2

Working with DELTA

This chapter is an introduction to a proposed human decision work process in which the DELTA method plays a central role. It is intended to serve only as an informal overview, introducing ideas and terminology enlarged on in Part II. The purpose is not to describe the mathematical or computational machinery necessary, but rather to give an intuitive feeling for how the method works and for its relevance to organisational decision-making. Another objective is to demonstrate that the suggested method is realistic to work with.

A feature of the method is that the decision-maker has to make his problem statements more visible than he would otherwise. This brings about a number of advantages. First, he must make the underlying information clear, and second, the statements can be the subject of discussions with (and criticism from) other participants in the decision process. Third, it can also be seen more clearly which information is required in order to “solve” the problem and within which areas some more information must be gathered before a well-founded decision can be made. Fourth, arguments for (and against) a specific selection can be derived from the analysis material. Fifth, the decision can be better documented, and the underlying information, as well as the reasoning leading up to a decision, can be traced afterwards. The decision can even be changed in a controlled way, should new information become available at a later stage.

Professional decision-makers in corporations as well as in public organisations today often use rather simple decision models to aid decisions. In many cases, decisions are made without employing any model at all. The decision might be based on rules of thumb or on intuition, or even be a repetition of a similar decision made earlier. Sometimes, decisions are made after listing the alternatives and discussing their consequences in an unstructured manner. These alternative–consequence lists may state the advantages and disadvantages of each course of action. When the special case of one action having all advantages and another all disadvantages does not prevail, it is often necessary to make a complicated comparison between the consequences of all alternatives. Other examples of well-known traditional decision aids include decision matrices and decision trees as discussed in Chapter 1. Many of them have the common disadvantage that they either do not handle probabilities at all, or else they require the decision-maker to make probability statements with precise numeric values, however unsure he is of his estimates.

Suppose a decision-maker wants to evaluate a specific decision situation. In order to solve the problem in a reasonable way, given available resources, a decision process such as the following could be employed, not necessarily in the exact order given.

- Clarify the problem, divide it into sub-problems if necessary
- Decide which information is a prerequisite for the decision
- Collect and compile the information
- Define possible courses of action
- For each alternative:
  - Identify possible consequences
  - For each consequence:
    - If possible estimate how probable it is
    - If possible estimate the value of it occurring
- Disregard obviously bad courses of action
- Based on the above, evaluate the remaining alternatives
- Carry out a sensitivity analysis
- Choose a “reasonable” alternative

The model described in the following should be seen in the context of such a decision process. The process is intuitively appealing, and numerous decision-makers unconsciously use a similar approach.

The Work Cycle

The decision process is carried out in a number of steps presented here in work-cycle form. A *work cycle* consists of six phases (Figure 2.1). The first step of the first cycle is special since there is much information to collect. The initial information is gathered from different sources. Then it is formulated in statements as indicated later in the chapter and entered into the DELTA Decision Tool (DDT, see Chapter 3).[1] Following that, an iterative process commences where step by step the decision-maker gains further insights and sometimes a conclusion. During this process, the decision-maker receives help in realising which information is missing, is too vague, or is too precise. He might also change the problem structure by adding or removing consequences or even entire alternatives as more decision information becomes available.

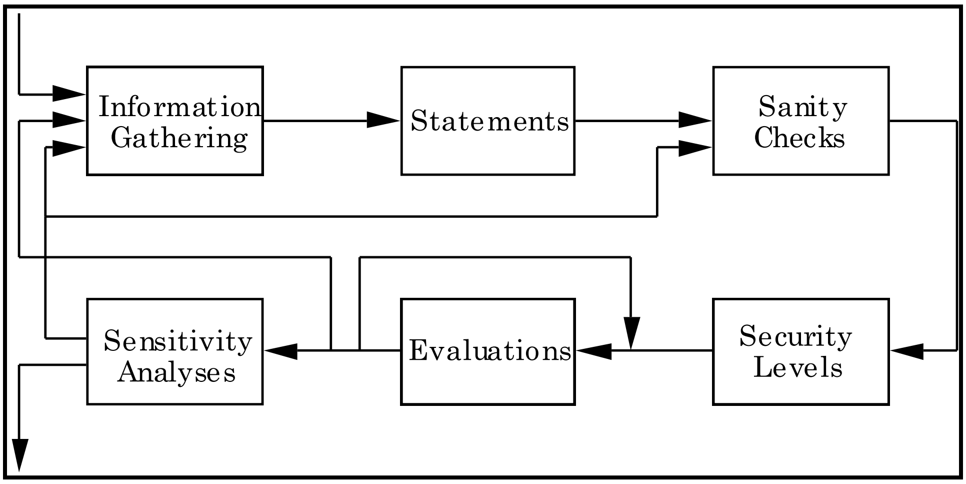

Figure 2.1 The DELTA work cycle

[1] The current version of DDT accepts numeric input by rulers, while future versions will accept linguistic input as well.

Information Gathering

In some cases, the first information collection phase can be a very long and tedious step. In larger investigations, it might take many man-years and result in documentation covering several meters of shelf space. In other cases, it might only require a few half-day discussions with experts and affected workers. It is impossible to describe any typical case because the situations are too diverse.

After the data collection phase, a filtering task commences where the decision-maker structures and orders the information. He tries to compile a smaller number of reasonable courses of action and identify the consequences belonging to each alternative. There is no requirement for the alternatives to have the same number of consequences. However, within any given alternative, it is required that the consequences are exclusive and exhaustive, i.e. whatever the result, it should be covered by the description of exactly one consequence. This is unproblematic, since a residual consequence can be added to take care of unspecified events.

Statements

Once the information is structured, it is entered into DDT in the form of statements such as *the probability of consequence C occurring is less than 40%*. For each new statement entered, the consistency of the information is checked.

The decision-maker's probability statements are represented by interval constraints and core intervals as further described in Chapter 4 on representation. Intervals are a natural form in which to express such imprecise statements. It is not required that the consequence sets are fixed from the outset. A new consequence may be added at a later stage, thus facilitating an incremental style of working. The collection of probability statements in a decision situation is called the *probability base*. Some elementary statements considered are the following.

- *The event H_1 is probable*
- *The event H_1 is possible*
- *The event H_1 is improbable*
- *The probability for event H_1 is a*
- *The probability for event H_1 is larger than a*
- *The probability for event H_1 is between a and b*
- *The event H_1 is as probable as H_2*
- *The event H_1 is more probable than H_2*
- *The event H_1 is much more probable than H_2*

A probability base is said to be *consistent* if it can be assigned at least one real number to each variable so that all inequalities are simultaneously satisfied.[2] The idea is that no meaningful operations can take place on a set of statements that have no variable assignments in common, since there is no way to take all the requirements into account. Note that the method deals with classes of functions of which there are infinitely many instantiations, and insists on at least one of them yielding consistent results.

Likewise, the values are expressed as interval statements. The translations of the value statements in a decision situation are called the *value base*. Some elementary statements considered in this thesis are the following.

- *The event H_1 is desirable*
- *The event H_1 is acceptable*
- *The event H_1 is undesirable*
- *The value of event H_1 is a*
- *The value of event H_1 is larger than a*
- *The value of event H_1 is between a and b*
- *The events H_1 and H_2 are as desirable*
- *The event H_1 is more desirable than H_2*
- *The event H_1 is much more desirable than H_2*

[2] For example $p(H_1) = 0.22$ and $p(H_2) = 0.39$.

Consistency is defined in the same way as for a probability base, and is also discussed in Chapter 4. The probability and value bases together with structural information constitute the *decision frame*.

When all statements in the current cycle have been entered, the data entry phase is over. As the insights into the decision problem accumulate during all the following phases, it is possible to add new information and alter or delete information already entered.

Sanity Checks

Thereafter, the work cycle goes into evaluating the alternatives. The first cycle begins by comparing the alternatives as they are entered. As the first evaluation step, the *sanity* of the decision frame is checked. Much information collected, especially in large investigations, runs the risk of being cluttered or misunderstood during the process. If some data in the frame is problematic, the decision-maker could consider leaving it out of the current cycle or recollecting it. Missing data is easily handled for later inclusion. For example, a missing consequence can be added at a later stage. If the set of consequences for some alternative is not exhaustive, a residual consequence can be temporarily added. Missing value constraints can be temporarily substituted with very wide intervals or just left out. Such possibilities have certain advantages as the results emerging at the outset of the evaluation may be viewed with greater confidence than if erroneous data is entered.

Security Levels

Many decisions are one-off decisions or are important enough not to allow a too undesirable outcome regardless of its having a very low probability. The common aggregate decision rules will not rule out an alternative with such a consequence provided it has a very low probability. If the probability for a very undesirable consequence is larger than some *security level*, it seems reasonable to require that the alternative should not be considered, regardless of whether the expected value shows it to be a good course of action. If the security level is violated by one or more consequences in an alternative and this persists beyond a predetermined rate of contraction (described below), then the alternative is *unsafe* and should be disregarded. An example of security levelling is an insurance company desiring not to enter into insurance agreements where the profitability is high but there is a very small but not negligible risk for the outcome to be a loss large enough to put the company's existence at stake. The security analysis requires some parameters to be set. This can often be done at an organisational level, and it will then have the effect of creating a policy within the organisation. Security levels is an important supplement to the expected value. It is more formally introduced in Chapter 5 and further discussed and exemplified in Appendix A.

Evaluations

After having taken security levels into account, which value does a particular decision have? In cases where the outcomes can be assigned monetary values, it seems natural that the value of the decision should be some kind of aggregation of the values of the individual consequences. One suggestion is to assign different weights to the consequences so that more probable ones are more influential than less probable ones. This line of reasoning leads to the expected monetary value (EMV), which is essentially the same construct as the general expected value discussed below. EMV shows the monetary result that would be obtained on average, should the decision situation reoccur a large number of times. Since not all decisions reoccur that often, some not at all, EMV should be interpreted as the average tendency prevailing in every decision situation.

There are a number of possible evaluation rules within DELTA, some of which are described in Chapter 5. Often, the final comparing rule of an evaluation in the DELTA method as well as in many other methods is the expected value (EV), sometimes instantiated as the expected utility or the expected monetary value. Since neither probabilities nor values are fixed numbers, the evaluation of the expected value yields quadratic (bilinear) objective functions of the form

$$EV(A_i) = p_{i1}v_{i1} + \dots + p_{in}v_{in}$$

where the p_{ik}'s and v_{ik}'s are variables. Maximisations of such expressions are computationally demanding problems to solve in the general case, using techniques from the area of quadratic programming [L89]. In Chapter 6 there are discussions about and proofs of the existence of computational procedures to reduce the problem to systems with linear objective functions, solvable with ordinary linear programming methods.

When a rule for calculating the EV for decision frames containing interval statements is established, the next question is how to compare the courses of action using this rule. It is not a trivial task, since usually the possible EVs of several alternatives overlap. The most favourable assignments of numbers to variables for each alternative usually render that alternative the preferred one. The first step towards a usable decision rule is to establish some concepts that tell when one alternative is preferable to another. For simplicity, only two alternatives are discussed, but the reasoning can easily be generalised to any number of alternatives.

> Alternative A_1 is *at least as good as* A_2 if $EV(A_1) \geq EV(A_2)$ for all consistent assignments of the probability and value variables.
>
> Alternative A_1 is *better than* A_2 if it is at least as good as A_2 and further $EV(A_1) > EV(A_2)$ for some consistent assignments of the probability and value variables.
>
> Alternative A_1 is *admissible* if no other alternative is better.[3]

If there is only one admissible alternative it is obviously the preferred choice. Usually, there are more than one since apparently good or bad

[3] This conforms to statistical decision theory [L59].

alternatives are normally dealt with on a manual basis long before decision tools are brought into use. All non-admissible alternatives are removed from the considered set and do not take further part in the evaluation. The existence of more than one admissible alternative means that for different consistent assignments of numbers to the probability and value variables, different courses of action are preferable. When this occurs, how is it possible to find out which alternative is to prefer?

Let $\delta_{12} = EV(A_1) - EV(A_2)$ be the differences in expected value between the alternatives. The *strength* of A_1 compared to A_2, given as a number $\max(\delta_{12}) \in [-1,1]$, shows how the most favourable consistent assignments of numbers to the probability and value variables lead to the greatest difference in the expected value between A_1 and A_2. In the same manner, A_2 is compared to A_1. These two strengths need not sum to one or to any other constant – the first might for example be 0.2 and the second 0.4. If there are more than two alternatives, pairwise comparisons are carried out between all of them.

Furthermore, there is a strong element of comparison inherent in a decision procedure. As the results are interesting only in comparison to other alternatives, it is reasonable to consider the differences in strength as well. Therefore, it makes sense to evaluate the *relative strength* of A_1 compared to A_2 in addition to the strengths themselves, since such strength values would be compared to some other strengths anyway in order to rank the alternatives. The relative strength between the two alternatives A_1 and A_2 is calculated using the formula

$$\text{mid}(\delta_{12}) = \frac{\max(\delta_{12}) + \min(\delta_{12})}{2} = \frac{\max(\delta_{12}) - \max(\delta_{21})}{2}$$

which is explained in detail in Chapter 5. The concept of strength is somewhat more complicated than discussed in this chapter. Alternative A_1 is said to strongly dominate alternative A_2 if $\min(\delta_{12}) > 0$, to markedly dominate if $\text{mid}(\delta_{12}) > 0$, and finally to weakly dominate if $\max(\delta_{12}) > 0$.[4] This is also explained in Chapter 5.

Only studying the differences in the expected value for the complete bases often gives too little information about the mutual strengths of the alternatives. Numbers close to any of the boundaries seem to be the least reliable ones when making the original imprecise statements. Hence, it would be advantageous to be able to study the strengths (or dominances) between the alternatives on sub-parts of the bases. If a dominance is evaluated on a sequence of ever smaller sub-bases, a good appreciation of the strength's dependency on boundary values can be obtained. This is denoted *contracting* the bases, and the amount of contraction is indicated as a percentage which can range from 0% to 100%. For a 100% contraction, the bases are contracted into single points, and the evaluation becomes the calculation of the ordinary expected value.[5]

The next chapter presents the DDT tool in some detail, complete with evaluation graphs. The results of the comparisons can be displayed either in a diagram for each pair of alternatives or as a summary for each alternative.

Sensitivity Analyses

After the evaluation, a *sensitivity analysis* is the next step. The analysis tries to show what parts of the given information are most critical for the obtained results and must therefore be given extra careful consideration. This is accomplished by varying a number of statements in desired ways, increasing or decreasing intervals, modifying structural information, etc. It also points to which information is too vague to be

[4] To be more precise, the DELTA method uses the concept of Δ-dominance as described in Chapter 5. It may colloquially be interpreted as the relative strength between the alternatives.

[5] The method uses the dual concepts of expansion and contraction as explained in Chapters 4 and 5, but the idea is the same as only contracting the bases. Since the core is not discussed in this chapter, neither is expansion.

of any assistance to the ongoing evaluation. Information identified in this way is subject to reconsideration, thereby triggering a new work cycle.

It is possible to regard the expansion and contraction procedures as automated kinds of sensitivity analysis. In order to maintain consistency, the expansion (contraction) increases (decreases) the bases in predefined ways. The decision-maker might, however, have other ideas of interesting modifications to make to the bases, like decreasing or even increasing selected intervals. He might have structural or problem specific information that leads him to manipulate certain intervals in special ways. A common strategy is decreasing intervals until only one alternative is admissible. This way further insights into the decision problem can be gained. It is simple to allow for this in the DELTA method and the procedures of expansion and contraction apply equally well to bases altered for reasons of sensitivity analysis.

Before a new cycle starts, alternatives found to be undesirable or obviously inferior by other information are removed from the decision process. Likewise, a new alternative can be added, should the information gathered indicate the need for it. Consequences in an alternative can be added or removed as necessary to reflect changes in the model. Often a number of cycles are necessary to produce an interesting and reliable result.

Decision Process Results

After the appropriate number of work cycles has been completed, both the decision problem and its proposed "solution(s)" in the form of preferred courses of action will be fairly well documented. Anyone interested and with access to the information can afterwards check, verify (and criticise) the decision based on the output documentation, which because all consequences are clearly presented shows how all the alternative courses of action have been valued. Also, during the decision process, the analysis is open for comments and can become the basis for further discussions. Another effect is that the decisions are less dependent on which employee handles a particular decision situation since deviations from corporate policy can be detected in the documentation after the process has been completed if not earlier.

This concludes the informal introduction to the DELTA method in a work process. The next chapter presents the DDT tool suitable for interactive use in a work cycle-based process. The chapters that follow in Part II go into considerably more detail in trying to present the representation and the evaluation procedures of the method.