

Risk management and integrated assessment

Carlo C. Jaeger

*Darmstadt University and Swiss Federal Institute for Environmental Science and Technology, EAWAG, Ueberlandstraße 133,
CH-8600 Duebendorf, Switzerland*

Received 3 February 1997; revised 2 February 1998

Integrated Assessment (IA) is the pursuit of a research program generated by the limitations of traditional forms of risk management. This claim can be justified by the following argument. Over the last decades, analysts and practitioners have brought to maturity a large array of tools for risk management. Most of them rely on combining judgments of utility with judgments of probability. This is the approach of the Rational Actor Paradigm (RAP). With many environmental problems, however, RAP-based tools have run into considerable practical and theoretical difficulties. In response to these difficulties, a series of alternative approaches to practical risk management and to the theoretical understanding of risk have been elaborated. They try to embed the rational choices of individual actors studied by RAP into a broader framework of social rationality. This task can be approached by distinguishing situations where an actor holds unambiguous judgments of preference and probability from situations characterized by ambivalent judgments. RAP can handle the former, but not the latter. Problems whose management requires a combination of widely differing scientific disciplines are especially likely to involve ambivalent judgments of probability. The study of such problems constitutes the research program of integrated assessment. It involves three main tasks: developing IA models which can represent ambivalent expectations and evaluations, developing IA models which use such representations to study non-marginal changes of social systems, and developing procedures of participatory IA which enable researchers to engage in an iterative exchange with various stakeholders.

Keywords: risk management, rational choice, integrated assessment

1. Introduction

Risk management is a well established practice in a variety of fields. Here we are interested chiefly in environmental risks. The management of these risks can draw on a wealth of experience in other areas. The most important ones are technology, insurance, and finance. Many successful practices in these areas of risk management can be understood in terms of RAP, the rational actor paradigm [19]. (A comprehensive treatment of the literature on risk analysis and on integrated assessment lies beyond the scope of the present paper. The references should help, however, to connect the argument of the paper to that literature.)

The emerging practice of integrated assessment might be seen as just another application of RAP to environmental problems requiring close cooperation between natural and social scientists. In fact, however, IA blends components from RAP-based risk management – like optimization schemes to obtain cost effective paths towards some desirable state – with approaches which deliberately move beyond the limits of RAP – like the use of different utility functions depending on different cultural perspectives which might be adopted by decision makers. Under these conditions, IA may develop in one of three ways. First, it could be reintegrated in the family of RAP-based procedures of risk management, with the more exotic components gradually being either dismissed or reintegrated in the established paradigm. Second, it could evolve as a conceptually incoherent mixture of approaches which coexist more or less peacefully in practice. And third, it could help

to design and realize a research program which gradually embeds RAP into a broader understanding of human choice in an uncertain world. The present paper explores the third alternative.

For this purpose, we will discuss how typical practices of risk management share an emphasis on evaluation (judgments of utility) and expectation (judgments of probability) with their combination into decision rules for risky situations (section 2). We will then analyze how these features are sustained by RAP, the rational actor paradigm, in a powerful theoretical scheme (section 3). Next, our discussion will highlight several practical and theoretical difficulties which the paradigm has met and which are relevant for environmental management (section 4). These difficulties have led to a series of participatory procedures for discursive conflict resolution (section 5). The way they work cannot be understood without embedding RAP into a broader theoretical framework, and this in turn has important consequences for the design of IA models (section 6). At a practical level, the field of IA holds promise to blend IA models and participatory procedures into forms of institutional mediation which could complement traditional forms of risk management in areas where these are deficient (section 7).

2. Probability, utility, and risk management

Contemporary risk management comprises a great variety of tools. Nevertheless, most of them are linked by strong family resemblances. Three of these are especially

pervasive: judgments of probability play a prominent role, so do judgments of utility, and these two kinds of judgments are combined to yield decision rules in the face of risk. In the present section, we will look at these features in three areas of risk management: technology, insurance, and finance.

With technological risks, the typical problem is a situation where some serious damage occurs with statistical regularity, e.g., fatal accidents in mining. The frequency of damage events can then be described in terms of objective probabilities. The probability of damage will often result from the interaction of various processes, each with its own probability of failure. The task of risk management is to lower these probabilities to some acceptable level. That level will usually depend on the financial costs of damage events. These costs express the utility of avoiding the damage in question.

Risk management of this sort depends on the combination of two kinds of judgment. On the one hand, there are judgments of probability. Here empirical observations, statistical analyses and scientific inquiries of varying sophistication come into play. They may yield not just one judgment of probability, but a whole system of such judgments, specifying how various hazards and stochastic processes interact. On the other hand, there are judgments of utility. In the case under consideration, more often than not they will refer to monetary costs. Nevertheless, purely physical indicators like numbers of casualties may be relevant, too. Judgments of probability and of utility are then combined into some decision rule. For example, a specific case of accident, say train derailments because of overheating axes, is to be kept below a certain level of probability that seems both morally justifiable and economically affordable. As a consequence, choices of materials, technologies, monitoring systems and use patterns will be shaped by this decision rule.

With insurable risks, there are two typical problems. One is the problem of how much to insure. The greater the insurance coverage, the smaller the remaining risks. As insurance is costly, however, one will keep some risks uninsured. Where to lay the threshold will depend on the insurance premium and on one's risk aversion. Risk aversion can be defined on the basis of purely subjective probabilities. If a person with given capabilities and habits has no statistical knowledge about the likelihood of fatal accidents for persons of her kind, she may still have rough subjective judgments about the likelihood of such accidents. Her degree of risk aversion may then be described as the loss of utility which she associates with the risk of incurring these accidents. Again, we are faced with the combination of judgments of probability and judgments of utility. Now, however, these judgments may have a very strong subjective component. Together, judgments of probability and of utility shape the propensity to take specific actions like buying an insurance. More generally, subjective probabilities can be inferred from the willingness of an actor to bet on certain outcomes of a lottery. And an important

way of betting is to take a specific action under risky conditions.

The concepts of probability and utility are also useful to frame the second insurance problem: what premium should an insurer charge for a given insurance? Given the huge uncertainties involved in actual insurance markets, this will depend on the insurer's guesses about possible revenues and costs. Demand for insurance together with competition by other insurers will shape the possible revenues at different levels of the premium. As for expected costs, for every level of the premium they must somehow represent an aggregate of possible events. Therefore, they depend on the insurer's expectation of damage events together with her own degree of risk aversion. The first may be described by objective or subjective probabilities, depending on the kind of information available, the latter may again be described by subjective utilities. In the case of the insurer, they will indicate, among other things, her willingness to risk bankruptcy. On the basis of these guesses, the insurer can then try to set the premium at a level which maximizes the difference between revenue and costs.

With financial risks, a typical problem is the investment of a sum of money in different assets, like shares of various firms and bonds of various governments. Each asset may be characterized by some expected return, μ , and a standard deviation, σ , from this return. Actors in financial markets are usually risk averse. Therefore, if an asset has lower μ than some other asset, it will be interesting only if it has a smaller σ . As a result, the investor is faced by a tradeoff between profitability and security. She then needs to define some level of σ which indicates her threshold of acceptable risk. However, this will not lead simply to choose assets characterized with this σ . An optimal portfolio will be diversified between various assets so that the acceptable level of risk is reached for the portfolio as a whole. This can be done if the correlation between the risks of various assets can be estimated. One will then try to avoid positive risk correlation. In this case, judgments of utility are judgments about rates of return and acceptable risk, while judgments of probability refer to expected values and standard deviations.

A more subtle problem is presented by sequential investment under uncertainty. Investments differ in the flexibility with which they can be modified. A bank account may yield a predefined interest with the option of withdrawing one's money at any time; building a production plant may yield an uncertain return while the option for withdrawal is much more limited. If uncertainty can be reduced thanks to new information in the course of time, keeping one's options open has a positive value. Again, judgments of probability and judgments of utility interact so as to yield decision rules for the problem at hand. The examples could be multiplied, but the list given so far makes it clear enough that there is a common conceptual pattern involved in a large variety of tools for risk management.

3. Risk management and RAP: The Rational Actor Paradigm

The judgments of utility which we have encountered again and again in the examples of risk management discussed in the preceding section point to the close links between risk management and the idea of a rational actor. In its basic form, this idea involves a well-defined set of alternatives and a preference ordering between these alternatives. Rational actors select alternatives according to their given preference orderings. This idea lies at the core of RAP, the Rational Actor Paradigm. While the idea looks simple and familiar, a closer look is in order. Environmental modeling and assessment requires intense cooperation between researchers with widely differing backgrounds. Under such conditions, there is a serious danger to delude ourselves that we share an understanding of important theoretical concepts, only to enmesh ourselves in jungles of misunderstandings and conceptual confusion thereafter. Avoiding this trap is all the more important as I want to argue that the limitations of RAP encountered in risk management lie at the origins of the research program of integrated assessment.

At the outset, we need to recognize that the strength of RAP does not lie in its description of an isolated actor. Quite the opposite: RAP provides an amazingly simple description of individual actors in order to develop sophisticated explanations and predictions for complex systems of such actors. The most important achievement in this regard no doubt is the analysis and modeling of interdependent competitive markets. Whether some model of decision making really represents an improvement over the basic RAP model, then, cannot be assessed without asking how the alternative model fares in studying isolated and interconnected markets as well as other social systems.

Usually, the preference orderings used by RAP can be represented by functions from the set of alternatives to the set of real numbers. Any monotonic transformation of such a utility function will be equally suitable to represent a given preference ordering. The rational actor is then bound to maximize her utility function over the set of alternatives. RAP is especially valuable when the set of alternatives can be described as a metric space in which several dimensions are connected by a constraint. This is the case, for example, when a limited resource has to be assigned to various activities which can be performed at different levels. (In such a space, alternatives with equal utility can be represented by indifference curves.) Increasing the utility from some use of a limited resource is then bound to decrease the utility from some other use. If a reference point in the set of alternatives is given, any change can then be analyzed in three steps (figure 1). First, only the dimensions in which the change in question leads to increasing utility are considered. The total utility increase from these dimensions is the benefit of the change in question. Second, the dimensions with decreasing utility are considered. The total decrease along these dimensions is the cost of the change.

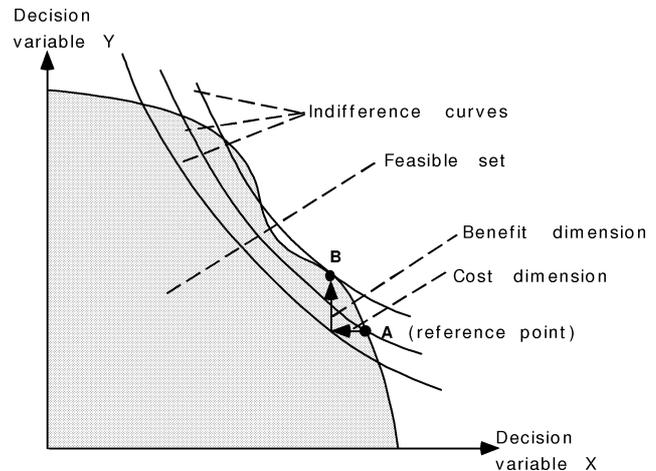


Figure 1. Benefits and costs of departing from a reference point.

Finally, benefit and cost are compared. A change where cost exceeds benefit makes no sense. On the other hand, if benefit exceeds cost, there may still be another change which would yield even better results.

This exercise lies at the basis of benefit–cost analysis. It does not require monetary evaluation, because it will yield the same result with different utility functions as long as these are linked by monotonical transformations. If the reference point is a maximum of a differentiable utility function (as with point B in figure 1), then for small changes benefit will be equal to cost. For slightly larger changes, cost will exceed benefit. As figure 1 also shows, this does not exclude the often neglected possibility that even larger changes lead to another utility maximum.

Each alternative under consideration may consist of an action and a set of possible but uncertain consequences. The distinction between actions and consequences leads to a whole range of new problems. Sometimes an action may lead to a unique consequence. A choice involving only such actions may be described as a choice under conditions of certainty. Other actions may lead to several possible consequences. In the literature, these actions are sometimes described as lotteries [28]. Choices involving at least one such action may be described as choices under conditions of uncertainty.

Now the question arises of how preferences between actions with unique consequences are related to preferences between actions with several uncertain consequences. It seems natural to assume that this depends on the perceived likelihood of the consequences. After all, if several lotteries involve the same set of possibilities, the only plausible reason to prefer one lottery to another is the belief that they differ in the likelihood of consequences among which the relevant actor is not indifferent.

This line of reasoning has led to the idea that the utility of a lottery should be computed as its expected utility, i.e., the sum of the utilities of the various consequences weighted with the probabilities of these consequences. In figure 2, a lottery is represented by points X1 and X2 on a continuum of possible outcomes whose utility in-

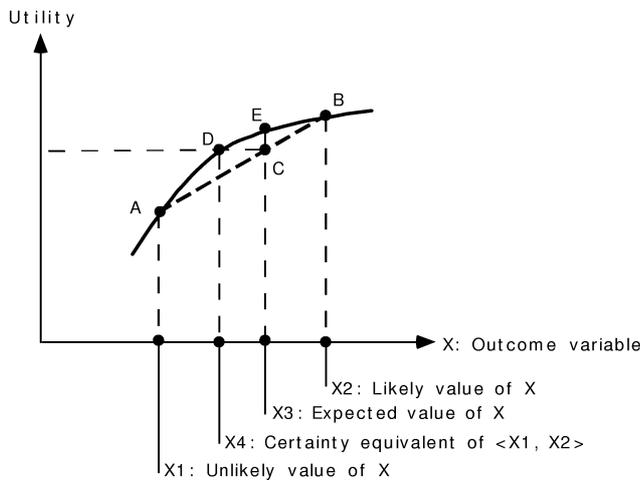


Figure 2. Risk aversion and cardinal utility.

creases with their size. I may judge that X2 is more likely than X1. My subjective expected value of the lottery is then a weighted average of X1 and X2, with the subjective probabilities as weights (point X3 in figure 2). If I am risk averse, my certainty equivalent of the lottery will be some point X4 which is smaller than the expected value. The utility of the lottery is equal to the height of C, the average of the utilities of X1 and X2 weighted with their probabilities. Now a utility function can be constructed in such a way that the utility of the lottery, read off at C, is equal to the utility of the certainty equivalent, read off at D, but smaller than the utility of the certain event X3 (point E).

The idea that utility functions are equivalent as long as they result from each other by monotonical transformations can now no longer be maintained. An explicit representation of uncertainty leads to cardinal utilities, i.e., utility functions for which only linear transformations are admissible. The shape of these functions becomes a representation of risk aversion. For a risk neutral actor, the utility function is a straight line. The widespread practice of using monetary magnitudes like GNP as measures of utility, therefore, misses two important features of most risk issues: the fact that many important actors are risk averse with regard to the relevant risks, and the fact that different actors usually have very different degrees of risk aversion with regard to the same risks.

So far, we have looked at the combination of utility and probability as introduced by RAP for the study of an isolated actor. The same combination can be used to study risk problems involving a conflict of interest between several parties. Such situations, which arise very often with environmental problems, can be studied with metaphors provided by game theory. While the metaphor of a lottery is used by RAP to analyze choices under conditions of uncertainty, the metaphor of a game applies RAP to choices made by interdependent agents. The consequences of an action now depend on the choices made by other agents.

Table 1 represents the skeleton of simple games involving two players each of whom can perform two actions. As a result of the combined action of the two players, one of

Table 1
Skeleton of simple games.

	Column's first action	Column's second action
Row's first action	α	β
Row's second action	γ	δ

Table 2
Utility functions for two simple games.

	α	β	γ	δ
U_{row}	-10	0	7	-10
U_{column}	-10	0	7	-10
U_{row}	10	0	7	-10
U_{column}	-10	0	7	10

four outcomes will result. Each player defines a strategy by assigning probabilities to his actions, while at the same time guessing at the probabilities of the opponent's actions. The probabilities may be one or zero, or any combination in between.

It is important to distinguish between outcomes and payoffs (which is why tables 1 and 2 differ). Outcomes may be any kind of events resulting from a combination of actions by the different players. In particular, outcomes may be risks, or lotteries in the terminology introduced above. Payoffs are utility values assigned to each outcome by a cardinal utility function which is specific for each player. The cardinality of the utility function is used to represent the risk aversion of the different players according to the method described above. In the world of RAP, a rational actor will now maximize her expected utility under the condition that her opponent will be equally rational. Therefore, each actor assigns probabilities to his own actions as well as to the actions of the opponent in such a way that both maximize their expected utility. Any situation meeting this condition is an equilibrium of the game in question.

Table 2 describes two different games by assigning utility payoffs to the skeleton provided by table 1. The first two rows describe a game involving multiple equilibria: β and γ designate two different situations where both actors maximize their expected utility under the condition that the opponent does the same. Clearly, equilibrium γ is superior to equilibrium β . However, rational actors operating in a RAP world may settle in equilibrium β , too: if historical accident or whatever reason has led them into that equilibrium, individual rationality will keep them there.

The last two rows of table 2 describe a slightly different game: two out of eight utilities have been increased with the result that γ is no longer an equilibrium. Whether this is due to different outcomes or to different evaluations of the same outcomes is irrelevant here. Now a social dilemma of the kind known as the prisoners' dilemma or, with more players, the tragedy of the commons arises. These are situations where rational actors are trapped in a sub-optimal outcome, in this case outcome β , definitely. They both have reasons to think about ways of achieving outcome γ ,

but with the game as it stands there is little they could do about it.

Social dilemmas are a major problem in risk management. For example, stalemate in a conflict about siting a landfill for toxic waste may lead to the waste staying in unsuitable places for long times. In principle, however, a social dilemma can always be eliminated by introducing an additional player with a suitable influence on the payoff structure of the other ones. With regard to risk management, three additional players are especially relevant: the market, government, and science. It may seem odd to consider the market as a player on a par with businesses, consumers, and other agents. Nevertheless, these latter agents experience the market as an independent force with which they are faced, while the market in turn is influenced by their actions. Therefore, the market can be and often is modeled as if its operations were driven by an agent interacting with the players involved.

For our present purposes, there is no need to look at the detailed representation of these very different agents in a game-theoretic format. Suffice it to stress a few salient features. Markets may assign prices to various goods so as to achieve an equilibrium between supply and demand. By so doing, they may avoid a social dilemma if they burden the actions which lead into the dilemma with additional costs. Government may reach the same outcome by legally constraining the choices of other actors, e.g., with taxes or with the establishment of new property rights. Finally, science may provide additional information which helps other actors to recognize and to avoid specific risks. Sometimes, this appears as a displacement of ignorance by knowledge, in other cases as a more gradual learning process of the kind first described by Bayes [5].

From a RAP perspective, then, risk management appears as a two-layered task. At one level, an isolated actor may try to cope rationally with some risk. Here RAP-based risk management draws on the metaphor of a lottery in order to organize judgments of utility and of probability and to combine them so as to reach a sound decision. This happens differently in different areas, but RAP still provides the common background for a wide array of tools and methods. At another level, interdependent actors try to handle a risk situation they share. In particular, they may struggle over the size and kind of risks they want to incur. Moreover, they may struggle over how to distribute the – positive and negative – consequences of incurring the risk in question. Here RAP-based risk management draws on the metaphor of a game, once more in order to organize judgments of utility and of probability and to combine them so as to reach sound decisions.

4. Limitations of RAP-based risk management

As we have seen, current practices of risk management have reached a high level of sophistication in fields as diverse as financial markets, technological risks, and insur-

ance for hazards ranging from illness to natural catastrophes. Moreover, RAP provides a common theoretical framework for these fields by combining utility analysis with probability theory. On this basis, lotteries and games provide powerful metaphors for the study of human decision making under conditions of risk and uncertainty.

While RAP-based risk management has enjoyed many impressive successes, it has also run into several practical difficulties. In particular, expert knowledge on probabilities and on the valuation of outcomes has been challenged in a wide variety of settings. This has led to substantial delays for risky projects, if not to insurmountable public opposition. Moreover, governmental decisions in risk management are often faced by considerable levels of mistrust, which may translate into serious difficulties for the implementation of those decisions. More generally, quite often lack of trust is a major obstacle to successful risk management, and RAP seems to be of limited help in practical attempts to establish trust relations. This is especially relevant in cases where scientific expertise is used by government to justify decisions in risk management. Finally, businesses involved in risk management cannot take much trust for granted, either. Toxic wastes, airplane crashes, even financial turmoil are seen by many as testimony to the danger that profit maximizing firms tend to confront human communities with unacceptably high levels of risk.

These difficulties are related to the possibility of overcoming social dilemmas with the help of an additional actor, be it science, government, or the invisible hand of the market. In contemporary society, this possibility is more limited than one might think. First, science is gradually losing its role as impartial provider of objective knowledge. Many risk controversies are characterized by the availability of contradictory expert knowledge. The case of nuclear waste is especially telling. Obviously, researchers have established many facts in this field – including the ones used for engineering nuclear reactions – which nobody would dare to contest. But crucial safety issues depend on views about which there is no strong consensus among experts. Similar situations arise in local controversies about the pros and cons for specific measures of risk management. Technically speaking, science is increasingly dealing with problems which would be completely untractable without drawing on subjective probabilities by the researchers involved [34]. As powerful stakeholders as well as the general public come to realize this state of affairs, they reduce their willingness to modify their own subjective probabilities in reaction to statements by scientific experts.

Second, government as we know it does not and cannot guarantee the realization of “social welfare” in dealing with environmental and other risks. Obviously, there are many instances where government may reasonably be said to foster a robust version of social welfare. But few doubt the fact that often governmental bureaucracies have interests of their own, and that they can be quite ruthless in pursuing them. Moreover, even if some bureaucracy honestly wanted to realize the public good, it would quite often find itself

in serious difficulties to define in operationally meaningful terms what this notion is supposed to designate. Divergent interests are a fundamental feature of pluralist societies, and votes and elections are a way of coping with this divergence, not of bringing them to an end. This means that even if government should take a clear and well-founded decision on some major risk issue, chances are that it will be hard, sometimes even impossible to realize it.

Third, markets have serious difficulties in dealing with long term risks because existing futures markets cover only a very restricted subset of the future possibilities we are aware of. In practice, futures markets exist mainly for certain financial assets and for commodities, and in both cases they cover a time-span of years, not of centuries. Risks like those of nuclear waste, of climate change, of genetic engineering, of toxic chemicals, etc. are simply not dealt with by futures markets of the appropriate range. This would necessitate futures markets for land, water, seeds, infrastructure, labor, and many other goods and services with a time horizon of several generations. It is hard to imagine how a workable system of such futures markets could ever emerge. The practical consequence of lacking futures markets is straightforward: the amazing efficiency of markets in coordinating human actions cannot be drawn upon to manage long-term risks.

Faced with these practical difficulties, RAP-based risk management is also plagued by serious theoretical problems. In fact, the extension of RAP to decision making with explicit uncertainties, pioneered by von Neumann and Morgenstern [36] and by Savage [49], led to a series of paradoxical results. Perfectly reasonable people, including researchers committed to RAP, were shown to act in ways which are incompatible with expected utility models. On these grounds, Allais [2] presented a powerful critique of what he called “the American school”, namely expected utility theory. The editor of “Econometrica” added an interesting note to his paper [2, p. 503]: “The problem discussed in Professor Allais’ paper is of an extremely subtle sort and it seems to be difficult to reach a general agreement on the main points at issue. I had a vivid impression of these difficulties at the Paris colloquium in May, 1952. One evening when a small number of the prominent contributors to this field of study found themselves gathered around a table under the most pleasant exterior circumstances, it even proved to be quite a bit of a task to clear up in a satisfactory way misunderstandings in the course of the conversation.” Allais’ basic argument was that reasonable people care about the dispersion of actual probability distributions in ways which cannot be handled by expected utility models. In particular, the question whether the probability of some very risky event is equal to 0 (the event in question being strictly impossible) or just slightly positive (the event being a real possibility, even if a very unlikely one) deserves special attention. In actual decision processes, this question plays a completely different role than the question whether some probability is 37.8 or 37.9%. In some cases people will treat probabilities close to zero as basically equivalent to

zero – as with the probability of train accidents – in other cases – as with failure probabilities of nuclear reactors – people will care very much about such differences.

Several related paradoxes have been published in the literature. In particular, Ellsberg [13] showed that most people prefer objective probabilities to subjective ones. For example, they prefer a lottery involving two outcomes with objective probabilities of 50% (as when the outcome depends on tossing a coin) to a lottery involving the same outcomes with subjective probabilities of 50% (as when no information about how the final outcome will be selected is available).

As a reaction to these results, the axioms leading to subjective expected utility models have been weakened so as to accommodate several of these paradoxes [14,32]. However, the seminal work of Kahneman and Tversky [25] has shown that the choices of reasonable and highly skilled people may depend critically on whether a problem is framed as one of avoiding losses or one of achieving gains. This holds with identical outcomes and identical probabilities. Therefore, it is hard to see how this result could ever be accommodated within RAP.

The practical and theoretical difficulties of RAP-based risk management have not led and will hardly ever lead to the demise of the management tools which have been elaborated in a RAP framework. But these difficulties cannot be accommodated by some minor amendments of RAP either. As we will see in the next section, they have already induced a whole series of fruitful practical innovations in risk management which in important ways go beyond RAP.

5. Alternative approaches to risk management

The difficulties of RAP-based risk management have led to a series of practical innovations for participatory risk management [45]. A common trait of these innovations is the use of discursive procedures for the purpose of conflict resolution. Important examples include Citizen Advisory Committees [18,31] as well as comparative risk projects [33], which are presently used in a variety of settings in the US. They also include various mediation procedures based on patterns of formalized negotiation [59]. For the present argument, it is important to check whether and how these innovations go beyond the framework of RAP. For this purpose, a short outline of some of the relevant methods is useful.

As a first example, consider the method of the “consensus conference” [22]. The method has been developed in the US, in Denmark, The Netherlands and the UK, mainly for purposes of technology assessment. Suppose a controversial topic of risk assessment arises with regard to some new technology. A steering committee of about half a dozen people representing a wide variety of stakeholders may then start a consensus conference. This requires a budget of about \$150,000 and a time frame of about one year for the whole process. A professional facilitator is

hired and administrative support made available. Advertisements in the mass media call for laypersons interested in the process to submit a written application. Out of the people applying, a panel of ten to twenty people is selected according to the following criteria: the panel should be roughly representative of the general population, it should not include experts on the topic under scrutiny and it should be interested in social and ethical dimensions of science and technology. Moreover, the steering committee prepares a list of about thirty experts in the field of interest. The panel meets for preparatory weekends where it is briefed on the general subject of the consensus conference and where it decides what questions it wants to pursue and which experts it wants to consult. These decisions are taken unanimously; additional experts may be named by the panel. The experts provide written answers to the questions submitted by the panel. The panel then meets again for the consensus conference proper, which lasts several days and is open to the general public. The conference starts with a hearing of the experts, then the panel retires and writes an opinion. The opinion states whatever conclusions the panel is able to reach unanimously. The conclusions may include open questions as well as policy recommendations. The report is then presented at a media conference.

A similar procedure is the “planning cell” [10]. Here about two dozen citizens chosen at random work for one to three weeks to produce an opinion about a controversial planning issue. The planning cell is usually called into being by some public authority which will make sure that the citizens in question obtain a leave from their jobs or other duties and obtain a financial reward for their efforts. A planning cell requires a budget of about \$30,000. Two facilitators direct the planning cell with the support of two or more experts in the field under study. Several planning cells may work in parallel on the same problem. Especially in Germany and Spain, planning cells have dealt with many local and sometimes national policy issues.

A simpler, but related tool is the “future workshop” [23]. Ten to thirty people meet for three days (or more) to discuss a planning problem. The workshop goes through three phases: criticism, imagination, implementation. Knowledge is generated from the future workshop itself under the guidance of a facilitator. In the last two decades a large number of future workshops has been run, especially in Germany, Switzerland and Austria.

The procedure of “rational discourse” [44] is another tool for cooperative risk management. Groups of laypersons are informed about a set of policy options. The group members may be chosen at random from the general population, but they may also be found by asking a large array of associations to send delegates. The participants are rewarded to avoid social bias and lack of commitment. Each group is directed by a facilitator and meets several times over a period of at least some months. The discussion groups are asked to define criteria for judging the policy options under consideration as well as weights for the different criteria. In a second stage, they assess the possible consequences of the

different options. For this purpose, they perform hearings with experts of their choice. Finally, they combine their assessment of possible consequences with their criteria to reach a consensus on a policy option which they are willing to support.

These examples may be sufficient to convey an impression of how practitioners in risk management are currently trying to deal with the limitations of RAP. It is useful to relate them to two management tools used in other areas: project groups and focus groups. The former are used in managing rapid organizational change. They are often deliberately formed so as to cut across established hierarchies. In the context of a well delimited task of limited duration they have far-reaching powers of inquiry. In particular, such a group may hold hearings with people who according to the standard organizational hierarchy are more senior than members of the project group. As for focus groups, they are widely used in marketing and in public opinion research [30]. The basic idea is to gather a group of people and expose them to some shared input, like a TV speech by a politician or information about a new product. The group is then given time to digest, as it were, the input in a common discussion. While individual interviews are useful to document individual attitudes and opinion, focus groups are useful to document social processes of attitude and opinion formation. This is especially relevant to assess the acceptability of new products and new policy options.

A common feature of these procedures is that they create the opportunity for a debate in which the participants can try to persuade each other of their respective judgments. To some extent, this may simply be a matter of negotiation in the narrow sense of bargaining over possible deals. This is a kind of interaction which fits nicely with RAP. But the debate may go well beyond that and lead to arguments about judgments of probability and preferability. As a result, the participants may reach common judgments which they did not share to begin with. This is a process of social learning which is hard to reconcile with a RAP framework.

6. Towards an understanding of social rationality

Recently, attempts to overcome the limitations faced by RAP in risk management have led to promising work on social theories of risk [24,29,55]. Important contributions range from the notion of social amplification of risk to the cultural theory of risk, from the use of the arena metaphor in risk analysis to the philosophy of post-normal science [15]. A common trait of these theoretical efforts is the attempt to embed the individual rationality considered by RAP into an understanding of rationality as a social phenomenon. For this attempt to be successful, it will be important to develop it in the realm of mathematical models elaborated by RAP. These models have become indispensable for a worldwide set of activities in research, policy advice and management practice. If these activities are to take the embeddedness of rational choices by isolated individuals

Table 3
Utility functions for a controversial game.

	α	β	γ	δ
Utilities of row player for consequences of outcome lotteries	10, -100	10, -10	10, -10	5, -200
Utilities of column player for consequences of outcome lotteries	5, -200	10, -10	10, -10	10, -100
Probabilities I for consequences of outcome lotteries	99%, 1%	50%, 50%	50%, 50%	90%, 10%
Probabilities II for consequences of outcome lotteries	90%, 10%	50%, 50%	50%, 50%	99%, 1%
Utilities of row player for outcome lotteries with probabilities I	$9.9 - 1 = \mathbf{8.9}$	$5 - 5 = \mathbf{0}$	$5 - 5 = \mathbf{0}$	$4.5 - 20 = \mathbf{-15.5}$
Utilities of column player for outcome lotteries with probabilities II	$4.5 - 20 = \mathbf{-15.5}$	$5 - 5 = \mathbf{0}$	$5 - 5 = \mathbf{0}$	$9.9 - 1 = \mathbf{8.9}$
Utilities of row player for outcome lotteries with probabilities II	$9 - 10 = \mathbf{-1}$	$5 - 5 = \mathbf{0}$	$5 - 5 = \mathbf{0}$	$4.95 - 2 = \mathbf{2.95}$
Utilities of column player for outcome lotteries with probabilities I	$4.95 - 2 = \mathbf{2.95}$	$5 - 5 = \mathbf{0}$	$5 - 5 = \mathbf{0}$	$9 - 10 = \mathbf{-1}$

into a broader framework of social rationality into account, RAP-based models need to be improved and generalized so as to overcome their current limitations.

This task constitutes a research program which can only be realized by a scientific community engaging in a long-term effort. In order to give a first outline of this research program, it is useful to consider a stylized version of a typical problem in risk management (table 3). Two parties struggle about how to deal with a risky choice, say, whether to build a specific landfill for toxic waste. Each outcome of their combined actions may be represented by a lottery. For example, outcome α may designate the case where the two parties agree to build the landfill; this may either lead to a positive result or to a highly undesirable accident. The two parties may differ, however, in how they evaluate the two possible consequences of this outcome, and they may differ even more in their judgment about the likelihood of these consequences. For example, the row player may represent the proponents of the landfill, the first possible consequence of outcome α may represent a regular operation of the landfill, and the second a possible accident. The column player, representing the opponents, considers the same possibilities, but associates them with different utilities.

To keep the numbers simple, no attempt has been made to provide empirically grounded figures for the various outcomes in table 3. One may imagine that outcome δ represents the case where both parties agree not to build the landfill. The first possible consequence may be a corresponding reduction in solid waste, the second consequence may be pollution from excess waste. Outcomes β and γ may represent different on-going quarrels about the landfill, each with its own uncertain outcomes.

The parties may well be insecure about their judgments of likelihood. Such a situation is excluded by the standard assumptions of RAP. In reality, however, ambivalent expectations are an extremely important feature of many social interactions. In table 3, they are represented by distinguishing two different sets of subjective probabilities (I and II), each of which may be assigned to each player. The proponents of the landfill may tend to think that the possibility of an accident is quite remote, but from time to time they may have doubts. So, they will in general hold probabilities I, but sometimes consider probabilities II, too. The converse holds for the opponents. Even from the highly arbitrary numbers provided it should be obvious how to represent various outcomes of a social conflict about risk with plausible numbers for utilities and probabilities.

We call such a game a controversial game. Both parties may now try to convince the other that a certain set of probabilities is the correct one. At the same time, they may use this controversy to make up their own mind about what they really want to believe. As table 3 shows, whether outcome α (building the landfill) will be an equilibrium of the resulting game depends on how these processes evolve. If the proponents manage to convince themselves and the opponents of probabilities I, building the landfill will be an equilibrium. Otherwise, other outcomes will become equilibria: if both players accept probabilities II, it will be δ ; and β or γ may become equilibria if the players do not stick to the same probabilities.

So far we have considered a situation where actors are ambivalent about their judgments of likelihood. Technically, one possibility to represent such situations is given by the method of imprecise probabilities, involving not a single probability but two probabilities designating an up-

per and a lower bound of the relevant range [43,54]. But once we acknowledge the importance of ambivalent expectations there is no reason to restrict them to probability judgments: judgments of preference may be ambivalent, too [52]. Technically, a notion of imprecise utilities, again characterized by an upper and a lower bound, may help to represent such situations.

The crucial point about situations of ambivalence is that here the different agents do not simply try to choose optimal strategies on the basis of their expectations and evaluations. Rather, they define their evaluations, expectations, and strategies in a shared social process. In fact, this may well mean that the identity of the agents is not given independently from their interactions (see [56] for a comprehensive analysis of this possibility).

Where ambivalent judgments of probability and utility arise, stakeholders may engage in fruitful arguments about these ambivalences, arguments which would be futile in their absence. How is a reasonable argument about ambivalent utilities and probabilities possible? RAP-based research usually assumes that judgments of utility lie beyond the realm of reason, while arguments about probabilities are restricted to Bayesian learning which combines prior probabilities with knowledge about empirical facts (see [6] for a recent discussion). Neither is of much use for arguing about ambivalent utilities and probabilities. To understand such arguments, a different approach is needed.

We consider the case of utility first. Here, it is a standard assumption that preference orderings can be represented by functions from the space of conceivable commodity bundles to the set of real numbers. People who orient themselves according to such preference orderings may be said to follow their self-interest, and they may be said to maximize their utility. This kind of preferences is closely associated with experiences of pain and pleasure. Such experiences are shared by humans and animals, and they may be seen as lying at the biological roots of preference orderings. We call this kind of preferences hedonic preferences. Hedonic preferences may be considered as brute facts, without a necessity or even a possibility of further justification. In such a view, human rationality shrinks to a remarkably simple affair. To be rational means, first, to display a preference ordering which fulfills some requirements of consistency and, second, to follow one's preference ordering in actual decisions.

Hedonic preferences, then, are not very instructive for arguments about ambivalent preferences. Ethical judgments are much more relevant here: ambivalence is very frequent in ethical matters and arguing about ethical judgments is a widespread practice. How then are ethical judgments related to preferences? They do generate preference orderings: an action may be judged preferable to another one on ethical grounds. Ethical judgment can be criticized and defended, and human beings spend a great part of their life in doing this. Hedonic preferences may be understood without asking how people account for them in social settings; after all, animals have hedonic preferences, too. Ethical prefer-

ences, however, cannot be understood in such a way. When it comes to ethical issues, accounting for one's preferences is a crucial part of the story.

One might expect that ethical preferences are justified by relating them to some first principles of ethics, like Kant's categorical imperative, Rawls' maximin principle for social structures, or Bentham's idea of maximizing utility across persons. These principles certainly deserve mindful consideration. Nevertheless, they miss a crucial feature of moral deliberations as they take place in real life. A careful look at the "moralities of everyday life" [44] will recognize "the primary locus of moral understanding as lying in the recognition of paradigmatic examples of good and evil, right and wrong: the typical cases of, for example, fairness or unfairness, cruelty or kindness, truth-telling or lying, whose merits and shortcomings even a small child 'knows at a glance'. In the history of culture and society, and in the child's own growing up also, this moral discernment has to be applied to new and more complex cases with progressively greater refinement" [21, p. 330].

This observation can be elaborated in an analysis of ethics as rule systems [20]. The relevant rules are to be understood not as a deductive system of imperative statements, but as open lists of paradigmatic examples. As Wittgenstein [57] has shown, a rule can be conveyed by offering a sequence of examples together with the injunction: try to continue it. By discussing attempts at such continuations, an understanding of the rule is developed. The point of this procedure is that there is no separate criterion behind the series in question: the series is the criterion.

Therefore, "An alternative to constructing ethical theories is to ground moral reasoning on central, paradigm cases. On these paradigm cases rests what coherence our ethical concepts can be said to have, and on these cases converges what ethical agreement we actually find between individuals and across cultures. On this model, rational persuasion replaces logical deduction, and cases replace principles as the focus of ethical reasoning. From paradigmatic cases moral argument proceeds outward, rather than downward from rational first principles and ethical systems" [12, p. 537]. According to this view, moral reasoning involves two kinds of arguments: first, arguments about what constitutes a paradigmatic example, a standard of justice; second, arguments about how different practical alternatives fare in comparison with such standards.

Still, ethical preferences may be represented by utility functions, as long as the notion of utility is used simply as a shortcut for describing preference orderings in a succinct way. The relevant preference orderings, however, differ in two important respects from the ones usually considered by RAP. First, if the very notions of justice, fairness and the like are defined in terms of paradigmatic examples, then there are standards of justice which cannot be surpassed. Technically, this can be represented by a utility function which assumes an infinite value at well-specified points. Such a utility function parts company with many RAP specimens of utility functions with regard to non-

satiability. In RAP, it is very common to assume that people can never be fully satisfied. Although this may be a sad truth about some people, it is a questionable assumption even for hedonic preferences (animals seem to experience moments of plain satisfaction). In the case of ethical preferences, however, the assumption of non-satiability misses the very point of ethical judgment.

Second, ethical preferences are necessarily incomplete. They are shaped by a past which has to be creatively interpreted in order to master the present. Simply to repeat past actions is not the way to follow ethical standards. In ethical matters, rationality includes the ability to develop a preference ordering for a novel situation which is a continuation of preference orderings developed in earlier situations. The issues raised by current biomedical research are a case in point. A physician may have a perfectly clear ethical orientation for the set of situations with which she is faced in her professional life, until some new technology raises new ethical problems. With regard to the new situation, she has no clear-cut ethical preferences. But she is not at a total loss, either. It may be quite appropriate to represent her situation by a utility function which is unequivocal over certain choice sets but which becomes ambivalent over other choice sets.

Ethical and hedonic preferences may be seen as two extremes of a continuum, with aesthetic preferences lying in between. For problems of risk management, ethical preferences are often very important. But in many human choices, e.g., with regard to consumption goods, aesthetic preferences are much more salient. Both ethical and aesthetical preferences, however, are often ambivalent in ways which again and again lead to exchanges of arguments along the lines sketched above.

A similar reasoning applies to judgments of likelihood. There are paradigmatic cases for what counts as a certain fact or as something strictly impossible ([41] provides a careful account of how human action relates to paradigmatic examples of certainty and to probability judgements). People talking to each other in a train will usually take their train ride as a fact beyond doubt and will consider a sudden jump of their train over the Pacific ocean as simply impossible. As with ethical standards, standards of certainty may evolve, too, and their practical use leaves plenty of room for situations of ambivalence.

So far, we have provided a general argument to the effect that reasoned debate about judgments of utility and of probability is in fact possible. Procedures for participatory risk management offer opportunities for such debates. Studying their dynamics and developing them as contributions to solving practical problems of risk management is a fascinating task. The task is twofold. On the one hand, there are debates starting with ambivalent judgments which end with unequivocal judgments. In these cases, RAP-based risk management is embedded into a broader procedure in such a way that the ambivalence which in the first stages precluded a straightforward application of RAP assumptions disappears later on. On the other hand, there are debates

which reinforce an initial ambivalence. In these cases, the decision rules provided by RAP must be modified, perhaps by generalizing them, perhaps by combining them with other decision rules, so as to solve decision problems in the face of persisting ambivalence. And as with standards of justice or certainty, decision rules evolve, too [9].

Exchanges of arguments about ambivalent expectations and evaluations lie at the core of social rationality. This is especially relevant with regard to the analysis of choices by social actors, i.e., actors involving a multitude of physical persons. Triggered by the seminal impossibility theorem of Arrow [4], an impressive body of literature struggles with the impossibility to define a reasonably democratic procedure that would aggregate arbitrarily given individual preferences into a social preference ordering (for a review, see Sen [50]). In conflicts about decisions under risk, widely varying preferences about actions with severe but uncertain consequences must be blended into a socially acceptable decision. But this aggregation problem does not arise only at the level of society at large. Organizations have a similar problem, and so do families and other social bodies. Human beings entertain a multitude of preferences and probability judgments which are consistent only in the frames of well-known choice situations. Novel and sometimes even familiar situations confront us with the ambivalence of our expectations and preferences. These situations trigger discussions about judgments of utility and probability which may lead to social choices by resolving critical ambivalences. Participatory tools of risk management create an open space for such discussions.

Both at the individual and the social level, expectations and preferences may evolve in the course of decision making. Sometimes they are frozen as assumed by RAP, on other occasions they are not. The definition of democracy as "government by discussion" implies that "individual values can and do change in the process of decision-making" [7, p. 120]. Risk managers need to acknowledge that "Values are established or validated and recognized through discussion, an activity which is at once social, intellectual, and creative" [27, p. 280]. And what holds for values, holds for expectations, too: they may evolve by passing through phases of ambivalence.

7. A research program for integrated assessment

As we have seen, the RAP approach to risk management runs into difficulties in situations characterized by ambivalent judgments of utility and probability. Experiences with participatory procedures as well as recent theoretical work suggest that these difficulties can be overcome by embedding the RAP approach into a broader understanding of social rationality. This perspective is especially relevant for the emerging field of Integrated Assessment (IA).

The subject of IA are decision problems whose management requires widely differing scientific disciplines. They occur remarkably often in human ecological systems, i.e.,

systems comprising human beings as well as their biophysical environment. But they are highly relevant in other fields as well, as with public health issues or in technology assessment. Ambivalent probabilities are pervasive in this class of problems because of the way boundaries between scientific disciplines have been organized in the past.

Scientific disciplines have evolved as strenuous efforts by specialized communities to provide reliable knowledge which can be trusted by the rest of society. Therefore, they have tended to define their research agendas so as to obtain results which either eliminate uncertainty or frame it in an unambiguous way. To the degree to which disciplines have been successful in this endeavor, they have relegated the inevitable areas of ambivalent probability judgment to the interstices between their respective areas of study. As IA focuses on problems in which these interstices play a prominent role, ambivalent probabilities are bound to be highly relevant, too.

Therefore, overcoming the limits of RAP is an indispensable task for successful IA. And while RAP is basically a product of the social sciences, overcoming its limits in IA is a joint venture by researchers from the social as well as the natural and engineering sciences. Maybe it is just this blend of capabilities which will enable the IA community to tackle the challenges involved in the very research program of integrated assessment. Before this background, we will now delineate three crucial features of the proposed research program.

The first feature is the development of computer models which can be used to represent ambivalent judgments of utility and probability. Computer modeling was essential for the field of IA since its beginning. The number crunching power of computers enabled researchers to use large data sets originating from very different disciplines; the techniques of computer simulation enabled them to make plausible statements about the dynamics of systems described by such data sets; taken together they enabled researchers to assess future prospects under the assumption of different policy decisions.

A particularly instructive experience in this regard is the integrated assessment of acid rain involving the RAINS model [1]. In political negotiations about international environmental policy ambivalent utilities play a crucial role, mainly via the difficulty to establish a set of agreed cost estimates. Under these conditions, it is a remarkable feature of the RAINS model that it enables its users to work with cost functions which are compatible with ambivalent utility judgments. This is due to the fact that the model computes abatement costs on the basis of a spectrum of techniques, each associated with its own cost estimate. Therefore, it is relatively easy to think through and actually model a situation where additional technologies are available, or where, to the contrary, some technologies are not available. Moreover, model users can set constraints on the technology mix employed. Under these circumstances, the overall cost estimates can be seen to depend on subjective evaluations

which may be geared to partisan interests. A model which would force users to "buy" an overall cost estimate as a scientific fact would not be very helpful for a complex negotiation process. As every negotiator knows, the ability to work with ambiguous beliefs and statements is vital in order to bring strongly dissenting partners to an agreement. On the other hand, a model which would avoid the cost issue would be of very limited use, too. After all, costs are a crucial ingredient of every major policy decision. Under such circumstances, models which enable the dissenting parties to think through and therefore to argue about different cost estimates linked to specific technologies may be of considerable help in establishing regimes to deal with environmental problems like acid rain.

While ambivalent utilities, at least in the implicit sense involved in the RAINS cost function, played a role in the early days of IA, the subject was somewhat neglected in the following years. However, much progress has been made towards dealing with ambivalent probability judgments. The familiar practice of handling decision making under uncertainty by developing best guesses for the relevant variables is still widespread, and on many occasions it is no doubt appropriate ([35] provides a wealth of examples). However, in IA this will rarely be a satisfactory procedure. The next obvious step is the use of sensitivity analysis to check the robustness of the best guesses elaborated in the first place (as done in [37]). A thoroughly probabilistic style of IA modeling has been proposed by Dowlatabadi and Morgan (see [11] for an illustration of their approach). The ICAM model developed by these authors enables the user to experiment with different subjective probability distributions, and by so doing it provides a remarkable opportunity to explore ambivalent probabilities.

Meanwhile, the link between uncertainty and preferences has been emphasized by authors drawing on the cultural theory of risk [42]. Without using the apparatus of probability theory and utility functions, Rotmans et al. have built such a link in their IA model TARGETS [46]. We are still a far cry from models which would represent ambivalent utilities and probabilities, but no doubt important milestones have already been set. In this regard, the research program of IA holds promise of further breakthroughs in the future, especially if its current focus on climatic risks should be enlarged to other problems.

As we have seen in section 6, developing models which represent ambivalent utilities and probabilities is a crucial step on the way to embed RAP into an understanding of social rationality. However, as we have seen in section 2, this can only be a first step: the real challenge is to use more comprehensive models of decision making so as to yield insights in systems involving a multitude of interacting decision makers operating in a shared environment. As a matter of fact, the great strength of RAP, on which so many techniques of risk management draw, is its ability to represent such systems.

The insights provided by RAP refer to marginal changes of the system under consideration. In particular, the kind

of market mechanisms analyzed by RAP operate in situations where the expectations and evaluations of a multitude of actors have crystallized in an unambiguous way. Under such circumstances, price signals can equilibrate supply and demand at the margin so as to avoid major discrepancies between them. IA, however, often deals with non-marginal changes. An example in point is given by the prospect of climate change as well as by the implications of policies which would keep climate change within narrow limits. The challenge of understanding the operation of price signals in non-marginal transitions is a second major feature of the research program of integrated assessment. This is a formidable task, as it requires modeling the interaction between two very different kinds of phenomena: ambivalent utilities and probabilities at the level of individual agents, prices and other social parameters at the level of markets and other institutions. Models representing this interaction will hardly ever lead to deterministic forecasts stretching over long time-spans. But they can lead to probabilistic forecasts and to insights into major dynamic patterns which are essential for policy making.

Is the research program of IA sufficiently characterized by the twin tasks of modeling ambivalent utilities and probabilities as well as modeling non-marginal changes in systems involving such utilities and probabilities? It is not, as a third task needs to be tackled, too: the design of new interfaces between the scientific community and society at large.

The scientific community has an astonishing record of processing doubts internally and providing society at large with consensual statements once these doubts have been settled. This means nothing less than that in the past the scientific community has been able to keep ambiguous judgments of probability within the boundaries of its internal debates while the statements which were passed across the interface with society at large were characterized by unambiguous probability – in fact, usually by a probability of one. And by reaching agreement on the probability of its statements, the scientific community could make a plausible claim to the effect that these probabilities were not a subjective matter, but due to objective properties of the topics under scrutiny. This is not to say that the scientific community did not communicate the unsettled questions with which it was struggling, but that it kept them well separated from the knowledge which it provided to society at large. The amazing trust with which society has endowed the scientific community (as shown by the large sums of taxpayers' money which this community is allowed to allocate by its own internal criteria) is a result of this practice. As for judgments of utility, the scientific community could credibly claim that its results were true independently of such judgments.

IA has emerged precisely because this way of defining the interface between science and society has become severely inadequate with regard to some of the most pressing problems of our time [58]. If the scientific community could only deliver statements characterized by unambiguous, consensual probabilities about, say, the prospect of cli-

mate change, it would have very little to say, indeed [39]. But this would be completely unacceptable for society at large, as the prospect of climate change needs to be dealt with one way or another long before scientists can eliminate the ambiguous probability judgments arising in this area.

Therefore, the interface between science and society has to be redesigned so as to enable the scientific community to communicate its current ambiguous probability judgments in an effective manner to society at large. This includes the ability to communicate the unavoidable subjective component in many of these judgments. Therefore, the scientific community needs to face the link between probability and utility judgments, and to engage in debates where ambivalence in both kinds of judgments are discussed. As we have seen in section 6, such debates are possible in a rational way. However, they are not a matter of scientific expertise. In particular, the scientific community has much to gain from the debates about ambivalent utility judgments which take place in society at large. For example, in studies about climate change it is important for scientists to know whether society at large judges coral reefs to be highly valuable or not, and whether this is more an aesthetic, a moral or an economic judgment. Society, on the other hand, can hardly develop such judgments about coral reefs without relying on science to enhance its perception of these reefs.

So far, debates about ambivalent probability judgments in IA have been carried out mainly in the frame of expert panels like the ones established by IPCC. In IA, expert panels are a vital complement to computer models. However, in the presence of ambivalent probability judgments IA needs processes of social learning which go way beyond expert panels. There is a need for an iterative process in which researchers exchange their judgments of probability and utility, ambivalent as they may be, with those of various stakeholders.

The challenge can be exemplified with the experience of Technology Assessment (TA), an especially interesting instance of risk management. "The original idea of TA was that an analyst would draft a report on the implications of new technological developments. Such reports would then help policymakers to design their policies" [16, p. 94]. The expectations raised by classical TA, however, have rarely if ever been fulfilled. In the countries which have made major TA efforts, these efforts have run into considerable problems of effectivity and legitimacy. "These disappointed expectations have led to the development of a new approach to TA in the various countries: assessments are supposed to result from an iterative process of debate among a variety of stakeholders, from joint analyses, and from analyses by the individual stakeholders" ([16, p. 95]; for further documentation of new approaches to TA see [8,17,40,53]).

IA in general shares with the particular instance of technology assessment the need for a structured debate involving both researchers and stakeholders [38]. Given the importance of computer models for IA, this debate needs to make IA models available outside the narrow limits of the scientific community. One possibility to do so is offered

by the procedure of IA focus groups [26]. As mentioned in section 5, focus groups are a well-established method of social research in which a group of people is exposed to some common input like a TV transmission, an oral presentation, and the like. The group is given the opportunity to digest the common experience in a shared conversation. The whole process is recorded and can be analyzed in order to register and understand the social dynamics of opinion formation about the common input. In an IA focus group, the common input is a presentation of IA models about a specific issue. These models are then used to trigger a process of opinion formation about the selected issue. A series of about five sessions provides a reasonable frame for this purpose.

The development of IA focus groups and similar procedures of participatory IA is a crucial part of the IA research program. It has serious implications for the development of IA models in at least two respects. First, the debates taking place in IA focus groups provide empirical material about how the ambivalence involved in many judgments of utility and probability is dealt with in actual debates. This is obviously useful for the task of modeling such ambivalence. Second, if core components of IA models have to be used in settings like IA focus groups, both the user interface and the internal structure of these models must be amenable to such use. The current generation of IA models is still designed mainly for a user community consisting of scientific researchers. Designing IA models which can be used by stakeholders is a serious, if fascinating, challenge.

The challenge is related to a major research question: what selection of stakeholders is appropriate in different IA procedures? The question has two sides. First, there is the issue of sampling. In principle, this is easy to handle as soon as the relevant populations have been defined. In practice, financial constraints will often require great skills to design a reasonable sampling procedure. Second, there is the problem of defining the relevant populations, namely the categories of stakeholders to be included. In designing a passenger aircraft, it seems sensible to include pilots, maintenance personnel, stewardesses, and passengers themselves in the design process. In siting a toxic waste landfill, representatives of the possible sites, of relevant industries and political organizations, as well as citizens living in the area under consideration may be included. Clearly, there is an important difference between stakeholders which bring some specific professional competence to the process and people who enter the process on the basis of plain common sense. Learning how to work with each category is an important part of the research program of IA, and so is learning how to combine the different categories.

Finally, there is the grand challenge of dealing with abuses of power and money along with their many dreadful consequences. IA research cannot avoid ambivalent judgments of probability, and for this very reason it cannot avoid ambivalent judgments of utility. As a consequence, it is faced with the opportunity, but also the necessity to design new interfaces with society at large. At these in-

terfaces, the scientific community has to share its own internal doubts while engaging about issues – namely ambivalent judgments of utility – where it has no particular competence. This may well make the scientific community highly vulnerable to abuses of power and money: external pressures can rather easily turn arguments into directions unwarranted by any actual research results.

Abuses of power and money certainly were a recurrent phenomenon in the traditional interface between science and society, too. In particular, the prevalence of military objectives in 20th century science funding offered plenty of scope for such abuses. However, the design of interfaces suited to communicate the ambivalences which are inherent in any IA effort offers opportunities to expose potential abuses of power and money to critical scrutiny. This is a strong reason to engage in an open dialogue between researchers and stakeholders, including ordinary citizens.

The potential of democratic problem solving is quite remarkable, as is highlighted by the “phenomenon that no major famine has ever taken place in any country with a multiparty democracy with regular elections and with a reasonably free press” ([52, p. 16], with references to further literature). If democracy were the outcome of a simple majority voting based on individual self-interest, this phenomenon would be unintelligible. If democracy were the outcome of an unconstrained consensus formed in disinterested debate, the same phenomenon would cease to be remarkable. In fact, however, democracy is based on the combination of formal decision procedures with public debate. Power and money do not necessarily distort collective problem solving, quite often they are essential ingredients of it.

If IA takes ambivalent expectations and evaluations into account, more often than not it will express conflicts of interests, quite often without being able to do much about them. But the very fact of providing such an expression together with opportunities to debate it may sometimes help to free the overall decision process at least from some errors and mischief. By providing computer models as well as participatory procedures to use them, IA may help to define procedures for decision making under conditions of ambivalent expectations and evaluation held by conflicting stakeholders.

Suitable procedures need to provide opportunities for a honest search for mutual understanding and attempts to account for one’s positions [3]. But they also need to provide opportunities for the negotiations and bargains without which democracy cannot work [48]. It is impossible to design a procedure which would guarantee a successful outcome, as this would contradict the very creativity required. But it is possible and often quite useful to explain why a shared debate among stakeholders, explicitly including laypersons, makes sense [51]. If and when such a debate takes place, scientists can offer IA models to share their knowledge and their doubts about plausible consequences of various actions.

8. Conclusion

RAP-based risk management works well where objective probabilities are given together with a single utility function. It can still work where different, but unambiguous subjective probabilities and utility functions are aggregated into a social outcome by competitive futures markets, including markets for insurance services. It runs into severe difficulties as soon as the relevant actors hold ambivalent probability and utility judgments. In these situations optimization procedures fail because no well-defined objective function exists. Such situations may be characterized as situations of complexity because they allow for a plurality of legitimate evaluations and descriptions. In these situations, which arise frequently in environmental management, RAP-based risk management can and should evolve into a practice of integrated assessment.

Non-optimizing problem solving is a normal practice in human life. In particular, standards of justice evolve in a non-optimizing way as a result of on-going debates among various stakeholders. Procedures of discursive conflict resolution provide opportunities for such debates, thereby assisting decision making under conditions of complexity. Recent advances in computer modeling greatly facilitate this task, because they enable the users to explore the implications of varying assumptions about the problem at hand. Further work may lead to models which confront the users with the fact that quite often several different alternatives are equally desirable according to the preferences and expectations held by those very users. Obviously, this would reinforce the importance of participatory procedures. Along these lines, crucial difficulties of traditional risk management can be tackled by the emerging practice of integrated assessment. This challenge marks a promising research program which may be realized in a joint effort of natural, engineering, and social scientists.

Acknowledgements

Research leading to the present paper was done in projects Nos. 4031-33525 and 5001-44593 of the Swiss National Science foundation and in the project "ULYSSES – Urban Lifestyles, Sustainability, and Integrated Assessment", financed by the fourth framework program of the European Union. An early version of the paper was given as a keynote address at the Conference "Ecology, Society, Environment" of the Society for Ecological Economics, Paris, June 23–25, 1996. Elements of the paper were presented – in collaboration with Ortwin Renn – at the conference of the Society for the Advancement of Socio-Economics, Geneva, July 12–14, 1996. Further elements were presented as a keynote address at the second open meeting of the Human Dimensions of Global Environmental Change Research Community, IIASA, Laxenburg, June 12–14, 1997. The Center for Interdisciplinary Technology Studies, ZIT, at Darmstadt University

provided an important source of inspiration. Discussions with H.R. Böhm, T. Burns, H. Dowlatabadi, O. Edenhofer, Y. Elkana, S. Funtowicz, B. Kasemir, G. Morgan, C. Pahl-Wostl, A. Rabl, S. Rayner, J. Ravetz, P. Reichert, O. Renn, E.A. Rosa, J. Rotmans, R. Schüle, M. van Asselt, M. Waechter, H. Wasmer, and T. Webler, as well as significant help by two anonymous reviewers are gratefully acknowledged. Responsibility for errors stays with the author.

References

- [1] J. Alcamo, R. Shaw and L. Hordijk, *The RAINS Model of Acidification: Science and Strategies in Europe* (Kluwer, Dordrecht, 1990).
- [2] M. Allais, *Econometrica* 21 (1953) 503–546.
- [3] K.O. Apel, *Diskurs und Verantwortung. Das Problem des Übergangs zur Postkonventionellen Moral* (Suhrkamp, Frankfurt, 1988).
- [4] K.J. Arrow, *Social Choice and Individual Values* (Wiley, New York, 1951).
- [5] T. Bayes, *Essai en Vue de Resoudre un Problème de la Doctrine des Chances* (Société Française d'Histoire des Sciences et des Techniques, Paris, 1988).
- [6] K. Binmore, De-Bayesing game theory, in: *Frontiers of Game Theory*, eds. K. Binmore, A. Kirman and P. Tani (MIT Press, Cambridge, MA, 1993).
- [7] J. Buchanan, *Journal of Political Economy* 622 (1954) 114–123.
- [8] J. Bunders, *Participative Strategies for Science-Based Innovations: The Case of Biotechnology for Small-Scale Farmers in Developing Countries* (VU University Press, Amsterdam, 1994).
- [9] T.R. Burns and T. Dietz, *International Sociology* 73 (1992) 259–283.
- [10] P.C. Dienel, *Die Planungszelle* (Westdeutscher Verlag, Opladen, 1992).
- [11] H. Dowlatabadi and G. Morgan, *Energy Policy* (1993) 209–221.
- [12] C. Elliot, *Inquiry* 34 (1991) 525–538.
- [13] D. Ellsberg, *Quarterly Journal of Economics* 75 (1961) 643–669.
- [14] P.C. Fishburn, *Nonlinear Preference and Utility Theory* (The Johns Hopkins University Press, Baltimore, MD, 1988).
- [15] S.O. Funtowicz and J.R. Ravetz, *Ecological Economics* 10 (1994) 197–207.
- [16] J. Grin and H. van de Graaf, *Science, Technology and Human Values* 211 (1996) 72–99.
- [17] E.G. Guba and Y.S. Lincoln, *Fourth Generation Evaluation* (Sage, Newbury Park, CA, 1989).
- [18] S.G. Hildebrand and J.B. Cannon, eds., *Environmental Analysis: The NEPA Experience* (Lewis Publisher, Boca Raton, FL, 1994).
- [19] C.C. Jaeger, O. Renn, G. Rosa and T. Webler, *Risk, Uncertainty, and Rational Action* (EAWAG, Zürich, 1996).
- [20] C.C. Jaeger and A.J. Rust, *Inquiry* 37 (1994) 65–84.
- [21] A. Jonsen and S. Toulmin, *The Abuse of Casuistry* (University of California Press, Berkeley, CA, 1988).
- [22] S. Joss and J. Durant, *Public Participation in Science: The Role of Consensus Conferences in Europe* (Science Museum, London, 1995).
- [23] R. Jungk and N.R. Müllert, *Die Zukunftswerkstatt* (Heyne, München, 1994).
- [24] D. Kahnemann, P. Slovic and A. Tversky, eds., *Judgment Under Uncertainty: Heuristics and Biases* (Cambridge University Press, Cambridge, 1982).
- [25] D. Kahnemann and A. Tversky, *Econometrica* 21 (1979) 503–546.
- [26] B. Kasemir, J. Behringer, B. De Marchi, C. Deuker, G. Dürrenberger, S. Funtowicz, A. Gerger, M. Giaoutzi, Y. Haffner, M. Nilsson, C. Querol, R. Schüle, D. Tabara, M. van Asselt, D. Vassilarou, N. Willi and C.C. Jaeger, Focus groups in integrated assessment: The ULYSSES pilot experience, ULYSSES Working Paper 97-4, ZIT, Darmstadt University of Technology (1997).

- [27] F. Knight, *Freedom and Reform: Essays in Economic and Social Philosophy* (Harper, New York, 1947).
- [28] D. Kreps, *Notes on the Theory of Choice* (Westview Press, Boulder, CO, 1988).
- [29] S. Krimsky and D. Golding, eds., *Social Theories of Risk* (Praeger, Westport, CT, 1992).
- [30] R.A. Krueger, *Focus Groups – A Practical Guide for Applied Research* (Sage, London, 1994).
- [31] F.M. Lynn and G.J. Busenberg, *Risk Analysis* 152 (1995) 147–162.
- [32] M. Machina, *Econometrica* 50 (1982) 277–323.
- [33] R. Minard, K. Jones and C. Paterson, *State Comparative Risk Projects: A Force for Change* (The Northeast Center for Comparative Risk, South Royalton, VT, 1993).
- [34] G. Morgan and M. Henrion, *Uncertainty. A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis* (Cambridge University Press, Cambridge, 1990).
- [35] N. Nakicenovic, W.D. Nordhaus, R. Richels and F.L. Toth, *Integrative Assessment of Mitigation, Impacts, and Adaptation to Climate Change* (IIASA, Laxenburg, 1990).
- [36] J. von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior* (Princeton University Press, Princeton, NJ, 1944).
- [37] W.D. Nordhaus, *Managing the Global Commons. The Economics of Climate Change* (MIT Press, Cambridge, MA, 1993).
- [38] S. O'Hara, *Ecological Economics* 16 (1996) 95–107.
- [39] C. Pahl-Wostl, C.C. Jaeger, S. Rayner, C. Schär, M. van Asselt, D.M. Imboden and A. Vckovski, in: *Views from the Alps: Regional Perspectives on Climate Change*, eds. P. Cebon, U. Dahinden, H.C. Davies, D.M. Imboden and C.C. Jaeger (MIT Press, Cambridge, MA, in press).
- [40] T. Pinch and W.E. Bijker, in: *The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology*, eds. W.E. Bijker, T.P. Hughes and T. Pinch (MIT Press, Cambridge, MA, 1987).
- [41] H. Putnam, *The Many Faces of Realism* (Open Court, LaSalle, IL, 1987).
- [42] S. Rayner, Cultural theory and risk analysis, in: *Social Theories of Risk*, eds. S. Krimsky and D. Golding (Praeger, Westport, CT, 1992).
- [43] P. Reichert, *Water Science and Technology* 5 (1997) 149–156.
- [44] O. Renn et al., *Policy Sciences* 26 (1993) 189–214.
- [45] O. Renn, T. Weblar and P. Wiedemann, *Fairness and Competence in Citizen Participation* (Kluwer, Dordrecht, 1995).
- [46] J. Rotmans and B. de Vries, *Perspectives on Global Change. The TARGETS Approach* (Cambridge University Press, Cambridge, 1997).
- [47] J. Sabini and M. Silver, *Moralities of Everyday Life* (Oxford University Press, Oxford, 1982).
- [48] G. Sartori, *Democratic Theory* (Wayne State University Press, Detroit, MI, 1962).
- [49] L.J. Savage, *The Foundations of Statistics* (Wiley, New York, 1954).
- [50] A. Sen, *American Economic Review* 851 (1995) 1–24.
- [51] K.S. Shrader-Frechette, *Risk and Rationality. Philosophical Foundations for Populist Reform* (University of California Press, Berkeley, CA, 1991).
- [52] N.J. Smelser, *American Sociological Review* 63 (1998) 1–16.
- [53] W.A. Smit, in: *Non-Provocative Defence as a Principle of Arms Reductions and Its Implications for Assessing Defence Technologies*, eds. M.L.A. Ter Borg and W.A. Smit (VU University Press, Amsterdam, 1989).
- [54] P. Walley, *Statistical Reasoning with Imprecise Probabilities* (Chapman and Hall, London, 1991).
- [55] C. Whipple and V. Covello, eds., *Risk Analysis in the Private Sector* (Plenum Press, New York, 1985).
- [56] H.C. White, *Identity and Control. A Structural Theory of Social Action* (Princeton University Press, Princeton, 1992).
- [57] L. Wittgenstein, *Philosophical Investigations* (Basil Blackwell, Oxford, 1953).
- [58] B. Wynne, *Global Environmental Change* 22 (1992) 11–127.
- [59] R. Zieschank, *Zeitschrift für Umweltpolitik und Umweltrecht* 1/91 (1991) 27–51.