Decision Theory

Mendoza, M.
Dept. Statistics, ITAM
Rio Hondo 1, San Angel.
México 01000 D.F. MEXICO

Gutiérrez-Peña, E.
Dept. Probability and Statistics, IIMAS-UNAM,
Ciudad Universitaria, México 01000 D.F.
Apartado Postal 20-726, Admón. No. 20, MEXICO

1. Introduction.

1.1. ***The Process of Decision-making.***

One of the most common activities of human beings is that of decision-making. Every person is constantly deciding on a wide variety of different subjects. There are easy as well as difficult decisions; there are important and irrelevant decisions; one must face personal and professional decisions. In the end, we all know, from our particular experience, that there are good and bad decisions, so a natural question arises. Are there any rules or procedures for decision-making which “guarantee” that the final result is a “good” decision?

A vast amount of effort has been devoted to explore this subject. Psychologists have studied how decision-makers work under different conditions. From a philosophical perspective, even the existence of such a thing as a “good decision” has been questioned. The logical approach has contributed to the understanding of the decision-making process. Mathematics, including Statistics, has played a major role in providing a formal structure for the process and defining criteria for optimality.

Under the heading of Decision Theory, the literature offers an account of the ways people actually make decisions and a discussion on the mechanisms underlying this behavior. This is called a “descriptive” decision theory. On the other hand, we can also find discussions about the principles to consider to make rational decisions. In this case, we have a “normative” decision theory.

In this article, we are concerned with a normative decision theory. We discuss the solution of a general class of decision problems and, briefly but not less importantly, comment on the relationship between Decision Theory and Statistical Inference.

1.2. ***Decision problems.***
A decision problem can be defined as a situation where a person or a group of people (the decision-maker) must select one and only one element (an action) from a given set $\boldsymbol{A} = \{a_1, a_2, \ldots, a_k\}$. The idea is to choose the "best" action, and thus the subjective nature of the solution arises naturally: The best action must be such for the specific decision-maker. For this purpose, every action in $\boldsymbol{A}$ is judged in terms of the consequence it produces. Therefore, another component of the problem is the set of consequences $\boldsymbol{C} = \{c_1, c_2, \ldots, c_k\}$ where c_i stands for the consequence from action a_i. If, for every action, the corresponding consequence is completely known and occurs every time the action is selected, then we have a decision problem without uncertainty.

Under these circumstances, choosing an action is equivalent to choosing a consequence and the best action will be that leading to the best or "most preferable" consequence. Thus, the decision-maker must define the set of actions $\boldsymbol{A}$, the set of consequences $\boldsymbol{C}$, and must also express her personal subjective preferences among the different elements of $\boldsymbol{C}$. Once this structure is defined, the problem is solved by choosing the action whose consequence is the most preferable. This seems to be a rather simple exercise; however, the importance of a careful definition of the problem cannot be overestimated. The set $\boldsymbol{A}$ must include all available actions, and these must represent actual alternative options.

Concerning the principles we referred to in Section 1.1, in the case of a decision without uncertainty, these are implicit and apply to the decision-maker's preferences. They assume all actions to be "comparable" in terms of preferences and the preference relation to be "transitive". If the structure of the problem is not well-defined or the axioms (coherence principles) are not fulfilled, the decision-maker will solve (if at all) a wrong problem.

If preferences are expressed in terms of a numerical score u such that, for every pair (i, j), c_i is less preferable than c_j if and only if $u(c_i) < u(c_j)$, then the problem reduces to maximizing the function u over $\boldsymbol{C}$. The score u is called the "utility function" and the solution is any action that maximizes the decision maker's subjective utility. Sometimes it is easier to use a "loss function" to represent preferences. If l is a loss function, then c_i is less preferable than c_j if and only if $l(c_i) > l(c_j)$. In this case, the best action minimizes the decision maker's subjective loss. As long as both functions, loss and utility, describe the same preferences, the optimal decision remains the same.

Decision problems without uncertainty are conceptually simple. Nevertheless, their solution can be difficult in practice. In particular, translation of preferences into a numerical score is not easy. However, the utility function is just a numerical representation of the preferences and thus these technical problems are usually due to the difficulties associated with the elicitation of preferences.

Recalling the question "Are there any rules or procedures for decision-making which guarantee that the final result is a good decision?", we can now say that, in the case of decision problems without uncertainty, if the problem is properly structured, the utility function really reflects the decision-maker's preferences, and these preferences obey the coherence principles, the answer is yes.

Unfortunately, real world decision problems are usually more complex than those without uncertainty. It often happens that once an action has been chosen, there is a set of possible consequences associated with it. Among them, one and only one will take place depending upon the occurrence of some uncertain event. The structure of a decision problem "under uncertainty" is described in Section 2.1.

1.3. *Historical Background and Related Fields.*

Ideas related to decision theory can be traced back to at least the 18th century. However, most of the formal developments emerged during the past 80 years and many disciplines have contributed to the methods of decision-making. The case without uncertainty belongs to the domain of Optimization and Operations Research. Game Theory provided the basic framework for the case under uncertainty. Also, modern Economic Theory deals with interesting decisions, such as the choice of an investment portfolio. In fact, this problem is the origin of an entire field known as Financial Economics.

The relationship between Statistics and Decision Theory deserves a special mention. Statistics can contribute to the solution of decision problems under uncertainty by means of methods which allow the decision maker to describe her own uncertainty. The result is known as Statistical Decision Theory, and different approaches to statistical inference have led to different statistical decision theories. Conversely, any situation where a statistical inference is required can be seen as a decision problem under uncertainty (an inference is just an assertion regarding the phenomenon of interest and is chosen among several alternatives). From this perspective, a Theory of Statistical Inference can be developed on the grounds of a specific Decision Theory. The most outstanding example is Bayesian Statistics, a statistical theory built upon the axiomatic decision theory described in Section 2.3. Bayesian inference, as a theoretical discipline, and Bayesian methods, as a set of tools for inference in practice, have been growing very rapidly in recent years.

Decision Theory is an interdisciplinary field with many relevant contributions published in Economics, Mathematics, Psychology and Statistics, among other areas. The basic concepts in Section 2.3 were mainly developed from the 1930's to the 1970's and now appear in a number of Bayesian Statistics texts.

1.4. *Decisions in Education.*

Decisions are omnipresent in the field of education. A specific syllabus must be chosen, among different alternatives, for each course. The admissions committee must decide whether or not an applicant should be accepted as a student of the program. As part of the grading process, for every exam, a cutoff value must be selected to decide whether a student fails. On the other hand, students face a decision problem when they choose a career; every time they select an elective course they are making a decision; when a student chooses an answer in a multiple response test, she is also deciding among a set of alternative options. These are only a few examples of decision problems in education.

For a long time, these problems were faced without any decision-theoretical backing. However, current research in education, specifically on test design, involves decision theory as a basic resource. As discussed by van der Linden (1991), this process is related to a change in the approach to design. Formerly, a test was considered as a measurement tool and research on tests was directed to explore the relation among the observed measurement and the attribute to be measured. More recently, it has been argued that the measurement is relevant insofar as it is useful to make decisions about the examinee. Therefore, an appropriate design must take into account the specific decision problem to be addressed. For examples of this idea we refer the reader to Novick and Lindley (1978), where some particular utility functions are used to identify cutoff scores; Sawyer (1996), who uses decision theory to validate course placement tests; and Vos (1999), where optimal sequential mastery tests are formulated using decision-theoretical arguments.

2. Decision Theory.

2.1. *Definition of a Decision Problem.*

A decision problem under uncertainty is defined by the following elements:

i) $\boldsymbol{A} = \{a_1, a_2, \ldots, a_k\}$,

ii) $\boldsymbol{E} = \{E_{11}, E_{12}, \ldots, E_{1m_1}; E_{22}, E_{21}, \ldots, E_{2m_2}; \ldots; E_{k1}, E_{k2}, \ldots, E_{km_k}\}$, and

iii) $\boldsymbol{C} = \{c_{11}, c_{12}, \ldots, c_{1m_1}; c_{22}, c_{21}, \ldots, c_{2m_2}; \ldots; c_{k1}, c_{k2}, \ldots, c_{km_k}\}$.

Here, $\boldsymbol{A}$ is an exhaustive and exclusive set of actions; $\boldsymbol{E}$ is the set of uncertain events where, for every action a_i, the collection of events $\boldsymbol{E}_i = \{E_{i1}, E_{i2}, \ldots, E_{im_i}\}$ is assumed to be a partition of the certain event Ω. Finally, $\boldsymbol{C}$ is the set of consequences and is such that, to each pair (a_i, E_{ij}), there corresponds a consequence c_{ij}. Both the space of actions and the set of uncertain events may contain an infinite number of elements.

In order to solve this problem, the simple optimality principle used in Section 1.2. no longer applies. We cannot substitute the choice of an action by the choice of a consequence. In fact, in this new setting, once the decision-maker chooses a_i, what she gets is the result of a "lottery" whose possible prizes are $\{u(c_{i1}), \ldots, u(c_{imi})\}$ where $u(c_{ij})$ occurs with probability $p_{ij} = P(E_{ij})$ and, for each i, $p_{i1} + p_{i2} + \cdots + p_{imi} = 1$. Here, there is no obvious way to transfer a utility score to the action a_i from the corresponding lottery. In fact, every method devised to solve a problem of this kind relies on a specific proposal of how to carry this out. Also, since the event E_{ij} is uncertain for the decision maker, the probability p_{ij} is a measure of her personal beliefs about the occurrence of E_{ij}. Thus, both utilities and probabilities are subjective.

A number of methods have been proposed to solve decision problems under uncertainty; all of them replace the original problem with another which, in a way, "does not involve any uncertainty". See Section 2.2.

Example. A Graduate Studies Committee must evaluate every application to a Ph.D. program and decide whether the corresponding candidate can be admitted. In a simplified version, only two actions are considered. Thus $\boldsymbol{A} = \{ a_1, a_2 \}$, where a_1 = "admit the candidate", a_2 = "reject the candidate". The consequences of each of these actions are uncertain. Rejection of a candidate would be appropriate if this action prevents admission of a bad student. On the other hand, it would be a mistake to reject a good student. If the set of uncertain events is defined as $\boldsymbol{E} = \{ E_1, E_2 \}$, with E_1 = "candidate's performance is good", E_2 = "candidate's performance is poor", then the consequences are given by $\boldsymbol{C} = \{ c_{11}, c_{12}, c_{21}, c_{22} \}$ where c_{ij} is the consequence obtained if the action a_i is chosen and the uncertain event E_j occurs. Here, c_{11} and c_{22} represent good outcomes whereas c_{12} and c_{21} are mistakes. A reasonable utility function describing these preferences must be such that $u(c_{ii}) > u(c_{ij})$ for $i, j = 1,2$. This problem has the structure described by Figure 1.

<Figure 1 near here>

2.2. *Intuitive Solutions to a Decision Problem.*

Here, we review two popular methods used to solve a decision problem under uncertainty. As many others, each of them proceeds in a two-step fashion. First, for every action, a utility score is derived from the corresponding lottery. In the second step, actions are compared as they would be in a problem without uncertainty so the best action is chosen as that maximizing the derived utility.

Minimax. For each action a_i, this method looks at the worst possible consequence and then acts as if this consequence will occur for sure. In other words, Minimax assigns the utility $u_M(a_i) = \min_j u(c_{ij})$ and this new utility is then maximized over the space of actions, so that the best action, denoted by a^*, is such that $a^* = \max_i \min_j u(c_{ij})$. The method is named after this last expression: Maximin, if we work in terms of utilities or Minimax if a loss function is used. This procedure completely ignores the probabilities that the decision-maker has assigned to the events in $\boldsymbol{E}_i$. In some sense, this is a pessimistic approach to decision-making and often yields to what is known as "opportunity loss".

Maximum expected utility. This method assigns, to each action a_i, the weighted average of the utilities assigned to the corresponding consequences, the weights being the respective probabilities, i.e. $u_E(a_i) = u(c_{i1})\, p_{i1} + u(c_{i2})\, p_{i2} + \cdots u(c_{im_i})\, p_{im_i}$. This expected utility is then maximized over the space of actions, so that the best action, denoted by a^*, is such that $a^* = \max_i u_E(a_i)$. This approach to decision-making takes into account all the information and is the only one that is consistent with the axiomatic theory described in the next section.

Example (ctd.).

Let $p_1 = P(E_1)$ and $p_2 = P(E_2)$. In this case, the expected utilities are computed as $u_E(a_1) = u(c_{11})\, p_1 + u(c_{12})\, p_2$ and $u_E(a_2) = u(c_{21})\, p_1 + u(c_{22})\, p_2$, and the action leading to the larger of these two quantities is chosen. For the sake of simplicity, let us consider the case where the preferences are described by means of a loss function. Moreover, let us assume that consequences associated with successes are assigned a zero loss so that $l(c_{11}) = l(c_{22}) = 0$ (with $l(c_{ij}) > 0$; $i \neq j$). In this setting, $l_E(a_1) = l(c_{12})\, p_2$, $l_E(a_2) = l(c_{21})\, p_1$ and then a_1 is chosen as the optimal action if, and only if, $l(c_{12})\, p_2 < l(c_{21})\, p_1$ or, equivalently, $p_1 / p_2 > l(c_{12}) /$

$l(c_{21})$. Here, the cutoff value depends on the relative importance of c_{12} and c_{21}. Note that, even if $p_1 \le p_2$, a_1 will still be chosen as the best action if $p_1 > (l(c_{12}) / l(c_{21}))\, p_2$.

2.3. *Axiomatic Decision Theory*.
Coherence axioms can be formulated in several different ways. One of the most intuitive discussions concerning the axiomatic basis for decision theory appears in Lindley (1972). For a more technical version the reader is referred to Bernardo and Smith (1994).

Let us represent an action a by $a = \{c_1 \mid E_1, \ldots, c_m \mid E_m\}$, thus indicating that an action is just a lottery where the decision-maker gets the consequence c_j whenever the event E_j occurs. Then it is clear that any consequence c is a particular case of an action since c is equivalent to $\{c \mid \Omega\}$, where Ω is the certain event.

Axiom 1 (Comparability). For every pair of actions a_1 and a_2 in $\boldsymbol{A}$, one and only one of the following conditions holds:

a_1 is less preferable than a_2	(denoted $a_1 < a_2$)
a_1 is more preferable than a_2	(denoted $a_1 > a_2$)
a_1 and a_2 are equally preferable	(denoted $a_1 \sim a_2$)

Moreover, it is possible to find two consequences c^* and c_* such that $c^* > c_*$ and $c_* \le c \le c^*$ for any consequence c. ($a_1 \le a_2$ means "a_1 is not more preferable than a_2"; $a_1 \ge a_2$ is similarly defined.)

Discussion. The comparability of actions implies both the comparability of consequences and the comparability of events. Denote by $\bar{E}$ the complement of the event E and let $a_1 = \{c^* \mid E_1, c_* \mid \bar{E}_1\}$ and $a_2 = \{c^* \mid E_2, c_* \mid \bar{E}_2\}$. Then comparing a_1 with a_2 is equivalent to comparing the likelihoods of the events E_1 and E_2: a_1 will be preferable to a_2 if and only if E_1 is more likely than E_2.

Axiom 2 (Transitivity). If $a_1 > a_2$ and $a_2 > a_3$, then $a_1 > a_3$.

Discussion. This axiom assumes that the set of actions can be ordered in such a way that a search for the most preferable element makes sense.

Axiom 3 (Substitutability and dominance). If $a_1 > a_2$ when the event E occurs and $a_1 > a_2$ when the event $\bar{E}$ occurs, then $a_1 > a_2$.

Discussion. This axiom implies that, in any given action $a = \{c_1 \mid E_1, \ldots, c_m \mid E_m\}$, a consequence c_j can be replaced by any action which is equivalent to c_j. This axiom also implies that, given two actions $a_1 = \{c_{11} \mid E_1, \ldots, c_{1m} \mid E_m\}$ and $a_2 = \{c_{21} \mid E_1, \ldots, c_{2m} \mid E_m\}$, if $c_{1j} \ge c_{2j}$ for all j, then $a_1 \ge a_2$. Moreover, if $c_{1j} > c_{2j}$ for some j, then $a_1 > a_2$.

Axiom 4 (Reference events). The decision-maker can conceive of a procedure to generate a random point in the unit square such that, given two regions R_1 and R_2 in the unit square, the event that the random point is contained in R_1 is regarded as more likely than the event that the random point is contained in R_2 if, and only if, the area of R_1 is larger than the area of R_2.

Discussion. This axiom defines a "unit of measurement" which can be used to describe the decision-maker's uncertainty concerning the occurrence of the various events which are relevant for the decision problem.

Axioms 1-3 are qualitative. They establish the rules that comparison of actions must obey. Axiom 4 is of a quantitative nature and allows the decision-maker to measure her degree of belief concerning uncertain events in terms of probability.

The axioms imply a three-step procedure to choose among actions. First, all preferences between consequences must be quantified in terms of a numerical utility (or, alternatively, a loss) function. Second, the uncertainty regarding any event affecting the consequences of an action must be quantified in terms of probability. Finally, the decision-maker's preferences about any pair of alternative actions must be described in units of expected utility and hence the best decision will be to choose the action that maximizes such expected utility.

In brief, the main idea behind this formulation is this: if the coherence axioms are acceptable for a decision-maker, then she has no other option but to use the maximum expected utility criterion to choose among actions. Any other method leads to the same optimal decisions or conflicts with the axioms. Stating the axioms explicitly allows anyone to evaluate if she is willing to act as a coherent decision-maker defined in this particular way. In addition, if the axioms are adopted, the existence of a probability function for the uncertain events, and a utility function for the consequences, is no longer an assumption. It is a fact, following from the axioms themselves.

This is a "subjective" method since the utility and probability functions, must describe the preferences and beliefs of the specific decision-maker. Interestingly, the subjective nature of the utility function has not been a matter of controversy. On the other hand, the idea of subjective probabilities led to a lively debate in the early years and even now appears as a relevant issue in some applications (see Schneider 2002, for an example).

When the decision-maker is faced with new information concerning the uncertain events, her beliefs must be revised in order to obtain an updated (posterior) probability measure describing all the available knowledge (which includes the original, prior beliefs as well as the new information). Another important consequence of the axioms is that this updating must be consistent with Bayes' rule.

When dealing with decisions under uncertainty, one cannot guarantee that the optimal action is a good decision in the sense that it will necessarily lead to the best possible consequence. Apart from guaranteeing compliance with the coherence axioms, the maximum expected utility method only leads to the action whose "expected" consequence is the most preferable.

Example (ctd.). To improve the results of the admission process to the Ph.D. program, the Graduate Studies Committee can ask the candidates to take a selection exam. The corresponding grade of the exam (X) is then used to update the probabilities $P(E_1)$ and $P(E_2)$. If the observed grade is x, then the new probabilities are defined as $P(E_1 \mid X = x)$ and $P(E_2 \mid X = x)$ and can be obtained via Bayes' formula as $P(E_i \mid X = x) = P(X = x \mid E_i)\, P(E_i) / P(X = x)$; $i = 1,2$.

If these updated probabilities are used to make a decision, a_1 must be chosen if, and only if, $P(E_1 \mid X = x) / P(E_2 \mid X = x) > l(c_{12}) / l(c_{21})$. Equivalently, a_1 is the best action if $P(X = x \mid E_1) / P(X = x \mid E_2) > (p_2 / p_1)(l(c_{12}) / l(c_{21}))$. Here, statistical knowledge is useful to propose the models $P(X = x \mid E_1)$ and $P(X = x \mid E_2)$ describing how likely a grade is under each of the scenarios E_1 and E_2. The decision is thus based on the likelihood ratio $P(X = x \mid E_1) / P(X = x \mid E_2)$. In this case, the cutoff value depends on both the ratio of the probabilities p_1 and p_2 and the relative importance of the possible mistakes. An interesting issue closely related to this problem is that of the "test design". If the classification exam is properly designed, it would be reasonable to expect high grades among those candidates who later prove to be good students and lower grades among bad ones. If the classification exam works well as a screening tool, the ratio $P(X = x \mid E_1) / P(X = x \mid E_2)$ should be a nondecreasing function of x and then $P(X = x \mid E_1) / P(X = x \mid E_2)$ will be larger than $(p_2 / p_1)(l(c_{12}) / l(c_{21}))$ only if x is larger than a cutoff given by $h((p_2 / p_1)(l(c_{12}) / l(c_{21}))$ where the function h depends on how well the exam discriminates between good and bad candidates, a characteristic that can be analyzed through statistical methods. This "decision rule" leads to the expected result: candidates with a high grade are admitted whereas those with a low grade are rejected. The more subtle issue is related to the cutoff value. How high should a grade be for a candidate to be admitted? Under this decision-theoretical framework, it is clear that the best rule is not necessarily to admit a candidate only if she passes the exam. Depending on the combination of the "prior" probabilities p_2 and p_1, losses $l(c_{12})$ and $l(c_{21})$, as well as on the function h, it might be optimal to admit some candidates who do not pass the exam or to reject some who pass the exam with a relatively low grade. This procedure proposes a solution for the choice of the cutoff value in a straightforward manner which makes it clear what the components of the problem are and how these should be taken into account.

3. Some Applications.

The literature offers a variety of applications of decision theory in education, mostly in the field of education research. In particular, van der Linden (1991,1997) provides a comprehensive review of applications to test theory. More recently, Vos (1999) and Vos and Glas (2000) propose a sequential decision procedure to develop adaptive mastery tests, whereas Segall (2004) applies ideas from decision theory to evaluate a new sharing item response theory model; see also van der Linden (1998). A common feature of these contributions (as well as many others of this type) is that the role of decision maker is played by the researcher, the instructor, or even the computer system used to administer the test.

From a different perspective, but also in the field of test design, Bernardo (1998) discusses several grading procedures for the case of multiple-choice examinations. There, the student, as the decision-maker, must devise a strategy to answer each item of the examination. In this context, however, the corresponding decision problem is defined by the instructor, who determines the set of possible answers *A*, the set of uncertain events *E,* the consequences set *C*, and, very importantly, the utility function (the grading score mechanism). If the student is a coherent decision-maker in the sense of Section 2.3, she must choose the answer maximizing her expected utility. This expectation is computed by combining the grading score function of the test and the subjective probability distribution describing the student's knowledge.

Bernardo shows how the utility function can be defined in order to discourage 'guessing'. He also discusses the cases where the aim of the student is either to get the highest mark or just to pass the examination. In addition, and elaborating on an original idea by de Finetti (1965), he explores the results obtained when the student is asked not only to mark the correct option but to report her entire probability distribution over the set of possible answers for each item of the test.

4. Current Developments.

Decision Theory is an active field of research. The list of problems under current investigation include: experimental studies of individual and group behavior; use of different methods to understand human judgments and decisions; discussion of alternative normative models; and applications of the theory to a wide variety of subjects.

Specifically, alternative sets of axioms are being explored to produce a structure able to encompass the widest range of real problems; Lipman (1999) and Dubois et al. (2003) are two examples of this type of investigation. Another line of research has to do with elicitation of both utilities and probabilities. Clemen and Reilly (1999) discuss the use of copulas for this purpose whereas Bleichrodt and Pinto (2000) explore a specific procedure in the context of medical decisions. In addition, for some decision problems maximization of the expected utility can only be accomplished through numerical or simulation methods (see Bielza et al. (1999), for instance). Beyond the limits of the theory described in this article, there are researchers who assume that consequences might include several facets, and preferences for each one of them are described through a different utility function. This framework is known as a multi-objective decision problem and there is no generally accepted solution in such a setting (see Dyer et al. (1992)). Another problem whose general solution is yet to be formulated is that of group decision. If the decision-maker is a committee where each member has a different utility function or a different probability distribution, how should these be combined to choose the action with the largest expected utility? Will this combination be an optimal solution in practice? Weerahandi and Zidek (1981) propose a solution for this problem; Hollenbeck et al. (1998) takes into account the structure of the group through multilevel theory.

Concerning Bayesian Statistics, the statistical ramification of Decision Theory, current research also includes alternative axiomatic formulations (see Karni (2007), for a recent example), elicitation techniques (Garthwaite et al., (2005)) and applications in an ever increasing number of fields. There is also a strong research effort directed towards statistical computing. Specifically, simulation techniques for the calculation of posterior distributions using Markov Chain Monte Carlo (MCMC) methods have been successfully employed over the last 15 years (Bhattacharya and Haslett (2007) provide an example).

References.

Bernardo, J.M. and Smith, A.F.M. (1994). *Bayesian theory*. Chichester: Wiley.

Bernardo, J.M. (1998). A decision analysis approach to multiple-choice examinations. In Girón, J. (ed.) *Applied Decision Analysis*. Boston: Kluwer.

Bhattacharya, S. and Haslett, J. (2007). Importance re-sampling MCMC for cross-validation in inverse problems. *Bayesian Analysis* **2**, 385-408.

Bielza, C., Mueller, P. and Rios Insua, D. (1999). Monte Carlo methods for decision analysis with applications to influence diagrams. *Management Science* **45**, 995-1007.

Bleichrodt, H. and Pinto, J.L. (2000). A parameter-free elicitation of the probability weighting function in medical decision analysis. *Management Science* **46**, 1485-1496.

Clemen, R.T. and Reilly, T. (1999). Correlations and copulas for decision and risk analysis. *Management Science* **45**, 208-224.

de Finetti, B. (1965). Methods for discriminating levels of partial knowledge concerning a test item. *British Journal of Mathematical and Statistical Psychology* **18**, 87-123.

Dubois, D., Fargier, H. and Perny, P. (2003). Qualitative decision theory with preference relations and comparative uncertainty: An axiomatic approach. *Artificial Intelligence* **148**, 219-260.

Dyer, J.S., Fishburn, P.C., Steuer, R.E., Wallenius, J. and Zionts, S. (1992). Multiple criteria decision making, multiattribute utility theory: the next ten years. *Management Science* **38**, 645-654.

Garthwaite, P.H., B. Kadane, J.B. and O'Hagan, A. (2005). Statistical methods for eliciting probability distributions. *Journal of the American Statistical Association* **100**, 680-700.

Hollenbeck, J.R. Ilgen, D.R. LePine, J.A., Coquitt, J.A. and Hedlund, J. (1998). Extending the multilevel theory of team decision making: Effects of feedback and experience in hierarchical teams. *The Academy of Management Journal* **41**, 269-282.

Lindley, D.V. (1972). *Bayesian Statistics, a review.* Philadelphia, PA: SIAM.

Lipman, B.L. (1999). Decision theory without logical omniscience: Toward an axiomatic framework for bounded rationality. *Review of Economic Studies* **66**, 339-361.

Karni, E. (2007). Foundations of Bayesian theory. *Journal of Economic Theory* **132**, 167-188.

Novick, M.R. and Lindley, D.V. (1978). The use of more realistic utlity functions in educational applications. *Journal of Educational Measurement* **15**, 181-191.

Sawyer, R. (1996). Decision theory models for validating course placement tests. *Journal of Educational Measurement* **33**, 271-290.

Schneider, S.H. (2002). Can we estimate the likelihood of climatic changes at 2100? *Climatic Change* **52**, 441–451.

Segall, D.O. (2004). A sharing item response theory model for computerized adaptive testing. *Journal of Educational and Behavioral Statistics* **29**, 439-460.

van der Linden, W.J. (1991). Applications of decision theory to test-based decision making. In: Hambleton, R.K. and Zaal, J.N. (eds.) *Advances in educational and psychological testing*, pp. 129-156. Boston: Kluwer.

van der Linden, W. J. (1997). Decision theory in educational testing. In Keeves, J.P. (ed.), *Educational research, methodology, and measurement: An international handbook* (2nd ed.), pp. 725-736. Oxford: Elsevier Science.

van der Linden. W. J. (1998). A decision theory model for course placement. *Journal of Educational and Behavioral Statistics* **23**, 18-34.

Vos, H.J. (1999). Applications of Bayesian decision theory to sequential mastery testing. *Journal of Educational and Behavioral Statistics* **24**, 271-292.

Vos, H.J. and Glas, C.A.W. (2000). Testlet-based adaptive mastery testing. In van der Linden, W.J. and Glass C.A.W. (eds.). *Computerized adaptive testing: Theory and practice*, pp. 289-310. Boston: Kluwer.

Weerahandi, S. and Zidek, J.V. (1981). Multi-Bayesian statistical decision theory. *Journal of the Royal Statistical Society* A, **144**, 85-93.

Further Reading.

Bernardo, J.M. and Smith, A.F.M. (1994). *Bayesian theory*. Chichester: Wiley.

Rao. C.R. and Sinharay S. (eds.) (2007). *Handbook of Statistics* (Vol. 26: Psychometrics), Amsterdam: North-Holland.

Lindley, D.V. (1972). *Bayesian Statistics, a review.* Philadelphia, PA: SIAM.