport decision) This is by far the best way of combining monetary objectives.

For a more detailed discussion of how to structure criteria see Belton and Stewart (2002, chapter 3), Clemen and Reilly (2001, chapter 3) and Keeney (1992). Criteria are sometimes described using words like *objectives, attributes, goals* – you will need to bear this in mind when reading these references.

There are likely to be different ways of structuring criteria leading to different trees. You need to choose a tree that is helpful for analysing your decision.

Step 3: Produce a table of consequences

This should be a table with criteria listed down the left hand side and alternatives across the top (or the other way round). You then write the consequences of doing each alternative from the point of view of each criterion in the list in the appropriate cell. Some of the entries may just be notes, others may be yes/no/don't know, and others may be a numerical scale (which should obviously be explained). If you have a criteria tree, focus on the criteria at the bottom of the tree

(for example, in the Transport decision there are entries for the subcriteria under journey quality, but no entries for journey quality itself).

Step 4: Recheck the options and criteria

Return to the options and criteria and see if you can think of any more. Creative thinking techniques may help here. Return to this process after each step below.

For example, in the Transport decision, the relatively high cost of going by car may suggest the idea of something like a car, but cheaper – such as a motor cycle or a car share scheme.

Step 5: Rank each option against each criterion

This will produce a table of ranks based on the consequences table. The ranks express the DM's preferences and show which options are preferred from the point of view of each criterion. For example, in my Transport decision under costs, the train is ranked 1 because it is the best (cheapest) and the taxis option is ranked 5 bcause it is the worst of the five alternative modes of transport from this point of view.

If an option is unacceptable against a criterion – regardless of the performance of the option against the other criteria – then this is indicated with a cross (X) instead of a rank. This applies to the duration of the walk option – which I judged to be unacceptable regardless of performance on the other criteria.

If your table of consequences is complicated producing this table of ranks should help with the next step. In other cases you may feel it's not necessary – in which case leave it out.

Step 6: Eliminate any dominated or unacceptable alternatives, and any criteria which rank the remaining options equally

The rankings table can be used to eliminate alternatives that are *dominated* by another alternative, or that fail to reach a *minimum acceptable level* on one or more of the objectives. An alternative would be dominated by another alternative if one alternative is as good as, or better than, the other, on *all* criteria. This would mean that there is no point in considering the dominated alternative.

In the Transport decision, walking is eliminated from further consideration because of the duration. None of the alternatives is dominated, although taking taxis is close to being dominated by the car option. Going by car is better than taking taxis in terms of cost, duration and flexibility, and the two alternatives are the same on all the other criteria *except* environmental damage – taking taxis is better according to this criterion. Having eliminated Walking, all remaining options are ranked equal first under Enjoyment – so this criterion is of no help in making the decision and so can be deleted.

Assuming that this leaves more than one option so that you still have a decision to make, there will be no option which is best on all criteria (otherwise this option would dominate the others), so you need to make a *trade-off* between the different criteria. The option you eventually choose will be worse on some of the criteria than other options. You now need a method of deciding whether a better performance on some criteria *compensates* for a worse performance on others. There are a number of ways in which this can be done. The simplest is to ask the DM for an intuitive decision based on the table of consequences and the table of ranks. Sometimes simply presenting the information in these formats may be sufficient help. More typically, however, we use a formal method for analysing the trade-offs and deciding on the best option. There are a

number of such formal methods – the one described in Steps 7 – 9 below is known as SMARTS, or using an additive value function (or a weighted sum or a weighted average value function).

The basic idea of SMARTS is to assess the value of the consequences of each option from the point of view of each of the criteria, and then use a weighted average of these values as a measure of the overall value of each option – where the weight represent the importance of each criterion to the DM.

Step 7: Convert consequences to value scores

The next stage is to measure the *value* (sometimes called *utility*) of each consequence from the point of view of each criterion. To make it easier to combine these values you need a similar scale for each criterion.

One approach is to use 0 for the worst performing alternative on each criterion, and 100 (or 10 or 1) for the best. Then 50 would represent a consequence halfway between the best and the worst.

Alternatively, you could define clearly the best and worst performances on each criterion and score these as (say) 100 and 0. This is the approach taken in the exercise on the houses below. The important thing is that best and worst are clearly defined.

Having defined the best and the worst, there are two ways you can derive the intermediate scores. If you have a measurable criterion (eg the area of an office) you can you can use this to define value scores. Normally we would assume a straight line (proportional) relationship between the criterion and the value.

If there is no measurable attribute, you will need to get the DM to rate each alternative directly.

In the Transport decision, I have assumed a straight line relationship between duration and its value score. The best duration is 20 minutes which scores 10. The worst is 50 minutes which scores 0. (Remember that Walking has been eliminated.) Taking taxis takes 30 minutes which is closer to 20 minutes than 50 minutes. To be more precise, the difference between 20 minutes and 30 minutes is 10 minutes, and remembering that the difference between 20 minutes and 50 minutes (30 minutes) is worth 10 value points, 30 minutes is worth 10 – (10/30)*10 or 6.7. This may be clearer if you sketch a graph of cost against value scores. A similar method is used for cost.

The other criteria are not numerical, so here I (as DM) have decided that "moderate" under Ability to carry luggage is worth 4 value points.

It is very important to distinguish between the criterion (e.g. duration) and the value the DM places on it. The assessment of value is, of course, subjective.

Step 8: Assess weights for value scores

Obviously some criteria will be more important than others. The weights represent the relative importance of each criterion *taking account of the value scoring system*. It is vital to consult the decision maker here. You are modelling preferences so everybody is likely to be different. For individual decision making this is not a problem, but for business decision making it may be if different stakeholders express different preferences.

It is important to ensure that the weights accurately reflect the DM's values. A good way of achieving this is to use the *swing* weighting procedure (see Goodwin and Wright, 1998; Clemen and Reilly, 2001; Edwards and Baron, 1994). The idea of this is to imagine an alternative which scores 0 (the worst score) on each criterion, and then imagine a swing to the best score for each criterion in turn. The DM is asked to compare these swings: the best is given

a provisional weight of 100, and so on. *It is important to be very careful that the DM has understood exactly what he/she is being asked to do.*

It is then convenient to normalise the weights so that they add up to 1 or 100%. This can be achieved by dividing each wieght by the total of the weights.

The problem with not using the swing weighting procedure, and simply asking the DM about the importance each criterion, is that you are likely to ignor the size of the swing on each criterion. If, for example there were only two options for the Transport decision – Take train and Ride bike – the cost difference would only be 2p and of very little importance. If I just asked the DM how important Cost is, I might forget to take into account the very small cost swing involved.

Step 9: Work out aggregate (total) weighted scores for each alternative, or plot costs against benefits

Either
Multiply value scores by weights and add up the total for each alternative.
Or
Multiply value scores by weights for all criteria except the cost, and then plot a graph of costs against aggregate value. This gives the DM the opportunity to see how costs relate to the other criteria and make a decision by inspoecting the graph. This may be useful if the DM finds it difficult to compare the importance of the cost criterion with the others. In this case of course, no weights or values are needed for the cost criterion.

You can use a graph like this to see which alternatives are dominated and so can be ignored. The alternatives which are not dominated, and so are worth considering, are known as the *efficient frontier* (see Goodwin and Wright, 1998, p. 33). In the Transport decision the efficient frontier comprises just one alternative – Take train – since this is better *and* cheaper than all the other alternatives.

Step 10: Check assumptions

This approach makes a number of important assumptions:

- The criteria tree is complete (nothing important left out) and without overlaps
- The list of options includes all viable possibilities
- The scores and weights can be meaningfully assessed by the DM. *Take particular care over the weights, and ensure that all value scores are on a similar scale (eg 0-10).*
- If value scores and weights are used, the DM must be prepared to let a high score on one criterion (or objective) compensate for a low score on another. This is known as a "trade-off", and the weights measure how much of an increase on one criterion is required to compensate for a given decrease on another.
- The value scores can be assessed independently of each other (described as "mutual preference independence" in Goodwin and Wright, 1998, and "judgmental independence" in Belton and Stewart, 2001). If the value scores for one criterion *depend on* another criterion this will not be the case. (Note that this is different from the issue of whether the variables are correlated or dependent on each other in a statistical sense.)

In the Transport decision, I was certainly not sure about the value scores, and the weights. On another day, I may have come to different conclusions! Also there is some doubt in my mind

about the independence of my assessments of the value of Duration and Ability to work on the journey[4].

If these assumptions are not reasonable, your conclusions are unlikely to be completely reliable. It may be possible to do better with another criterion tree. Or the analysis may have to stop with the table of consequences – which could still be useful for structuring the decision.

Step 11: Present the results to the DM and check they are sensible

It may help to use diagrams. Useful diagrams include:

1. Graph (scatter plot) of costs against benefits.
2. Option profiles (or performance profiles) to show how each alternative scores against each objective. The option profile in the Transport decision spreadsheet is perhaps a little too complicated – it might be more useful if the values were aggregated at the higher level (Cost, Journey quality, etc).

It is also a good idea to do a *sensitivity analysis*. Which scores or weights are least certain? What happens if you change them? Try it on the spreadsheet and see. Conclusions are often unchanged by large changes in scores or weights: demonstrating this will strengthen your conclusions.

It is easy to do a sensitivity analysis to look at the change in one or two inputs, but analysing the effect of changes in several inputs simultaneously is not easy[5].

Don't forget that different decision makes may have different values and criteria – their weights and possibly the value scores may differ – and so may come to different decisions.

Testing the model

If you want to test—for example—the distance recorder in a car you would compare the reading from this with another way of getting the same information—for example what it says on signposts. The problem with testing decision analysis models is that there is often no satisfactory alternative method of deciding which is the best option to choose. In many situations all you can do is

1. Check that the assumptions are reasonable (e.g. that there is no overlap between criteria, scores and weights based on credible information, etc), and
2. Show the analysis and the results to the DM and ask if he/she thinks they are helpful and/or right. (The difficulty is that your analysis may be so convincing that they are convinced by it even if it is wrong—but then if they think it's right then it must be right? Mustn't it?)
3. And …? Can you think of any other ways of testing decision models?

Techniques for enhancing creativity

(I will not be covering this in detail, but it may be relevant to your assignment.)
Creativity may be an issue when listing options or evaluation criteria (Steps 1 and 2 above), and in trying to foresee unintended consequences of potential options. Many techniques have been suggested - for example:

Brainstorming

In a group of 6-12 people ... each person produces an idea in rotation ... ideas put up on a big board ... *all* ideas accepted ... none criticized in the initial stages ... only later are those obviously not viable discarded.

Lateral thinking

Edward de Bono (De Bono, 1977) has put forward many ideas for encouraging new ways of looking at problems and situations. One such is to deliberately consider an apparently ridiculous idea: eg that cars should have square wheels.

Variety of inputs

It is also often a good idea to ensure that a variety of people from different backgrounds are involved; in particular outsiders who may bring a different perspective to bear on the situation may be useful.
See also Belton and Stewart (2002, chapter 3).

Exercises

1 Construct a criteria tree, and a list of possible courses of action, for a decision you will have to make soon. (For example, you could use a decision about a holiday or the purchase of a car.)

Are the criteria tree and the list of possible courses of action complete, or have you missed some things out? Is there any overlap? Is it helpful to structure it as a tree? If so, why?

2 Produce a consequences table for your decision in Q1. How would you assess the values on your value tree? Are they likely to be judgmentally independent? If they aren't it might be a good idea to change your value tree.

3 The CEO of a large automobile company is concerned about the design of a new version of a well-known model because there is a risk of fuel-led fire. The plan is to produce 4 000 000 cars. He has the following information:

- The probability of explosion is 1/100 000.
- Maximum compensation of 200 000 euros per death
- The cost of re-design is 9 euros per car

What criteria do you think are relevant to this decision? What do you think he should do?
(From a presentation by Marc Menestrel who got it from Mark Ashbar: "The Corporation")

4 What criteria do you think are relevant to setting speed limits on motorways? What are the trade-offs and how would you recommend they should be made?

5 An organisation is considering purchasing a new project management software package. They decide that it is important that the chosen package should have an extensive range of facilities available, and should be easy to use. They draw up a short list of four packages and rate each of them on these two criteria:

Package	*Range of facilities*	*Ease of use*
Package 1	62	100
Package 2	70	50
Package 3	60	60
Package 4	90	0

(This table is at http://userweb.port.ac.uk/~woodm/dra/package.xls .)

They also decide that the scores on the "range of facilities" should have a weighting of 2 and the "ease of use" scores a weighting of 1, because the former is "twice as important" as the latter.

What conclusions would you come to?

6. Choosing a house

This is the example- slightly edited - in the manual for the package, AIM (V Lotfi & S Zionts, 1994), reproduced here with permission from one of the authors. A more recent version is available from http://mgt2.buffalo.edu/webaim.

Consider a situation in which a house buyer wishes to evaluate potential houses for purchase on five criteria as follows:

1- The first criterion is the number of bedrooms and a house with more bedrooms is better than a house with fewer. However, an acceptable house *must* have at least two bedrooms and no more than eight bedrooms. All other things considered equal, houses with four or more bedrooms are the same.

2- The second criterion is the size of the garden in acres and you prefer larger gardens. However, an acceptable house *must* have at least 0.4 acres of land (an acre is about 4000 square metres). All other things considered equal, houses with two or more acres are the same.

3- The third criterion is the age of the house in years and you prefer newer houses. A house, that you consider acceptable, *must* be no more than 35 years old. All other things considered equal, houses which are 10 years or younger are the same.

4- The fourth criterion is the asking price in $000 and you prefer less expensive houses to more expensive houses. An acceptable house *must* not be more than $100,000.

Suppose that the buyer has examined the daily newspaper and has compiled a list of 100 houses. The first 10 are shown below (and at http://userweb.port.ac.uk/~woodm/dra/houses.xls). Note that while compiling the list, the buyer has decided to include houses that appeared promising but had values outside the *must* levels on one or more criteria. This is a "play safe" strategy and allows such houses to become feasible candidates if the *must* levels should be changed.

(a) Make a table of value scores. A score of 100 should indicate the best on each criterion, 0 should indicate the worst, and X should indicate a value that is unacceptable. You should assume that the graph of value against each attribute is a straight line. (Is this reasonable?[6].)

(b) Are any of the alternatives unacceptable? Are any of them *dominated* by another alternative?

(c) Would all buyers agree on the value scores? (This example is American where new houses tend be more highly valued than old ones. In England, the opposite is often the case.) Imagine it was *you* buying the house. What would your value scores be?

(d) Decide what the appropriate weights would be if *you* were buying the house (you should use the swing weighting procedure for this). Use these weights to work out

the best alternative among the 10 below. Now compare your answers with someone else. Who is right? (This is a silly question – but you should understand why!)

(e) What practical use do you think this model is? If you were buying a house, would you buy the top scorer on the model, or would you want to ask further questions?

LIST OF 10 AVAILABLE HOUSES				
	Bedrooms	**Acres**	**Age**	**Price**
HOUS1	5	0.25	48	290
HOUS2	5	0.4	22	90
HOUS3	3	0.6	25	92
HOUS4	2	0.3	45	42
HOUS5	2	0.25	16	47.5
HOUS6	2	0.2	34	87.5
HOUS7	4	0.6	12	95
HOUS8	7	1.33	40	180
HOUS9	3	0.3	45	55
HOUS10	3	0.4	30	100

7. An alternative way of analysing the house decision in Question 3 would be to produce a scatter diagram of benefits (aggregate value excluding costs) against costs. Try this and find the efficient frontier. Which method of analysis do you think is most useful?

The weights used to assess the importance of the three benefits criteria (bedrooms, acres, age) may be unreliable so a sensitivity analysis would be useful. Try increasing and decreasing each of your weights in turn. Does this make any difference to the best buy?

8. A British company has won an important contract to supply components regularly to Poland. Four methods of transport are being considered: (i) air, (ii) sea, (iii) road and ferry and (iv) rail and ferry. The company's distribution manager has identified four relevant attributes for the decision: Punctuality, Safety of Cargo, Convenience and Costs. She has also allocated weights of 30 to punctuality, 60 to safety of cargo and 10 to convenience.

The manager then rated the performance of each form of transport on the different attributes. The values she assigned are shown below together with the estimated annual cost of using each form of transport (and at

http://userweb.port.ac.uk/~woodm/dra/CompanyTransport.xls). (On all three value scales higher numbers are better than lower ones.)

Form of transport	***Benefits***			***Costs ($)***
	Punctuality	*Safety*	*Convenience*	
Air	100	650	60	150000
Sea	0	600	80	90000
Road and Ferry	60	300	100	40000
Rail and Ferry	70	800	0	70000

(a) The value scores for Safety range from 300 to 800, whereas the other scales range from 0 to 100. Explain why this is a problem.

(b) Adjust the Safety scale so that it ranges from 0 to 100 like the others.

(c) Determine the form of transport which has the highest valued overall benefits (excluding the costs).

(d) For each form of transport, plot the value of overall benefits against costs, and hence identify the forms of transport which lie on the efficient frontier.

(e) Discuss how she should choose between the options on the efficient frontier.

(f) Now supposed you have been asked to carry out a critical evaluation of this analysis. What would you check to ensure that the analysis is useful?

(This is an edited version of Question 4 in Goodwin and Wright, 1998, pp. 47-8.)

9 Suppose you are advising someone who wants to study for a business degree at a British university about which university they should apply for. What criteria do you think might be relevant? How would you find out?

Now look at the Times Good University Guide ranking at http://tinyurl.com/4dudl8 . How helpful is this? How could you make it more useful?

Now look at the explanatory notes (click on Key to the tables). You will see they use a "z-score transformation". This means that the scales are expressed in terms of standard deviations above the mean (e.g. if the mean was 80 and the sd was 20, the z score corresponding to 50 would be – 1.5). Why do you think they do this?

10 A management course is assessed by means of five examinations - in management of organisations, finance, marketing, statistics and computing. Each examination is marked "out of 100", the marks for the five subjects are added up for each student, and the pass mark is deemed to be 200. (Ie a student with a total of 200 passes, and one with 198 fails.)

Comment critically on this procedure as a means of deciding who should pass. How would you improve it?

Reading

There is a good account of this approach to multi-criteria decision analysis in Goodwin and Wright (1998, chapter 2), and a rather mode advanced account in Clemen and Reilly (2001, chapter 15 and 16), or in Clemen (1996). Baron (2008, pp. 332-3 and Chapter 14) covers similar ground. Hammond et al (1998 and 1999) give a simple description of a slightly different approach from a very practical perspective. There is more detail on the multicriteria problem in the books by Belton and Stewart (2002), Bouyssou et al (2000 and 2006) and Keeney (1992). There are also a number of websites – e.g. Hamalainen (n.d.).

References

Baron, J. (2008). Thinking and deciding. Cambridge University Pess. (in the library as an e-book)

Belton, V., & Stewart, T. J. (2002). Multiple criteria decision analysis: an integrated approach. Boston: Kluwer.

Bouyssou, D., Marchant, T., Pirlot, M., Perny, P., Tsoukias, A., & Vincke, P. (2000). Evaluation and decision: models a critical perspective. : Kluwer Academic Publishers.
(Especially chapter 6. Wide ranging content. Mix of good, intuitive explanations and abstract mathematics.)

Bouyssou, D ... [et al.] (2006). Evaluation and decision models with multiple criteria : stepping stones for the analyst. New York: Springer. (Rather more technical than the last one.)

Bowhill, B. (1994). Evaluating academic accomodation options in higher education. Property Management, 12(4), 4-11.

Clemen, R. T. (1996). Making hard decisions (2nd ed). Belmont, Calif: Duxbury Press. (Similar to the next reference except it does not include @risk.)

Clemen, R. T. & Reilly, T (2001). Making hard decisions with DecisionTools Suite Belmont, Calif: Duxbury Press. (This includes a disc with @risk and some other palisade software.)

De Bono, Edward. (1977). Lateral thinking: a textbook of creativity. Harmondsworth: Penguin. (On creativity, obviously. (There is an earlier edition of this book, and several others on a similar theme in the library.)

Edwards, W., & Barron, F. H. (1994). SMARTs and SMARTER: improved simple methods for multiattribute utility measurement. Organizational Behavior and Human Decision Processes, 60, 306- 325.

Goodwin, P., & Wright, G. (1998). Decision Analysis for Management Judgment (2nd edition). Chichester: Wiley.

Hammond, J. S., Keeney, R. L., & Raiffa, H. (1998, March). Even swaps: a rational method for making trade-offs. Harvard Business Review, 137-150.
(Very similar to a chapter in their book Smart Choices.)

Hammond, J. S., Keeney, R. L., & Raiffa, H. (1999). Smart choices. Boston: Harvard Business School Press.
(Relatively non-academic book aimed at practical decision makers.)

Hamalainen, Raimo P. (Project Director). (n.d.). *Multiple criteria decision analysis elearning site.* Retrieved October, 2005, from http://www.mcda.hut.fi/ . (There is also some software here, I think.)

Keeney, R. L. (1992). Value-focused thinking: a path to creative decisionmaking. Cambridge, Massachusetts: Harvard University Press.

Endnotes: more subtle points which can be ignored

[1] And many more. These methods are suggestions about how decisions should be analysed in the same way that statistical methods are suggestions about how statistical problems should be analysed. There is, however, an important difference between methods in these two areas. Statistical concepts and methods tend to be relatively complex and standardised. Concepts like the standard deviation and methods like analysis of variance depend on an extensive mathematical background and are widely used. They are difficult for users to adjust as they see fit because they may not appreciate the full mathematical background, and if they use non-standard concepts they are likely not to be understood by other workers, clients and colleagues. The situation is very different with decision analysis methods. The mathematical background is often (but not always – the

AHP is an exception) very simple, and the methods are not widely used. This means you should feel free to adapt methods to your particular purposes – provided you explain why.

[2] It is easy to assume that decision models simply *describe* these preferences. However, it is possible people may not have clear preferences and that the effect of modelling non-existent preferences may be to *build* preferences for the decision maker. In effect, the analyst may be telling the decision maker what he or she wants, not asking.

[3] If your decision is an important one with long term consequences – like choosing a career – it may be a good idea to include options like "spend a year trying several possibilities".

[4] The problem is that if I can work on the journey, I think I may prefer a *longer* duration. So the way I assess duration may depend on my assessment of Ability to to work on the journey.

[5] One approach would be to use Monte Carlo simulation to investigate the effects of changes in the inputs. The simplest way to do this would be to decide on maximum and minimum values for each value score and weight, and then use uniform or triangular distributions to run a simulation.

[6] This may not be reasonable. The DM may, for example, decide that a 3rd bedroom is more important than a 4th and any additional bedrooms. They might assess the value of 3 bedrooms as half-way between 2 and 8 bedrooms, so they might score the value of 3 bedrooms as 50 (see Goodwin and Wright, 1998, chapter 2 for more detail).