DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

POLICY REVERSALS IN RISK MANAGEMENT: THE EFFECT OF REFINED ANALYSES

Matthias Hild
California Institute of Technology &
Jet Propulsion Laboratory

SOCIAL SCIENCE WORKING PAPER 1116

March 2001

Policy Reversals in Risk Management: The Effect of Refined Analyses

Matthias Hild

Abstract

Reversals of policy recommendations occur in risk management when the social decision maker aggregates individuals' subjective utilities for the outcomes of a risky policy measure. The level of detail with which these outcomes are described can significantly affect the resulting policy recommendation. The choice of the level of detail on which we conduct an analysis therefore amounts to an implicit value judgement. Moreover, the power to fix the level of the analysis implies partial control over policy recommendations. We propose an alternative approach to decision–theoretically sound risk management.

Key words: group choice, risk management, social decision making

Policy Reversals in Risk Management: The Effect of Refined Analyses

Matthias Hild*

1 Introduction

The decision–theoretic approach to risk management structures the process of social decision making into two carefully separated parts and then uses concepts from decision analysis to arrive at a policy recommendation (Raiffa, 1968). In the first part, the highly professionalized discipline of probabilistic risk assessment procures probabilities for social decision making (Henley/Kumamoto, 1992). In the second part, a social utility function is constructed. We focus on applications, such as health care economics or environmental cost–benefit analysis, where a social utility is obtained on the basis of individuals' subjective utilities for the outcomes of a risky policy measure. In these applications, social utility is typically defined as the average of the individuals' utilities (Haddix et al., 1996). Decision–theoretic risk management then proceeds by maximizing the expected social utility relative to the estimated probabilities.

We show that this structure renders risk management decisions sensitive to a factor that has not been explicitly acknowledged in the literature. This factor is the degree of detail with which we describe the outcomes of a policy measure. Consider the decision problem of a group of directors who consider building a production plant in Europe. Suppose some member of the panel assigns a positive value to owning a plant in Europe. We can then make the individual's reasons for this evaluation more transparent by giving a finer description of this outcome. Consider only the two possibilities of, firstly, owning a production plant in Europe in a climate of an increasing Euro/USD exchange rate and, secondly, owning a production plant in Europe in a climate of a non–increasing exchange rate. The main phenomenon that drives our argument can now be stated easily: Different individuals may agree in their evaluation of a coarsely described outcome and yet disagree in their reasons for this evaluation. One individual may regard a rise in the exchange rate likely but only slightly beneficial, while another individual may regard a rise in the exchange rate unlikely but highly beneficial. The individuals' opposing reasons may thus lead to the same overall evaluation of the situation. Note that we have not made any

*Address for correspondence: California Institute of Technology, Mailcode 228–77, Pasadena, CA 91125 (USA). Email: matthias@hild.org. Website: http://www.hild.org.

new information available to the individuals but have only refined our description of the individuals' evaluations.

We shall study the effect that refinements of an analysis have on risk management decisions. Due to the structure of the decision–theoretic risk management process, this effect can in theory be dramatic and is in practice non–negligible. It can lead to an instability and a reversal of the policy recommendations. On a first level of description, policy a may be strictly preferred to policy b while on a second, more detailed, level of description b is strictly preferred to b. There is nothing in principle to stop this phenomenon to occur repeatedly. On a third, yet more detailed, level of description, policy a may again be strictly preferred to policy b and so forth. Throughout, the inclusion of additional parameters leaves the individuals' expected utilities unchanged. In other words, the refinement is irrelevant from the viewpoint of the individuals' decision making. We would have wished that such a refinement should also be irrelevant from the viewpoint of social decision making but, unfortunately, it is not. In statistics, similar phenomena are more familiar. Their theoretical possibility was first discussed by Simpson (1951) (hence the name 'Simpson's paradox') who pointed out that the level of the analysis can influence the conclusions about statistical correlations in the data. In a famous study on gender bias in the admissions process of Berkeley University, Bickel/Hammel/O'Connell (1977) found that an analysis of the overall admissions figures suggested a gender bias against women, while a more fine–grained analysis on a departmental basis found either no such trend or, in some departments, even a reverse trend. (For a more recent example from epidemiology, cf. Reintjes et al., 2000).

We give several illustrations of such phenomena in risk management. We then argue that there are no objective criteria for determining on which level a risk analysis should be conducted. This simple observation has two important implications. On the one hand, we can draw a methodological lesson for risk management: The choice of a relevant level of detail reflects a value judgement that ought to be make explicit in the analysis. On the other hand, we can draw a political lesson: The power to select a level of detail implies influence and partial control over the resulting policy recommendation. In all illustrations, we assume for simplicity that individuals and society maximize expected utility. A general statement of our result shows that these phenomena are not the artifact of a particular decision theory but a consequence of the very structure of the risk management process. After we expose the philosophical assumptions underlying the current practice of risk management, we suggest a concrete and practical approach to decision–theoretically sound risk management.

2 The Phenomenon

The structure of the risk management process

We focus on applications in which the construction of a social utility depends directly on individuals' subjective utilities for the possible outcomes of a policy measure. In such applications, social utility is typically obtained by averaging the individuals' subjective utilities, although the aggregation method may take a more complicated form. Given some method of aggregating individual utilities into a social utility, the structure of the risk management process can be studied abstractly. The theoretical literature on social choice refers to processes of this general type as 'ex post social choice rules' (Hammond, 1981). Rules of this type factor individuals' subjective evaluations of risky alternatives, i.e., individuals' expected utilities, into two separate components. These components are, on the one hand, the individuals' subjective probabilities and, on the other hand, their subjective utilities. These utilities for the outcomes of a risky policy measure are referred to as 'ex post' utilities. Processes of this class can be contrasted with 'ex ante social choice rules' that start from a different basis for social policy recommendations. The starting point of ex ante rules are not factorized probabilities and utilities but the compound evaluation that integrates these factors into the expected utility of a risky alternative. These expected utilities are also called 'ex ante' utilities. Examples of ex ante social choice rules are ubiquitous. Majority voting in a referendum or a committee is one example where individuals express ex ante evaluations of risky alternatives by casting their votes. Auctions for telecommunication bandwidths are another example where individuals express their ex ante evaluations by bidding with money. We emphasize that our phenomenon does not occur within processes of the ex ante type. Hence, the phenomenon is clearly a result of the factorization of ex ante evaluations into probabilities and ex post utilities.

The motivation for a factorization of ex ante evaluations into probabilities and (ex post) utilities is fundamentally a philosophical one. According to a common normative argument, risk managers ought to separate questions of fact (probabilities) and questions of value (utilities). The distinction between facts and values was most famously drawn by Hume (1739–40) and has been a central concept of ethics ever since (Hare, 1981).

Hume's maxim: Factorize beliefs and values!

Implicit is the view that beliefs enjoy a special status vis-a-vis values that allows us to apply the norms of rational discourse and of scientific method to disagreeing beliefs. The U.S. Nuclear Regulatory Commission (1975) was the first to put this conviction into practice, thereby creating the discipline of probabilistic risk assessment (PRA).

PRA's maxim: Be rational about probabilities!

Disagreements about values, on the other hand, pose a more difficult problem. It is commonly agreed that values can be the object of a rational analysis only in as far as they draw on factual assumptions. To use an example from environmental cost–benefit analysis, consider the decision to build a hydroelectric power station that runs the risk of cutting off the water supply of an extensive wooded area downstream. Environmental cost–benefit analysis routinely elicits the public's value for keeping some woodland intact. Members of the public can assign different values to the preservation of the forrest, for instance, because they hold different beliefs about the effects of acid rain. The factorization of the disagreement into a factual and an evaluative part is then intended to help arbitrate disagreements in a rational manner. PRA analyses the possible effects of acid rain, while any remaining disagreements about the utility of environmental preservation are settled by averaging individuals' subjective values (at least in the applications that we will consider). PRA's maxim is often overlaid by a somewhat paternalistic categorization of those topics of which the lay public is deemed a competent judge. Matters of risk and uncertainty are typically the reserve of expert professionals, while the evaluation of the ex post outcomes of a risky policy is often the privilege of the lay public.

Our punch line will be that the factorization of beliefs and values is an iterative, fractal–like process. What is a value on one level of analysis is a compound of beliefs and values on another level of analysis. Moreover, the recommendations obtained from decision–theoretic risk management are sensitive to how exactly compounded evaluations are factorized into beliefs and values. We must acknowledge that, within the current structure of the decision–theoretic risk management process, a value judgement is required to determine what counts as a relevant level of analysis or, equivalently, what counts as a relevant factorization of expected utilities into probabilities and utilities. As a solution to this problem, we will propose a concrete and practical process that is true to PRA's maxim but abandons Hume's maxim.

Example: Quality of Life

Quality Adjusted Live Years (QUALYs) are regularly estimated by presenting individuals with choices among lotteries over outcomes affecting the individuals' quality of life. In health care, for example, one wishes to estimate an individual's utility for an outcome C described in medical terms. The standard approach chooses the hypothetical medical states of perfect Health and certain Death as reference points in the estimation process.[1] The individual is then offered a hypothetical choice between the following options:

Option 1: Health with probability p and Death with probability $1 - p$.

Option 2: Outcome C with certainty.

For the sake of the argument, we accept the decision–theoretic assumptions embodied in this methodology (von Neumann/Morgenstern, 1947). We shall also not touch on the problem of hypothetical choices. If we then normalize the utility of Health to be 100 and

[1] The choice of Health and Death as reference points leads to an implicit normalization and interpersonal comparison of individual utilities.

the utility of Death to be 0, we obtain $u_i(C) = 100 \cdot p$ whenever individual i becomes indifferent between Option 2 and Option 1.

Quality of life measures can be applied to the problem of allocating donor kidneys to patients with kidney failure. The problem of kidney allocation is an example of risk management because the process of allocation involves several sources of uncertainty. For instance, kidney failures and therefore the arrival of patients on the waiting list form a stochastic process, just like the availability of donor organs and the success of transplantations (Hild, 2001e). Alternative allocation mechanisms therefore correspond to risky policy measures. In evaluating alternative allocation mechanisms, Zenios/Wein/Chertow (1999), for instance, use a quality of life measure (which in their case seems to be chosen purely ad hoc). We will argue that the use of such measures harbours an important problem that needs to be addressed by an explicit choice of a relevant level for describing the outcomes of the policy. Let C be the outcome that the patient receives a donor kidney for transplantation. Our inquiry starts from the trivial observation that the description of the outcome C allows several refinements and several interpretations because several additional attributes are left implicit in the outcome C. For instance, both donor organs and patients come in different medical types that vary in their compatibility. Currently, compatibility is measured in terms of HLA antigen matches (EBPG, 2000). This measure offers a good (although not optimal) proxy for the expected survival time of the kidney in the recipient's organism (Wujciak/Opelz, 1993). Moreover, donor organs are received only after a waiting period of variable length due to short supply. In the coarse–grained description C these attributes are not mentioned. Consider the following refined descriptions:

C_1: receiving a transplant with a good fit after a short waiting time.

C_2: receiving a transplant with a good fit after a long waiting time.

C_3: receiving a transplant with a poor fit after a short waiting time.

C_4: receiving a transplant with a poor fit after a long waiting time.

Individuals may now differ both in their probabilities and utilities for these additional attributes. We fix our time horizon at 1 year. For simplicity, we consider only two types of individuals that occur in equal proportions within the population. These types differ in their preferences for a good fit and a short waiting time. Individuals of type 1 put more emphasis on a good fit, whereas individuals of type 2 put more emphasis on a short waiting time. Type 1 refuses to accept an organ with a poor fit and, hence, assigns zero probability to outcomes C_3 and C_4. There is a 50–50 chance that an individual of type 1 will have to wait only for a short period of time. Type 2 has a strong preference for a short waiting time and therefore accepts any organ offered. For simplicity, we assume that type 2 has a zero probability for a long waiting time (outcomes C_1 and C_3). There is a 50–50 chance that an individual of type 2 will receive an organ with a good fit.

	transplant	Expected utility
p_1:	1	
u_1:	70	70
p_2:	1	
u_2:	70	70
p_0:	1	
u_0:	70	70

	good fit, short wait	good fit, long wait	poor fit, short wait	poor fit, long wait	Expected utility
p'_1:	.5	.5	0	0	
u'_1:	60	80	5	20	70
p'_2:	0	.5	0	.5	
u'_2:	20	80	5	60	70
p'_0:	.25	.5	0	.25	
u'_0:	40	80	5	40	60

Table 1: Kidney allocation on two levels of analysis.

The upper part of Table 1 shows the individuals' utilities u_i for receiving outcome C with certainty. Stacking the cards against our phenomenon, we assume that the individuals agree on these utilities and thus arrive at an identical social utility u_0 for receiving outcome C with certainty. On this level, individual and social probabilities p_1, p_2, p_0 are degenerate. The lower part of the same table shows the subjective probabilities p'_i and utilities u'_i of individuals in these types. The table also shows the social probabilities p'_0 and the social utilities u'_0.

The individuals' probabilities and utilities are chosen such that the refined description of the outcomes is fully compatible with the coarser description C. In other words, the refinement is irrelevant from the viewpoint of the individuals' decision making. Their preferences have not changed as a result of the refinement since their utility for the coarsely describe outcome C already incorporates their more fine–grained beliefs and utilities. In particular, the individuals have not received any new information. Formally, we have $u_i(C) = p'_i(C_1) \cdot u'_i(C_1) + \ldots + p'_i(C_4) \cdot u'_i(C_4)$ for each type i. Hence, in the above lottery, either type of individual would become indifferent between Option 1 and Option 2 at a p–value of 70% on either level of analysis.

We assume that the individuals' subjective probabilities are indeed the correct objective probabilities. Our argument does not hinge upon an assumption that the individuals are misinformed. We arrive at a social utility by averaging the individuals' subjective utilities. Our phenomenon occurs when we now calculate the social utility for the outcome C on different levels of description. On the first level, we obtain a social utility of

$u_0(C) = 70$. On the second level, we obtain a social utility of $u'_0(C) = 60$ (although the individuals' expected utilities remain unchanged). It is time to emphasize that the risk management techniques under consideration are of a quantitative nature and, moreover, highly sensitive to differences in quantitative assumptions. Hence, a difference of about 14% in the social utility can quickly lead to different qualitative recommendations. If we were interested in dramatization, we could easily conjure up numbers that lead to more extreme discrepancies (Hild, 2001a).

We may pursue the process of refinement yet further and include additional personal factors specific to each individual. An individual's evaluations of 'long' waiting times may again be compounded from beliefs and values. The individuals may, for instance, have probabilities for the event that their family becomes dependent on their physical support (e.g., if a child or a partner falls ill). A patient's priority for a short waiting time might then be explained by her high probability of her husband suffering a stroke. Depending on their circumstances, different patients may also have different utilities for being able to help needy family members (e.g., if the patient is a single mother of a small child or a grandmother in a family of 10). With this refined analysis, we again open up the possibility of changing social choices. On each level of refinement, we have to decide whether the new factorization of beliefs and values is relevant for the social choice problem. The inclusion of personal factors such as the family's special needs may or may not be judged relevant.

Our observation is not that the social decision maker needs to decide which parameters of a choice problem are relevant from a social perspective. Rather, the social decision maker needs to decide what is the appropriate form of factorizing individuals' beliefs and values. If QUALY measurements depend on how finely we describe the medical conditions of interest, an obvious idea suggests itself: Perhaps we should use finest description possible. We will discuss such suggestions in the next section. For now, we note that the level of analysis does, indeed, need fixing, by one method or another.

The Worldbank devised a measure of Disability Adjusted Life Years based on the evaluation of medical conditions by a panel of medical experts (cf. Worldbank, 1993 and Murray/Lopez/Jamison, 1994). This construction of a social utility fixes the level of description for the different medical conditions once and for all and thus avoids our problem. Unfortunately, the Worldbank developed this analysis for estimating the global distribution of diseases. For our problem of kidney allocation, this approach is of no help because events like the allocation of a kidney for transplantation are beyond the resolution of the Worldbank's analysis.

Example: Technological risk

We consider the same phenomenon in the perhaps more familiar context of technological risk management. We consider a study of solar energy as an alternative to nuclear energy. We have two alternatives:

C: use of nuclear energy in status quo.

D: develop solar energy.

In order to estimate the social utility of C, we might consider the individuals' willingness to pay for retaining the status quo. In other words, since it is very costly at the present time to replace nuclear energy by solar energy, we wish to estimate up to what cost individuals would be prepared to go along with a change in energy policy. Techniques based on willingness to pay are widely used in environmental cost–benefit analysis. Again, we take no issue with the problems inherent in this methodology since our purpose is to illustrate an unrelated point. We consider refined descriptions of outcome C. An obvious detail to include is the risk of a nuclear accident on the scale, say, of the Chernobyl accident and within the time horizon, say, 50 years.

C_1: status quo without serious accident.

C_2: status quo with some serious accident.

We again turn our attention to two types of individuals that occur with equal proportions within the populations. Let us assume that both types assign the same value to outcome C, namely $u_i(C) = 10$. Hence, social utility also equals 10 (i.e., $u_0(C) = 10$). These values are shown in the upper part of Table 2. (On the coarse–grained level, all probabilities are, again, degenerate.) Differences emerge only on the more fine–grained level. Individual of type 1 believe that nuclear accidents are unlikely but disastrous in their consequences. Individual of type 2 believe that nuclear accidents are more likely but localized and moderate in their consequences. Type 1 represents the 'old school' of nuclear risk management. The 'new school' of type 2 believes in non–linear model and a low response of cancer rates in adults to temporary radiation leakage (e.g., Jaworowski, 1999, and the International Atomic Energy Agency, 1996). In our example, both schools of thought arrive at the very same expected utilities of 10 for the benefits of nuclear energy in the status quo (shown in the lower part of Table 2). Moreover, individuals have complete agreement about the utility of the status quo without an accident (i.e., $u_1'(C_2) = u_2'(C_2) = 20$).

It is, of course, not surprising that expected social utility may differ from the individuals' expected utilities. What is surprising is that we again observe a change in social utility across these two levels of analysis. Note that we have again assumed that the correct social probability is the average of individual probabilities. This is not a necessary assumption and our general result allows social probabilities provided by experts to disagree strongly with the average layman's probability. This possibility only exacerbates our problem. By assuming that social probability is the average of individual probabilities, we have constrained our example in favour of a stable social choice because we have precluded the introduction of new information through the experts' analysis. Nonetheless, we were not able to guarantee stability. This shows that the effect is not due to an increase in information about matters of fact but only due to the way in which

	nuclear energy		Expected utility
p_1:	1		
u_1:	10		10
p_2:	1		
u_2:	10		10
p_0:	1		
u_0:	10		10
	accident	**no accident**	**Expected utility**
p'_1:	.1	.9	
u'_1:	−80	20	10
p'_2:	.2	.8	
u'_2:	−30	20	10
p'_0:	.15	.85	
u'_0:	−55	20	8.75

Table 2: Nuclear energy on two levels of analysis.

we factorize individuals' evaluations. Our general result will show that instabilities are completely independent of the manner in which social probabilities are constructed.[2]

A general result

On theoretical grounds, instability phenomena are an inherent feature of the way in which the decision–theoretic risk management process is structured. To arrive at general theorem, we need very few assumptions about the model that captures the individuals' decision making.[3] We make no assumptions whatsoever about the manner in which social probabilities are constructed. We thus allow the individuals' probabilities to be replaced by some experts' risk assessment for the purpose of social decision making. We shall assume that a social utility is constructed by averaging individuals' utilities. We make

[2] Again, we could consider a second refinement that analyzes whether the low–response hypothesis is true or false. This description would include details about the groups of individuals, such as children and adults, that are most affected by a nuclear accident. The International Atomic Energy Agency (1996) notes that children suffer proportionately more than adults. Even individuals who agree on their probability for the low–response hypothesis could therefore still disagree in their utility for equal and unequal distribution of negative effects among children and adults.

[3] Individual decision making may be modelled by any of the following theories: expected utility theory, expected utility theory with threshold, expected utility theory with non–additive probabilities (Choquet–expected utility theory), rank dependent utility theory, prospect theory, weighted utility theory, Machina's theory, regret theory or any ordinal decision theory based on 'precautionary principles' like leximin.

this unnecessarily strong assumption only to simplify the presentation. For technical details, we refer the reader to Hild (2001a).

Theorem 2.1 *If a social utility is constructed by averaging individuals' utilities, then there is an infinite sequence of increasingly refined descriptions of the outcomes of a risky policy such that the risk management process leads to contradictory policy recommendations ad infinitum although the individuals' beliefs and utilities are fully compatible across all levels of description.*

The proof of this theorem turns on the simple observation that individuals who have identical preferences for 'different reasons'. More precisely, the identity of two individuals' expected utilities does not imply the identity of either their probabilities or their utilities. One individual might judge a certain consequence of an act unlikely but highly desirable, another might judge the consequence more likely but less desirable, and yet both may agree on their overall evaluation of the act.

3 Evaluation

Minimal recommendations

We have found that policy choices may depend on the level on which a risk management study is conducted and on the manner in which individuals' evaluations are factorized into beliefs and values. If we must use the current structure of the decision–theoretic risk management process, then we should at least keep the following minimal recommendations in mind.

- Make the choice of the level of analysis explicit in study.
- Consider and discuss alternative choices.

Incidentally, since the current practice of risk management does not vary the level of analysis within a study, we cannot offer any concrete data that display our effect of policy reversals. The data that would allow us to locate such effects have simply not been collected.

We return to the idea to choose as finely–grained an analysis as possible:

- Let choice be guided by feasibility considerations, such as the degree of detail for which can we provide data.

There are several difficulties with this idea. Firstly, there is, of course, no reason to believe that an ultimate, maximally fine–grained, description of reality exists. More importantly, there is a second, epistemic, problem. Even if a level of analysis existed beyond which no instabilities of social choice occurred, we could never know whether our analysis has reached such a critical level. Yet, there might be a practical end to the refinability of our analysis. Beyond a certain level of detail, we may simply run out of statistical data for a probabilistic risk assessment. This is perhaps a stop–gap solution to our problem, but it is not fully satisfactory. Since the choice of a level of analysis can significantly influence the recommendation of a policy, there is room for political manipulation through a clever choice of a favourable level of detail. This problem is mitigated only by the difficulty to foresee what level of detail will yield which social choice. Nonetheless, we should take note of this political context.

- Contain the power to choose the level of analysis.
- Let choice be made by a disinterested party.

A decision–theoretically sound alternative

The current practice of decision–theoretic risk management factorizes social decision making into the two separate parts of probabilistic risk assessment and the construction of a social utility. This very structure leads to potential instabilities in policy recommendations. Alternative social choice rules such as majority voting in committees or the auctioning of telecommunications bandwidths do not have such defects. In this type of process, no instabilities occur (Hild, 2001a). We now propose an approach to risk management that retains the virtues of the decision–theoretic approach and is, at the same time, not culpable of its vices. The current decision–theoretic risk management process has two main attractions. Firstly, it professes to be decision–theoretically sound and, secondly, it embodies PRA's maxim to be rational about probabilities. We share the appreciation of these two properties. We object, however, to the implicit acceptance of Hume's maxim to factorize individuals' ex ante evaluations into beliefs and utilities. This assumption is an unnecessary addition to the first two properties.

This philosophical change of heart implies clear changes for the practice of risk management. We uphold PRA in order to let scientific method and rational discourse bear on social choice. But, instead of adopting the experts' risk assessment as a social probability, we communicate the results of PRA to the individuals concerned. Individuals are then given the opportunity to update their beliefs with this new information and thus to arrive at a better informed appreciation of the risk involved. This use of PRA puts a great responsibility on the communication of risks. Fortunately, the actual practice of risk communication and experimental studies in psychology are providing us with a growing understanding of how best to communicate risk. Assuming that probabilistic information is communicated in a psychologically tractable format, social choice then proceeds in an ex ante fashion along the familiar lines of majority votes, auctions etc. The important contribution of PRA is to base risk management choices on the ex ante evaluations of better informed individuals. To coin a term, we call this feedback process *interactive risk management*. We summarize the steps of this process:

Step 1: Probabilistic risk assessment.

Step 2: Communicate PRA to individuals concerned.

Step 3: Individuals update their beliefs.

Step 4: Social choice by individuals' ex ante evaluations.

We note with satisfaction that interactive risk management also avoids the many empirical problems with the elicitation of individual utilities. Ex post aggregation is intrinsically bound to the use of detailed quantitative information about the individuals' ex post utilities. In the previous section, we have accepted without questions that such information can be reliably obtained. In fact, this is an extremely problematic (that is to say, false) assumption. Many empirical studies of individuals' decision making show that they systematically deviate from the theoretical assumptions that underly any methodology for measuring their ex post utilities.

Interactive aggregation is most attractive for small expert–like groups of individuals. Consider the situation of a board of directors of a company that has to decide on a corporate strategy. Studies are conducted trying to gain insights into the probable effects of a strategy. The results of these PRA studies are communicated to the members of the board. Remaining uncertainties and disagreeing values are discussed by the board and, finally, a group choice is reached by voting or informal mitigation processes. This process can be applied directly to the *Jet Propulsion Laboratory*'s objective of selecting a landing site for a future Mars mission (cf. http://www.jpl.nasa.gov). Different scientific interests typically favour the selection of different landing sites. An additional factor in the management's choice of a landing site is the probable success of the mission. With missions to Mars, the current landing technology is a source of considerable uncertainty that lead to a loss of the space probe, especially in rugged terrain. The challenge is to optimize the expected value of the mission. Some of the suggested solutions to this problem follow the ex post mould of decision–theoretic risk management and are therefore sensitive to the level of detail on which the analysis is conducted (Miles, 2001). Interactive ex ante aggregation takes the following approach to the problem: A PRA study of the landing phase is conducted and its results are then communicated to the different science teams. After the teams have updated their beliefs with this information, a managerial decision is made on the basis of the scientists' updated ex ante evaluations of the alternative landing sites. This decision can be achieved through informal means, through votes or the design of an auction (Ledyard/Porter/Wessen, 2000). Our current research adapts existing work on auction theory to the specific context of risk management.

Another place where this approach is put into practice is our own proposal to solve the kidney allocation problem that was presented in the previous section (Hild, 2001e). In this proposal, the ex ante evaluation of the uncertainties arising from the allocation process is the privilege of the patient. We communicate to the patient in a psychologically manageable format all available information about the expected survival times of organs depending on the donor's and the patient's medical types. We also offer active decision support and make tools for decision analysis available to the patient. We then choose an ex ante allocation mechanism that avoids any form of dependence on the description of medical outcomes. These concrete and practical implications show that it can at times be well worth questioning one's philosophical assumptions.

Acknowledgements: I wish to thank Ralph Miles for his help.

References

[1] P.J. Bickel, E.A. Hammel, and J.W. O'Connell. Sex bias in graduate admissions: Data from Berkeley. In W.B. Fairley and F. Mosteller, editors, *Statistics and Public Policy*, pages 113–130. Addison–Wesley, Reading (Mass.), 1977.

[2] EBPG Expert Group on Renal Transplantation. European best practice guidelines for renal transplantation (part 1). *Nephrology, Dialysis, Transplantation*, 15, Supplement 7:1–83, 2000.

[3] Henley E.J. and Kumamoto H. *Probabilistic Risk Assessment: Reliability Engineering, Design, and Analysis.* Institute of Electrical and Electronics Engineers, Inc., Piscataway, 1992.

[4] A.C. Haddix, S.M. Teutsch, P.A. Shaffer, and D.O. Dunet, editors. *Prevention Effectiveness: A Guide to Decision Analysis and Economic Evaluation*, Oxford, 1996. OUP.

[5] P. Hammond. Ex–ante and ex–post welfare optimality under uncertainty. *Economica*, 48:235–250, 1981.

[6] R.M. Hare. *Moral Thinking: Its Levels, Methods and Point.* Clarendon, Oxford, 1981.

[7] M. Hild. Stable aggregation of preferences. *California Institute of Technology, Social Science Working Paper*, 1112, 2001a. Available at www.hild.org.

[8] M. Hild. Fair kidney allocation based on waiting time. *California Institute of Technology, Social Science Working Paper*, 1117, 2001e. Available at www.hild.org.

[9] D. Hume. *A Treatise of Human Nature*. 1739–40. Modern edition by E.C. Mossner, London: Penguin, 1984.

[10] International Atomic Energy Agency. One decade after Chernobyl. *www.iaea.org*, 1996.

[11] Z. Jaworowski. Radiation risk and ethics. *Physics Today*, 52:24–29, 1999.

[12] J. Ledyard, D. Porter, and R. Wessen. A market-based mechanism for allocating space shuttle secondary payload priority. *Experimental Economics*, forthcoming, 2000.

[13] R. Miles. A decision–theoretic approach to risk–adjusted mission value. unpublished, 2001.

[14] C.J.L. Murray, A.D. Lopes, and D.T. Jamison. The global burden of disease in 1990 — summary results, sensitivity analysis and future directions. *Bull. World Health Organ.*, 72:495–509, 1994.

[15] H. Raiffa. *Decision Analysis*. Addison–Wesley, Reading, 1968.

[16] R. Reintjes, A. De Boer, W. Van Pelt, and J. Mintjes de Groot. Simpson's paradox: An example from hospital epidemiology. *Epidemiology*, 11:81–83, 2000.

[17] E.H. Simpson. The interpretation of interaction in contingency tables. *Journal of the Royal Statistical Society, Ser. B*, 13:238–241, 1951.

[18] G. Opelz T. Wujciak. Computer analysis of cadaver kidney allocation procedures. *Transplantation*, 55:516–521, 1993.

[19] U.S. Nuclear Regulatory Commission. *Reactor Safety Study : An Assessment of Accident Risks in U.S. Commercial Nuclear Power Plants*. Nuclear Regulatory Commission, Washington, D.C., 1975.

[20] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, 1947.

[21] Worldbank. *World Development Report*. OUP, Oxford, 1993.

[22] S.A. Zenios, L.M. Wein, and G.M. Chertow. Evidence-based organ allocation. *American Journal of Medicine*, 107:52–61, 1999.