Background Documents on
Risk Assessment in Engineering

Document #6

Utility, Preferences, and Risk Perception in Engineering Decision Making

JCSS
Joint Committee of Structural Safety

Date: November 2008

JCSS authors: **Marc A. Maes, University of Calgary and Michael H. Faber, Swiss Federal Institute of Technology**

Utility, Preferences, and Risk Perception in Engineering Decision Making

Marc A. Maes, and Michael. H. Faber

1 INTRODUCTION

Decision theory as described in Raiffa and Schlaifer (1961) and the axioms of utility as proposed by von Neumann and Morgenstern (1943) are generally accepted as the rational and operational framework for the support of engineering decision making. This fact is reflected by the adoption of decision theory as the basis for the probabilistic model code (PMC (2001)) by the Joint Committee on Structural Safety (JCSS). When we realize that "risk assessment" if conducted properly is yet another word for decision analysis, it becomes obvious that the implementation of this framework for the purpose of decision making in engineering application is almost universal.

In the practical application of decision theory it is often argued that risk perception needs to be taken into account in the formulation of the utility function. The reason for this being that the utility function must be able to represent the "real" behavior of the decision maker in regard to her/his preferences in a given situation. The amount of research invested into the experimental investigation and the mathematical modeling of the behavior of human decision makers is vast, see e.g. Pratt (1964), Arrow (1971) and Kahneman and Tversky (1979). It clearly points to the basic characteristics of the perception of risks under different conditions, it also points to a whole set of problems related to the consistent formulation of utility functions as summarized in Camerer and Weber (1992) and also discussed in some detail in Maes and Faber (2003). For this reason the axioms of utility theory proposed by von Neumann and Morgenstern (1943) have been heavily disputed during the last decades and various competing formulations suggested, see e.g. Kahnemann and Tversky (1979).

However interesting the theoretical mathematical implications of the effect of risk perception may be in regard to the formulation of utility functions, it is imperative not to loose focus on the characteristics of the decision problems at hand. Here it is suggested to differentiate between two different situations, namely 1) the situation where the purpose is to predict and represent the behavior and the attitudes of decision makers and 2) the situation where the purpose is to provide support for rational decision making.

The present paper is concerned about the appropriate representation of preferences and risk perception in risk based decision making in engineering problems, i.e. utility modeling. Problems are considered where the purpose is to provide support in the decision making for owners of engineering activities, owners of structures and authorities. The basic premise of this paper is that the risk aversion intrinsic to non-linear utility functions can almost always be explained by the non-inclusion of certain "follow-up" consequences. "Follow-up" consequences are, generally speaking, triggered by extreme losses, such as excessive business losses, loss of reputation or other indirect or so-called intangible losses. The non-inclusion of such losses occurs either voluntarily or involuntarily. Although in principle the use of appropriate non-linear utility functions and the inclusion of "follow-up" consequences are mathematically equivalent and may lead to identical decisions, only the latter approach leads to risk-consistent rational decision making.

Within the framework of risk based decision making, one is often faced with the problem of integrating risk perception into the formulation of the utility function. Often it is felt that a standard risk-neutral utility function is inappropriate as it fails to express the increasing undesirability of large consequences. In practical risk assessment various proposals for non-linear utility functions have been suggested and applied. Some of these questions prompted us to investigate the modeling and the use of utility functions in decision making. Utility theory has its roots in both cognitive psychology (as related to the modeling of the preferences of an individual) as well as macro-economics (with respect to a society's effective dealings with wealth, risk and policy making). As structural engineers striving to design and operate safe and cost-effective structural systems, we have borrowed extensively from these fields. In many applications this approach is justifiable, in others it may not be so clear, raising questions such as: "can we do better?" and "is what we are doing really justified?".

In subsequent sections, we will investigate how risk perception influences the modeling of utility functions. Specifically, we will be interested in learning which type of functions are suited to certain types of engineering decision making. We will also look at how the theory of preferences is linked to utility modeling. The elements of "rational" decision making, if such a concept truly exists, will be identified based on the behavior of individuals (including experts and non-experts) and on the behavior of groups, organizations, and the society. The fuzzy area between what is called risk aversion and complete description of consequences including follow-up consequences will be unraveled. Finally a formulation of generalized linear utility functions is provided which highlights the importance of including follow-up consequences into the preference representation as well as all prevailing aleatoric and epistemic uncertainties.

2 ENGINEERING DECISION MAKING

Decision making in civil engineering can be seen as being equivalent to participate in a game with nature acting as the main opponent, see also Ditlevsen and Madsen (1996). Considering Figure 1 the illustrated constituents of the decision problem system can be considered equivalent to the constituents of a game.

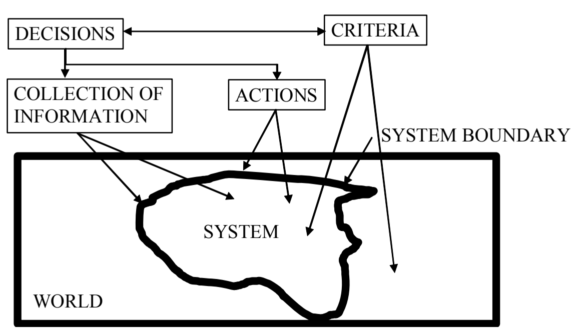

Figure 1. Main constituents in risk based decision analysis.

Knowing the rules of the game, i.e. the (success or acceptance) criteria, the system, the boundaries of the system, the possible consequences for the system and how all these factors are influenced by the world outside the system, is essential in winning the game. For this reason a very significant part of risk based decision making in practice is concerned about system identification/definition as well as the identification of acceptance criteria, possible consequences and their probabilities of occurrence. Playing the game is done by "buying" physical changes in the system or "buying" knowledge about the system such that the outcome of the game may be optimized.

2.1 Decision making based on expected utility

As outlined in Maes and Faber (2003) it can be stated that most decision makers and risk engineers would agree to the basic principle of ranking alternatives A based on their expected utility $E\left[U(A)\right]$ (von Neumann and Morgenstern, 1944)

$$E\left[U(A)\right] = \sum_{i=1}^{n_O} p(i|A)u(A,O_i) \tag{1}$$

where $E[\cdot]$ is the expectation operator, n_O is the number of possible outcomes associated with alternative A, $p(i|A)$ is the probability that each of these outcomes will take place (given A) and $u(A,O_i)$ is the utility associated with the set (A,O_i). Equation (1) assumes a discrete set of outcomes but can straightforwardly be generalized to continuous sample spaces.

As stated earlier the analysis of the utility function can serve two purposes; either for the prediction of the behavior of decision makers, or as the basis for rational decision making. Whereas this distinction from a mathematical point of view has no implications, it provides a useful guideline for maintaining the focus on decision making in engineering.

3 RATIONALITY AND IRRATIONALITY OF HUMAN BEHAVIOUR

Rational decision-making presumes first of all that the decision-maker behaves rationally. While the concept of human rationality has been basic to most economic analysis, it has not been as successful in the face of high-uncertainty situations. In fact, rationality has turned out to be a rather weak hypothesis, easily refuted and therefore not always useful as an axiomatic explanation of the theory of decision-making. The rationality or irrationality of choice has become a leading interest in the branch of psychology called cognitive psychology. This field of enquiry studies the capacity of human beings for perception and judgment. Of particular interest is the field of risk perception. Recent controversies over technological risks ranging from nuclear power to amalgam tooth fillings, from global warming to the consumption of mercury-saturated seafood, have sharpened interest in the way individuals form judgments about risk and act on them.

A stunning series of experiments on the so-called preference reversal phenomenon (Lichtenstein and Slovic 1971) have completely challenged the notion of rational decision-making. Preference patterns typical for gambling individuals are now explicitly accounted for in the insurance industry, in stock market behavior, and in advanced macroeconomic modeling. Perhaps then, it is necessary for engineers to demystify the ideal rational decision-maker pictured by von Neumann and Morgenstern. It can be countered that the irrationality of the individual cannot and, indeed, should not be extrapolated onto preference making by organizations and society (Rackwitz 2003).

In the more narrow context of decision making, rational behavior is, generally speaking, interpreted as decision making which maximizes expected utility (see Section 3.1), with utility being modeled such that it fairly and professionally represents the preferences of the stakeholders (Raiffa and Schlaiffer 1961). Given the origins of decision making theory, the type of decisions most often referred to in literature are related to economic and investment risks, and also to questions from game theory, gambling and strategy making. But as engineers, we are mainly concerned with two types of decisions:

– Operational decisions, where, for certain available resources, an optimal action is sought to counter a specific set of hazards; this includes applications ranging from operational and environmental safety, to optimal design decision making and life-cycle cost optimization.
– Strategic decisions, which, in addition to the above factors, also involve decisions regarding one's level of preparedness or anticipation; this includes decision making regarding maintenance planning or optimal inspection strategies.

In each case, it is important to define the stakeholders whose utility it is we wish to identify. In view of the complexity and the high degree of technical content of the decisions we face, it is important to communicate properly and effectively with (a) the stakeholders so as to ensure that we are aware of their preferences; and (b) with expert groups so as to ensure that our own insights are complete and correct. This is discussed further in Section 4.

4 RATIONALITY IN ORDERING PREFERENCES

4.1 What is Rational Decision Making?

The logical foundations of rational decision-making have been explored by a considerable number of scientists. The main axiomatic foundation in the case of pure traditional "frequentist probability" theory is due to von Neumann and Morgenstern (1944), and in the case of subjective probability theory, due to Savage (1954).

Three simple axioms form the foundation of expected utility analysis in which the optimal decision is the one that maximizes the utility of all alternatives. An excellent overview of the axiomatic basis is given in the appendix of a recent paper by Ditlevsen (2003). As pointed out by Machina (1982) the key axiom that addresses "rationality" is the so-called independence axiom. This axiom gives the expected utility theory its main empirical content by placing a restriction on the functional form of the preference, implying that it, or some monotonic transformation of it, must be "linear in the probabilities" and hence representable as the mathematical expectation of a von Neumann and Morgenstern utility function defined over the set of pure outcomes. The independence axiom states that "the alternative A is preferred (or, indifferent) to the alternative B, if and only if a p to (1-p) lottery between A and C, is similarly preferred (or, indifferent) to a p to (1-p) lottery between B or C, for arbitrary positive probabilities p and arbitrary alternatives A, B, and C."

4.2 Violations of the independence axiom

On first inspection, the independence axiom, seems to be reasonable and universally applicable, but three types of systematic violations have been uncovered and documented in recent years:

– the common consequence effect: its most famous example is the so-called Allais paradox (1952, p.89). A simple illustration of this paradox is given in Appendix A.

– the common ratio effect: an effective example is the so-called Bergen paradox (Hagen, 1979) supported by empirical evidence by Kahneman and Tversky (1979). A simple illustration of this type of violation is given in Appendix A.

– oversensitivity to tail changes or to changes in small probability events: this is related to the individual's perception of outlying and/or extreme events (Machina, 1982). The further a specific undesirable consequence such as a structural failure due to an extreme environmental event, is removed from the range of "normal" consequences or outcomes, the more sensitive the decision maker becomes to changes in the probability of this outlying consequence (with respect to changes in the probability of the more "normal" events"). In many instances this effect can be viewed as a factor in explaining an individual's behavior when facing choices involving the first two effects (common consequence and common ratio) as can be seen clearly from the small probabilities that sneek into the examples of these effects given in Appendix A. The fact is that individuals do react differently to small probability consequences. A further illustration of this type of violation is given in Appendix A.

Slovic and Tversky (1974) have fairly convincingly refuted specific counter-arguments raised by Savage (1954). Even so, the sheer strength of the expected utility theory, both in terms of its convincing simplicity, its predictive power, and its analytical elegance, seems to be almost fanatically supported by a large majority of users. The main thinking is one of "if it ain't broke, don't fix it": there is little point trying to rationalize specific "irrational" behavior using more advanced descriptive utility methods. In other words, any behavior that is not consistent with the axioms of the von Neumann-Morgenstern utility theory, while acceptable and "real" from a preference perception point of view, should be accepted as an interesting curiosum, but not as an aberration or an "irrationality".

4.3 Preference reversal and problem formulation

In addition to the above exceptions to the rational behavior fundamental to the maximum expected utility theory, we should also mention the well-known phenomenon of preference reversal first identified and studied by Lichtenstein and Slovic (1971). This apparent paradox refers to certain suitably chosen decision trees with two alternatives, having the following characteristics: when individuals are asked to choose between the two, they express preference for one or the other, but when they are subsequently asked to state the amount of money that, if given with certainty, would be indifferent to each alternative, they decide on amounts that are ranked in opposite order to the expressed preferences. This rather upsetting finding has been investigated widely by several researchers, but based on an analysis by Arrow (1982) it seems that this effect does not have as critical an impact on the fundamentals of utility theory as the objectives raised in Section 2.2.

Similarly, many psychometrists have observed that preference ranking can depend on how the several alternatives are presented or defined. For instance, it is well known in medical circles that patients and physicians alike, who are offered two or more choices of treatment for a specific disease, may come to different conclusions depending on whether empirical data are presented in terms of survival probability or probability of dying (McNeil et al. 1982). In engineering decision making, we similarly observe that a detailed description of some undesirable outcomes (for instance, the effects of an oil spill on sea birds) may affect the decision maker's risk perception and his assignment of probabilities.

5 NON-LINEAR UTILITY MODELS

5.1 Linearization of preference ordering

All we need to satisfy the axioms of completeness and transitivity for preference ordering is a real-valued functional $M(.)$ that is such that if alternative a_2 is preferred to a_1, then $M(.|a_2)>M(.|a_1)$. The outcome or consequence space is either discrete or continuous and therefore, each alternative *i* can be characterized either by a set of discrete probabilities $\boldsymbol{p_i}$ or else by a distribution function F_i . Without loss of generality, consider the continuous case. Hence the functional *M(F)* allows for a_2 to be chosen over a_1, if and only if $M(F_2) > M(F_1)$.

Let us now assume that *M* is differentiable in *F*, in other words *M* is smooth with respect to probabilities. Then we can apply a first-order Taylor expansion to *F(x)* about any distribution $F_0(x)$. But the first term of this expansion must be linear in the argument $(F-F_0)$, and hence linear in probabilities which means that it can be written as the expectation of a function $u(x|F_0)$, with respect to $(F-F_0)$. This function turns out to be the utility function, so that:

$$M(F) - M(F_0) = \int u(x \mid F_0)\left[dF(x) - dF_0\right] + 0(\| F - F_0 \|) \tag{2}$$

where $\| \cdot \|$ is the L' norm of the distance between two functions. In other words, the ranking of various alternatives *a* characterized by *F(x)* in the vicinity of F_0 occurs on the basis of

$$M(F) = \int u(x)dF(x) = E_F(u(x)) \tag{3}$$

Of course, this result corresponds to Von Neumann and Morgenstern's (1944) finding that linearity in probabilities (a direct consequence of the independence axiom discussed in Sections 3.1 and 3.2) is equivalent to maximization of expected utility. This is the basis of "linear utility theory" (LUT). Note that we ought to realize that "linear" means linear with respect to probabilities, irrespective of the curvature of the utility function *u(x)*.

LUT seems to be restrictive in as much as we clearly need higher order terms in (2) to be small with respect to the linear term. However, in most of engineering decision making the practical choice is between alternatives that alter probabilities over consequences, only slightly. For instance, increasing a design parameter by 20% will have a very small effect on the probability of dying of an individual who is possibly exposed to structural collapse. So, in most cases, linearization of some "general" preference functional *M(F)* or, equivalently, expected utility maximization, are excellent first-order tools.

5.2 Generalized Models

Several generalizations of LUT have been proposed. A first step up the ladder would be to consider a quadratic utility model by adding to (2) a second order term in the probabilities (Machina 1982). A general quadratic extension would require a rather awkward two-dimensional expression of the kind

$$M(F) = \int\int u(x,y)dF(x)dF(y) \tag{4}$$

It is more practical to consider a univariate simplification requiring two "utility functions" u_1 and u_2 in the expression for the preference functional:

$$M(F) = E_F(u_1(x)) \pm \tfrac{1}{2}\left(E_F(u_2(x))\right)^2 \tag{5}$$

which can be seen to be equivalent to a *u* function in the LUT model (2) which is not just a function of *x* but also of *F(x)*.

$$u(x \mid F) = u_1(x) \pm u_2(x)E_F(u_2(x)) \tag{6}$$

Allais' (1952) mean-and-variance model is in fact a special case of this quadratic utility model. It assumes that preferences among alternatives are based not just on the expected value of *u(x)* but also on its variance, resulting in:

$$M(F) = E_F(u(x)) - kE_F\left[u(x) - E_F(u(x))\right]^2 \tag{7}$$

which is equivalent to a LUT with the following *F*-dependent expression:

$$u(x \mid F) = u(x) - ku^2(x) + 2ku(x)E_F\left(u(x)\right) \tag{8}$$

The idea of using second moment expressions for *u* is also discussed by Ditlevsen (2003) in a discussion of the example described in Section A3 of Appendix A.

Other generalizations include the so-called "anticipated utility" models (Segal 1987) based on expressions of the type:

$$M(F) = \int u(x)d(h \circ F)(x) \tag{9}$$

showing the distribution *F(x)* to be transformed by a strictly increasing surjective function $h:[0,1] \rightarrow [0,1]$. The discrete equivalent would be $u(x)f(p)$. In other words, we are now "distorting" the probabilities in addition to distorting the consequences using *u*. The same idea is reflected in Puppe's (1991) utility-dependent probability distortion model with discrete expressions of the type $u(x)p^{g(x)}$. All of these high-order models are somehow successful in addressing the "inconsistencies" raised by the Allais, the Bergen, and other paradoxes (see Sections 3.2 and 3.3) but they are, of course, rather awkward to handle in practical situations.

5.3 Risk aversion and follow-up consequences

Ever since Bernoulli argued that individuals tend to display aversion to the taking of risks (with the notable exception of gambling) the standard way of modeling consequence ranking (Section 3) is the use of risk averse utility functions. Risk aversion can be defined as follows (a) if a sure gain $E_F(X)$ is always preferred to the distribution F_X itself; or (b) if starting from a position of certainty, the risk averter is unwilling to take a bet which is actually fair. Both concepts are equivalent to the concavity of a LUT utility function *u(x)*.

Although concavity $u''(x)<0$ reflects risk aversion, it is not a valid measure of it, as can be seen from the fact that it is not invariant under a linear transformation of the utility function (which should not alter preference ordering). Two appropriate measures of risk aversion can be obtained by defining the absolute risk aversion (ARA) and the relative risk aversion (RRA) factors, the latter being the elasticity of the marginal utility (Arrow, 1965):

$$\text{ARA} = -\frac{u''(x)}{u'(x)} \quad \text{and} \quad \text{RRA} = -\frac{\partial u'(x)}{u'(x)} / \frac{\partial x}{x} = -\frac{x\,u''(x)}{u'(x)} \tag{10}$$

It can easily be checked that a negative exponential function (Pratt 1964) and a simple generic power function (Arthur 1980) are strictly risk averse utility functions having constant ARA and RRA, respectively:

$$u(x) \sim -\exp(-\text{ARA}\, x) \quad \text{and} \quad u(x) \sim x^{1-\text{RRA}} \quad 0<\text{ARA};\ 0<\text{RRA}<1 \tag{11}$$

These functions, if properly scaled and valued to actual monetary units can easily be used in lifetime utility modeling and lifecycle cost optimization (Maes and Stewart 2004).

From the discussion in previous sections it should be clear that rational behavior requires that the decision maker is willing and able to identify all consequences, including all indirect and follow-on consequences. This requires considerable effort but it forces the decision maker to "explain" his specific aversion for large losses. In very many cases, groups as well as individuals who feel some sort of subconscious need to attribute more negative utility to large losses (Section 3.2), are in reality avoiding the challenging task of identifying large follow-up consequences. In the case of groups and organizations the preparedness to play a "fair game" against nature, should be, quite naturally, matched by the responsibility to include all indirect and future follow-on consequences in the decision analysis. Ultimately, the need to include risk aversion would disappear. The introduction of risk perception into societal decision making is often misunderstood in risk based decision making. It is as an example not unusual to see that risk acceptance criteria incorporate a certain effect of risk perception by being more restrictive for high consequence events. However, risk perception is a phenomenon which is highly dependent on the not only the severity of the consequences but also the scenario leading to the consequences. Therefore risk perception should be included in the risk based decision process not integrated into the risk acceptance criteria but rather as an integral part of the consequence analysis.

6 "FOLLOW-UP" EVENTS AND GENERALIZED UTILITY FUNCTIONS

In practical risk based decision making, various proposals for non-linear utility functions have been suggested and applied. But the problem remains which type of non-linear utility function would be appropriate in a given situation and whether it would lead to a rational decision. To overcome this problem (following an idea in Faber (2002)) a Bayesian approach is proposed here, which includes explicitly into the formulation of the utility function, the marginal utility of all possible outcomes which may occur as a consequence of the occurrence of other outcomes. The aim is to arrive at a formulation of the utility function that is (possibly) more complete, more transparent to the decision maker and one that explicitly takes into account the effect of epistemic uncertainties.

In Equation (12) an expansion or generalization of the utility function from Equation (1) is given.

$$\begin{aligned} E\left[U(A)\right] &= E_{\mathbf{e}}\left[\sum_{i=1}^{n_O} p(i\,|A,\mathbf{e})u(A,O_i) + \sum_{j}^{m} p(\mathbf{O}_j\,|A.\mathbf{e})u_{FO}(A,\mathbf{O}_j)\right] \\ &= \sum_{i=1}^{n_O} p(A,i)u(A,O_i) + E_{\mathbf{e}}\left[\sum_{j=1}^{m} p(\mathbf{O}_j\,|A,\mathbf{e})u_{FO}(A,\mathbf{O}_j)\right] \end{aligned} \tag{12}$$

In Equation (12) the probabilities $p(i|A,\mathbf{e})$ are aleatory probabilities, conditional on decision A and the outcome of the epistemic uncertainties $\mathbf{e}$. An additional term has been included to take into account marginal "follow-up" consequences. In this term m is the number of different combinations $\mathbf{O}_j$ of one or more of the n_O outcomes associated with the alternative A, $p(\mathbf{O}_j|A,\mathbf{e})$ is the probability that this combination occurs and $u_{FO}(A,\mathbf{O}_j)$ is the corresponding marginal utility. Notice here that the difference between the utility functions given in Equation (1) and Equation (12) is that m "follow-up" events are included which occur as conditional events with marginal utility $u_{FO}(A,\mathbf{O}_j)$ given the occurrence of at least one of the n_O outcomes. The probabilities $p(\mathbf{O}_j|A,\mathbf{e})$ may be assessed by means of probabilistic analysis conditional on the epistemic uncertainty $\mathbf{e}$ but may also be purely subjective in which case, however, the expectation operation becomes obsolete. In the latter case the subjective probabilities may be updated using a Bayesian framework at the same rate as evidence becomes available.

Utility U for civil engineering decision making purposes can be formulated in terms of monetary benefit B

$$U = B = I - C \tag{13}$$

where the benefit B is expressed in terms of the difference between committed incomes I and expenditures C.

For a structural design decision problem Equation (13) may be expanded as

$$U = B = I - C = I - C_D - C_{OM} - C_F - C_C \quad (14)$$

where C_D, C_{OM}, C_F and C_C denotes design costs, operating and maintenance costs, failure costs and decommissioning costs respectively.

In accordance with the axioms of von Neumann and Morgenstern (1944) decisions shall be based on the maximum expected utility:

$$E[U] = E[B] = E[I] - E[C] = E[I] - E[C_D] - E[C_{OM}] - E[C_F] - E[C_C] \quad (15)$$

In some cases expected value decisions based directly on formulations as given in Equation (15) do not reflect the complete picture for a given decision maker. One of the typical reasons for this is that decisions often are made subject to certain boundary conditions, e.g. the decision making is constrained by budgetary limits or integrated into the complex problem of portfolio management. This essentially implies that in some cases e.g. the costs C, as well as the income can include additional "follow-up" consequences.

As an example consider the situation where a decision maker needs to identify an optimal decision subject to budget constraints. Assume that if the given budget is exceeded a bank loan must be obtained to cover the budget overrun (loss). Following the formulation given by Equation (12), Equation (14) can be written as:

$$\begin{aligned} U &= I - C_D - C_{OM} - C_F - C_C - w(I - C_D - C_{OM} - C_F - C_C) \\ &= I - C_D - C_{OM} - C_F - C_C - w(L) \end{aligned} \quad (16)$$

where $w(\cdot)$ is some monotonically increasing function of the loss L. Equation (16) yields the following expected utility now:

$$E[U] = E[I] - E[C_D] - E[C_{OM}] - E[C_F] - E[C_C] - E[w(L)] \quad (17)$$

Although in principle the use of appropriate non-linear utility functions and the inclusion of "follow-up" consequences are mathematically equivalent and may lead to identical decisions, only the latter approach leads to risk-consistent rational decision making.

7 DO INDIVIDUALS NEED TO BE EXPERTS TO MAKE RATIONAL DECISIONS?

7.1 Introduction

Two fundamental assumptions of maximum expected utility theory are that:

– decision makers are willing and able to calculate the probabilities or probability distributions needed for their analysis. This raises the issues of incomplete information, experience bias, the use of epistemic uncertainty, and Bayesian probability updating. These issues are discussed in Sections 7.2, 7.3, 7.4, and 7.5 respectively.

– decision makers are willing and able to update the utility functions they use. In other words, they must realize that utility functions are expressions of preference rankings, and therefore they must change utilities in accordance to their tastes and in pace with their learning curve. This discussed in Section 7.6.

7.2 Incomplete information

Individuals do not always have all the information that may be available to the society as a whole. An individual or an organization making a decision on the basis of limited information is often aware of the incompleteness of their information. Seeking assistance comes at a price, and the benefit gained from requesting expert input, may or may not be worth the expense. Rational decision makers would voluntarily try to delay the making of an important choice until professional advice is available, just as individuals turn to physicians or counselors in turn of need. But in decision making, the resistance to requesting or admitting expert advice may be considerably greater, not just due to the cost of engaging experts, but also due to the attitude of overconfidence (Lichtenstein and Slovic 1971). Decision makers often feel that they alone are the supreme experts of their own problem: they are in control, and they will be facing the consequences of their decision.

7.3 Bias towards one's own experience

At a later stage, individuals may eventually come to use and trust expert opinion, but they will still be largely biased by their own experience, correctly or incorrectly. A typical example is the workplace, where decision makers are often influenced by their own experience (Viscusi 1979) in dealing with occupational hazards and small technological risks. Personal experience frequently takes precedence over more or less objective data in the form of injury data, or component failure data.

7.4 Difficulties in using epistemic uncertainty

Consider now more experienced decision makers who, after collecting all possible data and accounting for expert opinion, still feel that their information is limited. They are fully prepared to accept the limitations of the analysis and they wish to include epistemic probabilities in their decision making. Most engineers are sufficiently familiar with Bayesian thinking to agree that in making optimal decisions, there is no point in differentiating between first-order physical uncertainty and second-order epistemic uncertainty. In other words, probabilities, whether objective or subjective are what they are, probabilities and it is meaningless to assign a degree of uncertainty to uncertainty (de Finetti, 1990; Marschak 1975).

In Savage's (1954) maximum expected utility theory there is in fact no room for someone who talks about reliable probabilities and unreliable probabilities. Every decision maker is expected to express in each probability the best possible state of belief. However, rational informed behavior apparently is difficult when decision makers do not feel that they know the rules of the game they are playing. This is reflected clearly by the Ellsberg's paradox constructed as a criticism of Savage's axioms by Ellsberg (1961): an example is given in Appendix A, section A4. As summarized in e.g. Camerer and Weber (1992) several approaches have been suggested for the treatment of ambiguity averseness reaching from a modification of the utility subject to epistemic uncertainty over a non-linear representation of second order probabilities to the concept of minimum subjective expected utility and non-additive probabilities. Even though such approaches may be applicable to represent a certain behavior of decision makers in given situation it is still highly questionable whether this is relevant – assuming that the decision maker would act rationally provided the right information. Furthermore the introduction of such approaches may also give rise to a whole spectrum of inconsistencies in the even larger field of subjective or Bayesian probability.

7.5 Problems with Bayesian updating of probabilities

At a more advanced level we find decision makers who are fully aware that their probabilities are based on incomplete data, and who are prepared to account properly for epistemic uncertainties. Furthermore they are prepared to abide by Bayes' theorem to process information as it becomes available. The problem is, however, that Bayesian updating is not always applied correctly. Extensive psychometric studies (Tversky and Kahneman 1974) have shown that generally speaking, too little weight is given to prior information, and too much importance is given to new data. Newsworthiness seems to systematically win over long accumulated prior knowledge. Similarly, many media-dominated present-day societies seem to be characterized by excessive volatility in policy making: regulations, laws, policies, preferences are changing quickly in response to "new" findings notwithstanding decades or centuries of valuable data.

7.6 Bayesian updating of utility functions

Other than subjective probabilities (representing knowledge and belief), the decision maker must also be willing and able to update subjective utilities (representing preferences and tastes). In models of adaptive utility (Witsenhausen 1974; Cyert and DeGroot 1980) the notion of learning is applied to utility as well as probabilities. The situation where this most critically applies is when certain consequences are new and have not been fully appreciated by the stakeholders. For instance, it is very hard to express one's preference between mango and plums if one has never tasted a ripe mango before. Similarly, it is hard to construct appropriate utility functions capturing the loss of fish and sea birds as a result of pollution, if one has never witnessed this outcome before.

For a n-dimensional vector of consequences $\boldsymbol{x}$, we can express preferences using a multi-attribute utility function $u(\boldsymbol{x},\boldsymbol{\theta})$ where $\boldsymbol{\theta}$ represents a vector of uncertain "taste" parameters. Consistent with Bayesian thinking, we can assign a prior joint pdf to these parameters. As preference ranking data become available subsequently, this prior distribution may be updated using likelihood functions involving $u(\boldsymbol{x},\boldsymbol{\theta})$ expressing how likely it would have been to prefer certain alternatives a_i over others given a vector of taste parameters. For instance, Cyert and DeGroot (1980) use a very flexible:

$$u(\boldsymbol{x},\boldsymbol{\theta}) = \sum_{i=1}^{n} \theta_i \log x_i \quad \text{where} \quad \theta_i > 0 \quad \text{and} \quad \sum_{i=1}^{n} \theta_i = 1 \tag{18}$$

to update utilities conditional on stated preferences regarding the choices a_i. The authors are not aware of the use of utility updating in engineering applications, but the basic idea is certainly interesting.

REFERENCES

Allais, M. 1952. The foundations of a positive theory of choice involving risk and criticism of the postulates and axioms of the American School. In Allais & Hagen 1979.

Allais, M. & Hagen, O. 1979. *Expected utility hypotheses and the Allais paradox*. D. Reidel, Holland.

Arrow, K. J. 1965. *Aspects of the theory of risk-bearing*. Lecture 2. Yrjö Jahnssonin Säätio, Helsinki.

Arrow, K.J., 1971. *Essays in the Theory of Risk Bearing*, Chicago, Markham.

Arthur, W. B. 1980. The economics of risks to life. *The American Economic Review*.

Arrow, K. J. 1982. Risk perception in psychology and economics. *Economic Inquiry* 20:1-9.

Camerer C. and Weber, M. 1992. *Recent Developments in Modeling Preferences: Uncertainty and Ambiguity*, Journal of Risk and Uncertainty, Kluwer Academic Press, pp. 325-370.

Cyart, R. M. & DeGroot, M. H. 1980. Learning applied to utility functions. *Bayesian Analysis in Econometrics and Statistics* 1:159-168.

de Finetti, B. 1990. *Theory of Probability*. John Wiley & Sons.

Ditlevsen, O. and Madsen, H.O., 1996. *Structural Reliability Methods*, Wiley, Chichester.

Ditlevsen, O. 2003. Decision modeling and acceptance criteria. *Structural Safety* 25:165-191. Elsevier Science.
Ellsberg, D. 1961. Risk ambiguity and Savage axioms. *Quarterly Journal of Economics* 75:57-67.
Faber, M.H., 2002. *Risk Assessment and Decision Making in Civil Engineering*, Workshop on Reliability-Based Design and Optimization: RBO'02. September 23-25, 2002, Warsaw, Poland.
Faber, M.H., 2003. *Uncertainty Modeling and Probabilities in Engineering Decision Analysis*. In proceedings to the International Conference on Offshore Mechanics and Arctic Engineering, OMAE2003 in Cancun, Mexico.
Hagen, O. 1979. Towards a positive theory of preferences under risk. In Allais & Hagen 1979.
JCSS, 2001. *Probabilistic model code*. The Joint Committee on Structural Safety, http://www.jcss.ethz.ch/.
Kahneman, D. & Tversky, A. 1979. Prospect theory: an analysis of decision under risk. *Econometrica* 47 263:291.
Lichtenstein, S. & Slovic, P.1971. Reversal of preferences between bids and choices in gambling decisions. *Journal of Experimental Psychology*, 89:46-55.
Machina, M. J. 1982. Expected utility analysis without the independent axiom. *Econometrica* 50: 277-323.
McNeil, B. J., Pauker, S. G., Sox, H. C. Jr. & Tversky, A. 1982. On the elicitation of preferences for alternative therapies. *New England Journal of Medicine* 306:1259-1262.
Maes, M. and Faber, M.H. 2003. *Issues in Utility Modeling*, Proceeding to the 11'th IFIP WG 7.5 Working Conference on Reliability and Optimization of Structural Systems, November 2-5, 2003, Banff, Alberta, Canada.
Maes, M.A. & Stewart, M. Optimizing structural safety levels on the basis of lifetime utility objectives of the individual. To be published in *Proc. SEMC* 2004. Cape Town, South Africa.
Marschak, J. 1975. Do personal probabilities of probabilities have an operational meaning? Theory and Decision 6:435-453.
Pratt, J. W. 1964. Risk Aversion in the Small and in the Large. *Econometrica* 32, 1-2: 122-136.
Puppe, C. 1991. *Distorted probabilities and choice under risk*. Lecture Note in Economics and Mathematical Systems 363. Springer-Verlag.
Rackwitz, R. 2003. Risikowahrnehmung and rationale Riskobewaeltigung. *Draft manuscript, private communication*.
Raiffa, H. & Schlaifer, R. 1961. *Applied Statistical Decision Theory*. Division of Research, Graduate School of Business Administration, Boston, Harvard University.
Sahlin, N-E. 1983. On second order probabilities and the notion of epistemic risk. *Foundations of Utility and Risk Theory with Applications, Editors: Stigum & Wenstop, D. Reidel publishing company 95-104.*
Savage, L. J. 1954. *The foundations of Statistics* John Wiley & Sons.
Segal, U. 1987. The Ellsberg Paradox and risk aversion: an anticipated utility approach. *International Economic Review* 28:175-202.
Slovic, P. & Tversky, A. 1974. Who accepts Savage's axiom. *Behavioral Science* 19:368-373.
Tversky, A. & Kahneman, D. 1974. Judgment under Uncertainty: heuristics and biases. *Science* 185:1124-1131.
Von Neumann, J. & Morgenstern, O. 1944. *Theory of Games and Economical Behavior*. Princeton University Press (also: 3rd Edition, 1953, John Wiley & Sons).
Viscosy, W. K. 1979. *Employment Hazards: An investigation of market performance*. Harvard University Press, Cambridge.
Witsenhausen, H. S. 1974. On the uncertainty of future preferences. *Annals of Economics and Social Measurement* 3:91-94.

APPENDIX A: EXAMPLES OF FOUR TYPES OF DOCUMENTED VIOLATIONS OF THE AXIOMATIC BASIS OF THE EXPECTED UTILITY MODEL

Four types of empirically verified behaviour inconsistent with the axioms of the Von Neumann-Morgenstern utility theory are discussed in Section 3.2. The first example A1 illustrates the "common consequence effect", the second example A2 illustrates the "common ratio effect" and the third example A3 shows the importance of perceived tail effects. Example A4 pertains to an inconsistency that arises in Savage's (1954) extension of the Von Neumann-Morgenstern theory to epistemic probabilities. In all of the following m represents a large sum of money, for instance m=\$1,000,000.

A1: The Allais (1952) paradox.
Rank the following two alternatives (a_1, a_2):

$$a_1: \ 100\% \text{ chance of winning } m$$

$$a_2: \begin{cases} 10\% & \text{chance of winning } 5m \\ 89\% & \text{chance of winning } m \\ 1\% & \text{chance of winning } 0 \end{cases}$$

Also rank the following two alternatives (a_3, a_4):

$$a_3: \begin{cases} 10\% & \text{chance of winning } 5m \\ 90\% & \text{chance of winning } 0 \end{cases}$$

$$a_4: \begin{cases} 11\% & \text{chance of winning } m \\ 89\% & \text{chance of winning } 0 \end{cases}$$

For a rational decision maker with an arbitrary utility function u, the difference in expected utility between the alternatives a_1 and a_2 is equal to:

$$\Delta_{1-2} = u(m) - [0.1u(5m) + 0.89u(m)] = 0.11u(m) - 0.1u(5m)$$

whereas the difference between the expected utilities of a_3 and a_4 is exactly the opposite of this amount:

$$\Delta_{3-4} = 0.1u(5m) - 0.11u(m) = -\Delta_{1-2}$$

Consequently, a rational expected utility maximizer will prefer either a_1 and a_4, if her utility function is such that Δ_{1-2} is positive, or else a_2 and a_3, if her utility function is such that Δ_{1-2} happens to be negative. Either way, independent of the functional form of the utility function, the "rational" choice disagrees with the experimental finding (Allais 1952; Slovic and Tversky 1974) that a large majority of individuals prefer a_1 over a_2, and a_3 over a_4. This type of violation of the independence axiom has been generalized to other examples involving choices between pairs of alternatives containing "common consequences".

A2: The Bergen paradox (Hagen, 1979)
Rank the following two alternatives:

$$a_5: \begin{cases} p & \text{chance of winning } m \\ 1-p & \text{chance of winning } 0 \end{cases}$$

$$a_6: \begin{cases} q & \text{chance of winning } km \\ 1-q & \text{chance of winning } 0 \end{cases}$$

Also rank the following two alternatives:

$$a_7: \begin{cases} \alpha p & \text{chance of winning } m \\ 1-\alpha p & \text{chance of winning } 0 \end{cases}$$

$$a_8: \begin{cases} \alpha q & \text{chance of winning } km \\ 1-\alpha q & \text{chance of winning } 0 \end{cases}$$

where $p > q, 0 < \alpha < 1, \text{and } k > 1.$ Note that the ratios p/q is "common" in both ranking tests, hence the name "common ratio effect". It is clear from the independence axiom that a rational expected value maximizer would rank a_5 and a_6 exactly the same as a_7 and a_8. But this contradicts ample experimental evidence. For instance, in the Kahneman and Tversky (1979) tests with p=0.9, q=p/2, k=2, it appeared that while 86% of their subjects preferred a_5 over a_6, only 27% preferred a_7 over a_8 if a reduction factor

$\alpha = 0.0022$ was applied so that $(p = 0.9, q = 0.45)$ would become $(\alpha p = 0.002, \alpha q = 0.001)$. Other tests with different parameters p, q and α showed similar results.

A3: Tail effects and sensitivity to outliers

An interesting example is provided by Ditlevsen (2003) by contrasting two loss probability functions. The first one is a long-tailed gamma PDF with α=1/9, β=3 and a lower bound of 1. The second represents a discrete loss with a 90% probability of 1 and a 10% probability of 13/3. The means and standard deviations of both losses are identical, but clearly, the tails are not. A strictly risk neutral decision maker (linear utility function) would be indifferent with respect to expressing a preference between the two loss profiles. A risk averse individual wishing to avoid facing a large loss of 13/3, would supposedly prefer the continuous gamma type loss. However, the use of a continuous risk averse utility function of the type $u(x) \sim x \exp(-kx)$ rather irritatingly shows that the greater the concavity of the utility function, and hence the greater the aversion to large losses, the greater the difference in expected utility in favor of the opposite choice (the discrete loss function).

A4: The Ellsberg (1961) paradox:

Consider an urn known to contain 20 yellow balls and 40 green and blue balls, the latter in an unknown proportion. A single ball is drawn from the urn. Rank the following two alternatives:

$$a_9 : \begin{cases} \text{receive } \$100 & \text{if the ball is yellow} \\ \text{receive } \$0 & \text{if the ball is green or blue} \end{cases}$$

$$a_{10} : \begin{cases} \text{receive } \$100 & \text{if the ball is blue} \\ \text{receive } \$0 & \text{if the ball is yellow or green} \end{cases}$$

Also rank the following alternatives:

$$a_{11} : \begin{cases} \text{receive } \$100 & \text{if the ball is yellow or green} \\ \text{receive } \$0 & \text{if the ball is blue} \end{cases}$$

$$a_{12} : \begin{cases} \text{receive } \$100 & \text{if the ball is blue or green} \\ \text{receive } \$0 & \text{if the ball is yellow} \end{cases}$$

It can easily be seen that the decision maker following Savage's maximum expected utility theory would include the epistemic uncertainty associated with the unknown ratio of green and blue balls in the analysis and would be indifferent to choosing between a_9 or a_{10}, and between a_{11} or a_{12}. But Ellsberg (1961) observed that a majority of individuals tested, systematically preferred a_9 to a_{10} and a_{12} to a_{11}.