

Hypothetical Retrospection

Sven Ove Hansson

Accepted: 04 October 2006 / Published online: 3 January 2007
© Springer Science + Business Media B.V. 2007

Abstract Moral theory has mostly focused on idealized situations in which the morally relevant properties of human actions can be known beforehand. Here, a framework is proposed that is intended to sharpen moral intuitions and improve moral argumentation in problems involving risk and uncertainty. Guidelines are proposed for a systematic search of suitable future viewpoints for hypothetical retrospection. In hypothetical retrospection, a decision is evaluated under the assumption that one of the branches of possible future developments has materialized. This evaluation is based on the deliberator's present values, and each decision is judged in relation to the information available when it was taken. The basic decision rule is to choose an alternative that comes out as morally acceptable (permissible) from all hypothetical retrospections.

Key words hypothetical retrospection · regret · retrospection · risk · uncertainty · expected utility · prudence

1 Introduction

Uncertainty and risk are pervasive features of practical decision making,¹ It is in fact difficult to find an example of decision-making in real life that does not contain a component of uncertainty. In spite of this, moral theorizing has mostly focused on idealized, “deterministic” situations in which the morally relevant properties of human actions can be known beforehand.² It seems to be tacitly assumed that if moral philosophy delivers solutions for the deterministic

¹Following convention I use the term “risk” to denote such lack of knowledge that can be expressed in probabilities, and “uncertainty” to denote lack of knowledge that cannot be so expressed. For simplicity, the word “risk” will sometimes be used in lieu of “risk or uncertainty.”

²The idealization used in most of moral theory is in fact much stronger than determinism. The consequences of an agent's actions are assumed to be not only determined but also knowable at the point in time of deliberation. This corresponds fairly well with the standard decision-theoretical notion of decision-making under certainty. (Luce and Raiffa 1957, p 13) For reasons of convenience, the term “deterministic” is used here to denote such conditions for (moral) decision-making.

cases, then these solutions can be used by decision theory to derive solutions for the more realistic, indeterministic cases. However, this division of labour between the two disciplines is far from unproblematic. Since decision theory operates exclusively with criteria of rationality, such a procedure leaves out the moral aspects of risk or uncertainty itself. In other words, it cannot deal with moral issues that arise in the uncertain situation but not in any of the alternative (deterministic) chains of events in terms of which the uncertainty is characterized. This precludes a moral account of actions such as taking or imposing a risk that do not take place in the idealized deterministic situations to which moral analysis is restricted in this model. (Hansson 2001) A further complication is that the more common non-utilitarian ethical theories are not easily combined with a decision-theoretical framework.

Whereas ethics is virtually silent on risk and uncertainty, risks are normatively evaluated in other disciplines, primarily risk analysis and risk-benefit analysis (a form of cost-benefit analysis). It would therefore be natural to look to these disciplines for guidance on how risk and uncertainty can be dealt with in ethics.

The standard procedure in these disciplines is expected utility maximization, according to which we should choose an action with the highest probability-weighted average of the values of the possible outcomes. Hence, provided that lost lives are valued equally and additively, the risk associated with a probability of 1 in 1,000 that 1,000 persons will die is considered equivalent to certainty that one person will die.³ This model excludes the use of decision-weights that are not proportional to probabilities, such as risk-averse or cautious rules that give special weight to the avoidance of improbable but very large catastrophes.⁴ Furthermore, person-related moral considerations such as voluntariness, consent, and justice are disregarded. These limitations of the model are prominent in the many conflicts that its application gives rise to – conflicts that are usually depicted as failures in the communication between experts and laypersons but are often better described as failures of a thought model that excludes legitimate normative issues to which members of the public attach great importance. (Hansson 2005a)

Furthermore, this application of expected utility theory is unstable against the actual occurrence of a serious negative event that was included in the calculation. This can be seen by studying the post-accident argumentation after almost any accident. If the expected utility argumentation were followed to the end, then many accidents would be defended as consequences of a maximization of expected utility that is, in toto, beneficial. However, this type of reasoning is very rarely heard in practice. Seldom do we hear a company that was responsible for an accident say that the accident was acceptable and only to be counted with as part of the maximization of total utility. Instead, two other types of reactions are common. One of these is to regret one's shortcomings and agree that one should have done more to prevent the accident. The other is to claim that someone else was responsible for the accident. It should also be noted that accident investigation boards are instructed to

³ See Hansson (1993) for some of the arguments against this way of evaluating losses in lives. Although it does not follow from expected utility maximization that deaths should be evaluated in this way, this evaluation method is almost universally applied in risk analysis and risk-benefit analysis.

⁴ One well-known decision rule that employs such weights is Kahneman's and Tversky's ([1979] 1988) prospect theory. In that theory, the objective probability $p(A)$ of an event A is replaced by $\pi(p(A))$, where π is an increasing function from and to the set of real numbers between 0 and 1. In prospect theory, $\pi(p(A))$ takes the place that $p(A)$ has in expected utility theory. However, the composite function $\pi(p(\cdot))$ does not satisfy the axioms of probability, and therefore prospect theory is not a form of expected utility theory but an alternative to it.

answer the questions “What happened? Why did it happen? How can a similar event be avoided?,” not the question “Was the accident defensible in an expected utility calculation?”⁵ Once a serious accident has happened, the application of expected utility maximization appears much less satisfactory than what it did before the accident. In this pragmatic sense, expected utility maximization is not a stable strategy.

We should therefore not expect to obtain an adequate ethics of risk and uncertainty by transferring the thought model that dominates in risk analysis to the more general application area of ethics. Instead, a systematic account needs to be developed that allows us to deal explicitly with the ethical aspects of risk, such as risk-imposition and risk-taking. It should also be stable against the occurrence of serious negative events. The purpose of the present paper is to provide an outline of such an ethical account of risk. Risk management cases will often be used as illustrations, but the intended scope of application is moral argumentation in general.

2 The Basic Idea of Hypothetical Retrospection

The risks and uncertainties under consideration here concern what can happen in the future. Therefore it is a good starting-point to consider the future-related arguments that we appeal to in supposedly deterministic settings. The most basic such argument can be called the “foresight argument.” It consists in a reminder of the effects that a certain act or behaviour will have at some later point in time. As an example of this, some of the consequences of drinking excessively tonight can, for practical purposes, be regarded as foreseeable. Thinking about these consequences may well be what deters a person from drunkenness. Colloquially, the argument is often stated in terms of regret: “Do not do that. You will regret it.” As we will see, predicted regret is only a first approximation that needs to be replaced by a more carefully carved-out concept. However, the basic insight is that a prediction of how one would evaluate one’s actions at some later point in time can be a useful component of moral reflection.

In indeterministic cases there are, at each point in time, several alternative “branches” of future development.⁶ Each of them can be referred to in a valid moral argument about what one should do today. As a first approximation, we wish to ensure that whichever branch materializes, a posterior evaluation should not lead to the conclusion that what one did was morally wrong.⁷ If this is achieved, then so is the stability that was seen in Section 1 to be missing in standard applications of expected utility maximization.

⁵ <http://www.havkom.se/eng/pdf/AboutSHK.pdf>. Cf. <http://www.notisum.se/mp/SLS/LAG/19900712.HTM>.

⁶ I use the noun “alternative” to denote an option than can be chosen in a particular decision. A “branch” (or “branch of possible future development”) is one of the possible developments after a particular event (typically after the choice of an alternative in a decision). For the present purposes, a branch is not followed indefinitely into the future but only to a certain point in time that, in combination with the branch itself, constitutes a “viewpoint” from which evaluations can be made.

⁷ Much of this is applicable to self-regarding decision-making as well, but here the focus will be on decision-making that has morally relevant effects on others than the decision-maker.

Related ideas have been expressed in terms of regret-avoidance.⁸ However, regret is often unavoidable for the simple reason that it may arise in response to information that was not available at the time of decision. Suppose you decline an offer to invest your savings in a high-risk company, only to find out half a year later that the company has boomed and made all its shareholders rich. You may then regret that you did not invest in the company, while at the same time recognizing that at the time of the decision you did the right thing, given what you then knew.⁹ Therefore, predicted regret is not a good decision guide.¹⁰ This is not a new insight. Several authors seem to have approached criteria related to our notion of hypothetical retrospection, but have backed off due to the implausibility of a regret-avoiding decision strategy.¹¹ In order to make a strategy of hypothetical retrospection workable, we need to specify a type of retrospection that avoids these problems, so that it can achieve the decision stability we require.

In hypothetical retrospection, each evaluation should refer to a branch of future development from the decision up to the moment at which the retrospection is enacted. This means that the evaluation should not be restricted to the endstate (i.e. to effects remaining at the point in time of the evaluation), but also cover the process leading up to it. Furthermore, in order to be decision-guiding, hypothetical retrospection has to refer to the decision that one should have made given the information (actually) available at the time of the decision, not the information (hypothetically) available at the time of the retrospection. The decision-relevant moral argument is not of the form “Given what I now know I should then have...,” but rather “Given what I then knew, I should then have...” The purpose of hypothetical retrospection is to ensure serious consideration of possible future developments and what can be learnt from them, not to counterfactually reduce the uncertainty under which the decision must be taken.

Since hypothetical retrospection refers back to a situation in which one did not know which branch of future development would materialize, arguments may be used that refer to the other branches. Consider a factory owner who has decided to install an expensive fire

⁸ The idea of hypothetical retrospection is an extension and regimentation of patterns of thought that are prevalent in everyday moral reasoning. Unsurprisingly, numerous instances of related ideas can be found in the philosophical literature. Careful consideration of one’s future interests was recommended by Plato’s Socrates (Prot. 356a–e). In classical accounts of prudence, the moral perspective from future hindsight was a key component. (Mulvaney 1992; Vanden Houten 2002) John Rawls’s concept of deliberative rationality, ascribed by him to Sidgwick, includes a notion of a rational plan of life and a requirement that “a rational individual is always to act so that he need never blame himself no matter how things finally transpire” (Rawls 1972, p 422; for comments see Williams 1976 and Ladmore 1999). Similar ideas by Nagel (1986, p 127) have been interpreted by Dickenson (1991, p. 51) as a remorse-avoiding strategy. In decision theory, regret-avoiding strategies have been given exact formal interpretations. (Bell 1982; Loomes and Sugden 1982; Sugden 1985) A proposal to use regret-avoidance as a moral principle was put forward by Ernst-Jan Wit (1997). Jeffrey’s (1983) criterion of ratifiability should also be mentioned, since it recommends what can be described as hypothetical retrospection with respect to probability assignments.

⁹ Weber (1998, pp 105–106) distinguishes between outcome regret that refers to “how things turned out” and “decision-making regret” that requires “that one can, in hindsight, think that one had an available reason at the time of choice to choose other than as one did.” In hypothetical retrospection, only the latter form of regret should be (hypothetically) elicited.

¹⁰ The same applies, for similar reasons, to several related notions such as a predicted wish to have acted differently.

¹¹ Williams (1976, pp 130–131) criticized Rawls for not paying attention to preference changes when proposing that a rational individual should act so that he will never need to blame himself. Humberstone (1980) and Weirich (1981) both argued convincingly, but with different arguments, that it may be rational to do something one knows that one will regret. More recently, Ladmore (1999) has again emphasized that our decisions can never be immune to regret since our later self judges earlier choices on the basis of changed preferences.

alarm system in a building that is used only temporarily. When the building is taken out of use, the fire alarm has yet never been activated. The owner may nevertheless consider the decision to install it to have been right, since at the time of the decision other possible developments (branches) had to be considered in which the alarm would have been life-saving. This argument can be used, not only in actual retrospection, but also, in essentially the same way, in hypothetical retrospection before the decision. Similarly, suppose that there is a fire in the building. The owner may then regret that he did not install a much more expensive but highly efficient sprinkler system. In spite of this regret he may consider the decision to have been correct since when he made it, he had to consider the alternative, much more probable development in which there was no fire, but the cost of the sprinklers would have made other investments impossible. Of course, this argument can be used in hypothetical retrospection just like the previous one. In this way, when we perform hypothetical retrospection from the perspective of a particular branch of future development, we can refer to each of the alternative branches and use it to develop either counterarguments or supportive arguments. In short, in each branch we can refer to all the others.

Hypothetical retrospection aims at ensuring that whatever happens, the decision one makes will be morally acceptable (permissible) from the perspective of actual retrospection. To accomplish this, the decision has to be acceptable from each viewpoint of hypothetical retrospection.¹² What makes this feasible is of course that although each hypothetical retrospection takes a viewpoint in one particular branch of future development, from that viewpoint it deliberates on what one should have done, given the knowledge available at the point in time of the decision, and therefore it also takes into account the need to be prepared for the other branches.¹³

Arguably, the criterion of full acceptability in all branches is not always achievable, in particular not in conflict-ridden situations with high stakes. If a serious accident that we believed to be very improbable nevertheless happens, then we are almost sure to learn something from it, and see what we did before in a new light.¹⁴ Only rarely would it be appropriate to say that enough had been done to prevent the accident and that nothing should have been done differently. In order to achieve acceptability in this branch, we may have to take measures that would be so costly and cumbersome that they are unacceptable in at least some of the (perhaps much more probable) branches in which no accident of this type takes place.¹⁵ This is foreseeable, and therefore we can also foresee that full acceptability in every branch is impossible to obtain. At least for the sake of argument let us assume that there are cases like this in which full acceptability in all branches cannot be achieved.

Such cases can be described as a (generalized) form of moral dilemmas. According to a common view, a moral dilemma consists in a situation in which there are several alternatives to choose between, but none of them is considered to be morally acceptable. Similarly, in the cases referred to here, there are several alternatives to choose between, but none of them is morally acceptable in all branches of future development. As I have argued elsewhere (Hansson 1999), even if a moral dilemma cannot be solved, we can make an optimal moral choice when confronting it. To achieve this we have to choose an option that is not normatively inferior to any other option. (In some dilemmas, two or more alternatives satisfy this criterion, in others only one.) Similarly, if no alternative is available that is

¹² It is not assumed that the decision has to be morally optimal, or satisfy moral requirements maximally. Hence, scope is left for choice between alternatives that are all acceptable but not all of the same moral value.

¹³ Obviously, errors of prediction cannot be avoided. This is a problem shared by all decision rules.

¹⁴ As one example of this, the fact that the accident took place gives us a reason to reconsider whether we were justified in believing it to be highly improbable.

¹⁵ On the use of probabilistic arguments in hypothetical retrospection, see Section 4.

acceptable from every future viewpoint, we should choose an alternative that comes as close as possible to that ideal. This means that we should determine the lowest level of unacceptability that some alternative does not exceed in any branch, and choose one of the alternatives that do not exceed it.¹⁶

In the decisions that I make at a particular point in time I am committed to apply the moral values that I have at that point in time, since these are the only values that I am then bound by. There is no reason to apply moral values that one expects to acquire but does not presently endorse. Therefore, decision-guiding hypothetical retrospection has to be specified so that it (contrary to predicted regret) refers to the moral values at the time when the actual deliberation takes place, not the time at which the hypothetical retrospection is staged.¹⁷ In practice, the values at the time of the deliberation will normally coincide with the values at the time of the decision. The hypothetically retrospective judgment is a (counterfactual) statement about the judgment that one would make at a future point in time if one's values did not change. Since it is difficult to imagine how one would reason if one had different moral values than one has, this proviso does not seem to complicate the process in practice, although it may be seen as philosophically complicating.

In order to make systematic use of hypothetical retrospection we need to identify branches of future development, and points in time in these branches, at which retrospections will be enacted. A selection is necessary since it will not be possible to attend to all possible branches. This selection should be based on a search that aims at identifying branches that can be predicted to have an influence on the decision. We aim at choosing an alternative that is as defensible as possible in all branches of future development that can follow after it has been chosen. Therefore, we should attempt to find, for each alternative, those among its possible subsequent branches in which it is least defensible. In other words, for each alternative we should make a search among the branches that can follow after it, trying to identify those among these branches in which the choice of this alternative will be most difficult to defend in hypothetical retrospection.¹⁸ This refutationist approach can be summarized as a simple rule of thumb: "For each alternative, find out when it would most difficult to defend."

¹⁶ To make this more precise (although admittedly somewhat overexact), let there be a set A of alternatives and a set V of viewpoints. Each viewpoint is constituted by a point in time in a branch of possible future development, as explained above in footnote 6. Let $f(a,v)$ denote the degree to which the alternative a violates moral requirements as seen from the viewpoint v . If $f(a,v)=0$ then a does not at all violate any moral requirements as seen from the viewpoint v . We have a moral dilemma if and only if for all $a \in A$ there is some $v \in V$ such that $f(a,v) > 0$. The proposed solution is to choose an a such that for all $a' \in A$:

$$\max_{v \in V} (f(a', v)) \geq \max_{v \in V} (f(a, v))$$

Although this rule does not explicitly mention probabilities, it is not a non-probabilistic decision rule in the same sense as for instance the maximin rule in its standard version that leaves us without resources for making use of probabilistic information. In contrast, for each a and v , $f(a,v)$ represents an evaluation in which the probabilistic information available at the time of decision should be taken into account.

¹⁷ In contrast, it is often reasonable to take predicted or expected future changes in preferences into account. Admittedly, the distinction between moral values and personal preferences is not always crystal-clear.

¹⁸ In the language of footnote 16, this corresponds to identifying, for each a , the value of

$$\max_{v \in V} (f(a, v)),$$

which is exactly what we need to apply the decision rule described there.

For obvious practical reasons, the search for future viewpoints suitable for hypothetical retrospection cannot reach indefinitely into the future. Fortunately, the proposed refutationist search procedure incorporates a mechanism that tends to restrict the time perspective. As we go further into the future, the effects of choosing one alternative in preference to another will become more and more uncertain. With this increasing uncertainty, there will in general be less substantial differences in value between alternative branches. Therefore, the refutationist approach will in typical cases preclude excursions into the far-off future.

An additional criterion needs to be applied in the search procedure. When different alternatives are compared, they should if possible be compared in terms of future developments (branches) that are specified in the same respects.¹⁹ Consider a choice between (1) receiving 10,000 euros 1 month from now and (2) receiving a lottery ticket that will with 50% probability win 30,000 euros a year from now and with 50% probability nothing. Hypothetical retrospection that goes only 1 month ahead in time will lead to asymmetry between the branches in terms of available knowledge. In order to achieve symmetric comparisons, the stage for hypothetical retrospection should instead be set 1 year ahead, when the relevant information is available in all the branches.

The basic principles for decision-guiding hypothetical retrospection can be summarized as follows:

A hypothetical retrospection is an evaluation of a decision in relation to its alternatives. It is hypothetically enacted at some future point in time in any of the branches of possible future developments following after the decision, but it is based on the deliberator's present values. Its outcome is an evaluation of the decision as seen in relation to what the agent was justified in believing at the time when it was performed.

Moral deliberation under conditions of risk or uncertainty should include a systematic search for future viewpoints for hypothetical retrospection. The major guiding principle in this search should be to find, for each alternative, the future developments under which it would be most difficult to defend morally in hypothetical retrospection.

If there is an alternative that comes out as morally acceptable (permissible) in every hypothetical retrospection that is enacted from a viewpoint at which this alternative has been chosen, then such an alternative should be chosen. Otherwise, an alternative should be chosen that does not in any hypothetical retrospection exceed the lowest level of inacceptability that some alternative does not exceed in any hypothetical retrospection.

It might be contended against the present proposal that hypothetical retrospection adds nothing of importance to the process of moral evaluation. Since the retrospective evaluation concerns the decision that one is about to make, and assumes the knowledge and the values that one actually has, why take the trouble of going mentally "forwards and backwards" in time to make this evaluation? Why not reason directly about the alternatives instead of pondering how one would judge them in a perspective of hindsight?

My answer is that the realism induced by hindsight adds seriousness and concreteness to the process of moral reflection, and can therefore change its outcome. The purpose of hypothetical retrospection is to make moral evaluations more foresightful by simulating

¹⁹ This is an instance of the general requirement for rational comparisons that they should refer to the same aspects of the comparanda. This is only possible if the same characteristics are available for the different comparanda. Hence, it would be difficult to compare two opera performances if you have only heard a sound recording of one of them and seen a silent film of the other.

these effects of afterthought. To the extent that this is achieved, hypothetical retrospection has substantial effects on the outcome of moral reflection. (These effects are to some extent analogous to those of considering one's moral decisions from the perspective of other concerned individuals.)

Examples from risk management can be used to illustrate this. Many severe accidents have been caused by negligent or careless acts (or omissions) that stand out as indefensible after the accident has happened. Such acts can be avoided by means of moral reflection that includes hypothetical retrospection, perhaps in the simplified risk manager's version "make a decision that you can defend also if an accident happens." Arguably, this simple recipe can provide more useful guidance to risk managers than many applications of cost-benefit or risk-benefit analysis.

3 A Framework for Moral Argumentation

The proposed procedure of hypothetical retrospection represents a style of ethical thinking that differs from traditional ethical theories. This is not a proposal for a complete moral theory. Instead, it is a framework for argumentation that is intended to sharpen moral intuitions with respect to one particular aspect that has been neglected in moral philosophy, namely the moral problems of risk and uncertainty.

A comparison with John Rawls's contractarian theory can be used to clarify the intended function of this proposal. The two major components of Rawls's theory are a framework for reasoning (the original position) and a set of criteria (primarily his "two principles of justice") with which alternative courses of actions can be morally evaluated.²⁰ The intended function of hypothetical retrospection corresponds to that of the first of these two parts of Rawls's theory (whereas expected utility maximization corresponds to the second). However, there are also important differences. The framework for moral reasoning proposed here is intended to be applied directly to the moral problem at hand, not as a mean to derive rules that will in their turn, in a second phase of moral discourse, be applied to real moral problems. Furthermore, although the present framework makes use of the reasoner's capacity for imagination and hypothetical reasoning, it does so only in order to invoke fully realistic situations for moral appraisal, namely situations that represent possible future developments.²¹ Therefore, in contrast to major trends in current moral philosophy (of which hypothetical social contracts is but one example²²) this proposal represents a step from the abstract to the concrete, an attempt to apply moral intuitions in a systematized way directly to the objects of our deliberations. However, the form of moral reasoning represented by this proposal should not be confused with an atheoretical approach that does not go beyond what is immediately given. In order to make full deliberative use of our moral appraisals of realistic situations, we need (1) systematic procedures for finding relevant and realistic scenarios, such as the refutationist search procedure described above, and (2) an account of the various types of moral arguments that can validly be used within

²⁰ Rawls himself emphasized the independence of the two components, and noted that "[o]ne may accept the first part of the theory (or some variant thereof), but not the other, and conversely." (Rawls 1972, p 15) In subsequent discussions the distinction between the two components has sometimes been less clear.

²¹ An exception must be made if we extend the procedure of hypothetical retrospection to future viewpoints in which the agent is no longer present as a capable reasoner. In such cases, a hypothetical evaluator with the same moral standards as the agent can be used as a heuristic device.

²² On the use of unrealistic examples in moral philosophy, see Lucey (1976), Lackey (1976), and Ward (1995).

this framework for comparing and combining the insights gained from different scenarios. Several major types of arguments will be introduced briefly below in Sections 4 to 5.

The concreteness gained through hypothetical retrospection has the advantage that our moral deliberations will be based on “the full story” rather than on curtailed versions of it. More specifically, this procedure brings to our attention the interpersonal relations that are essential in a moral appraisal of risk and uncertainty, such as who exposes whom to a risk, who receives the benefits from whose exposure to risk, etc. It is only by staying off from this concreteness that standard utility-maximizing risk analysis can remain on the detached and depersonalized level of statistical lives and free-floating risks and benefits.

In the absence of empirical data from applications of hypothetical retrospection it has to be left as an open issue, for the time being, whether or not this procedure can also facilitate agreement in controversial decisions. However, experience from the risk management field indicates that it may have such an effect. Although there is often disagreement after an accident has happened, this is mostly not disagreement about what should have been done but about who should have done it. If this effect of actual retrospection can to some degree also be attained in hypothetical retrospection, then the procedure may contribute to more timely agreement about what should be done.

4 Moral Argumentation about Uncertainty and Probability

This and the following section are devoted to a brief outline of argument types that can be used in moral deliberation that employs the framework of hypothetical retrospection.

Uncertainty about the effects of one’s action or decision changes the moral situation and can therefore be a decisive factor in the evaluation of alternatives. Although the moral effects of uncertainty can go in different directions, there is a general tendency for uncertainty to increase the moral leeway. In other words, uncertainty widens the range of acceptable, or morally permitted, alternatives that are open to the agent. This is a simple precept that does not seem to have been stated previously as a general principle. I propose to call it *uncertainty transduction* since it consists in uncertainty being transduced from the empirical to the moral realm. For an example, consider a man who is going to buy a theatre ticket as a birthday present for his wife. There are two plays to choose between. If he knows for sure which of the plays his wife would prefer, then he will probably feel obliged to choose that play. However, if he only has a very uncertain idea about her preferences, then he may feel free to choose the play that he prefers himself. The uncertainty present in the latter case can be said to increase the moral leeway, or scope of permissible actions.

In this example, no probability estimates were involved. However, uncertainty transduction also operates on uncertainty about probabilities. Generally speaking, a moral argument that is based on a probability estimate is weakened if it can be shown that this estimate is uncertain.²³ Consider a physician’s choice between two drugs that she can prescribe for a patient. According to the best available estimates of an expert committee, the probability that drug A will cure the patient is 70%, whereas for drug B the corresponding probability is 80%. If these probabilities are known to be well-founded and highly relevant for the patient in question, then a good case can be made that the physician is morally

²³ This is so even if the uncertainty is symmetric in the sense that it gives no reason to either raise or lower the estimate.

required to choose drug B.²⁴ On the other hand, if these estimates are uncertain, then the physician has a much better claim to moral leeway in the choice between the two preparations. (I leave it to the reader to spell out this example in detail in terms of hypothetical retrospection.) In a clinical trial, patients are randomized between treatments, typically between an experimental treatment and one with reasonably well-known properties. According to an influential account of the ethics of clinical trials, what makes this procedure morally permissible is the prevailing uncertainty about which of the treatments is in fact best for the patient.²⁵

Although moral argumentation in the present framework is not governed by probability weights, it may include arguments that refer to probabilities or degrees of plausibility. Suppose that we consider taking a measure that is useful only in one particular branch of future development. If this branch can be shown to be very likely, then this is a valid argument in favour of the measure. Conversely, if the branch can be shown to be unlikely, then this is a valid argument against the measure. Clearly, such argumentation may or may not refer to quantitative (estimates of) probability. Some possible events are so uncertain that no meaningful probability assessments can be made, but non-numerical comparisons may nevertheless be serviceable.²⁶

In cases when expected utilities can be meaningfully calculated, they are often important decision-guides. This implies in particular to decisions in which a major objective is to maximize a specified type of (aggregated) outcome. Consider a decision whether or not to make seat belts obligatory in a certain country. The expected number of deaths in traffic accidents is 300 per year if safety belts are compulsory and 400 per year if they are optional. Suppose that our aim is to reduce the total number of traffic casualties. Then, if these numerical estimates are accurate, there is a safe way to achieve this aim. Due to the law of large numbers, this can be done by choosing the alternative with the highest expected utility (lowest expected number of deaths). In hypothetical retrospection from the viewpoint of having chosen not to make seat belts compulsory we will have to face the fact that about 100 more persons have died than if the other choice had been made. Under the given assumption about the policy aim, this is a strong argument in favour of compulsory seat belts.²⁷ However, the validity of this argument depends on the large number of road accidents, that levels out random effects in the long run. The same type of argument cannot be used for case-by-case decisions on unique or very rare events.²⁸

5 Risk and Personal Relations

Relations between persons have a central role in most moral discussions, including those that refer to risk and uncertainty. I will not attempt here to cover the full range of interpersonal ethical issues that arise when risks are taken or imposed. Instead, I will focus on one fundamental problem, namely the extent to which a risk for one person can be

²⁴ For simplicity we may assume that there are no other differences, such as side effects, that can influence the decision.

²⁵ For a discussion of this argumentation, see Hansson (2006b).

²⁶ See Hansson (2004b) for some examples of this.

²⁷ The most common argument against compulsory seat belts is that such legislation is paternalistic. On paternalism and anti-paternalism in risk-related issues, see Hansson (2005b).

²⁸ See Hansson (1993). – This also illustrates two essential differences between the present framework and that of expected utility maximization. The latter but not the former is committed to (1) assign exact probabilities to all possible events and (2) use these probabilities as weights in all decisions to be made.

outweighed by a benefit for another person. This is of course a variant of the more general issue whether a disadvantage to one person can be outweighed by an advantage to another person, but the risk variant has specific features that call for our attention.

Everyday moral reasoning does not in general allow gains for one person to cancel out losses for another. I am not allowed to inflict even a minor loss on you against your wish in order to achieve a larger gain for myself or for some third person. This can be expressed as a prima facie no-damage principle: *Every person has a prima facie moral right not to be exposed to negative impact, such as damage to her health or her property, through the actions of others.* Since this is a prima facie right, it can be overridden, but we tend to give it up only reluctantly and to require strong reasons for doing so.

Standard risk-benefit analysis has a radically different approach. This discipline proceeds by weighing the sum of all individual risks (i.e. probability-weighted negative outcomes) against the sum of all individual benefits. An option is acceptable if the first sum is smaller than the latter. In such a calculation, it makes no difference to whom a risk or a benefit accrues. Therefore, one person's benefits can easily cancel out another person's losses. (Hansson 2004a, 2006a) This feature of risk-benefit analysis is difficult to defend, since programmatic disregard for persons and their relations is equally implausible in situations of risk as in deterministic situations. Hence, it makes a significant moral difference if it is my own life or that of somebody else that I risk in order to earn a fortune for myself.

Hypothetical retrospection induces us to consider risks in terms of the concrete situations that will evolve if the risk eventuates. In such hypothetical situations, persons and interpersonal relations should ideally be as concrete as they are in actual situations. Therefore, hypothetical retrospection provides an impetus to treat the personal aspects of risk in accordance with how they are dealt with in ordinary moral reasoning, rather than in the same way as in risk-benefit analysis. This approach can be codified in the form of a prima facie no-risk principle that is an extension of the no-damage principle for deterministic situations: *Everyone has a prima facie moral right not to be exposed to risk of negative impact, such as damage to her health or her property, through the actions of others.*

Since this is a prima facie principle, it can be overridden. Arguably, it has to be overridden in many more situations than the no-damage principle. Social life would be impossible if we were not allowed to expose each other to certain risks: your life as a pedestrian is riskier since I drive a car in the town where you live, the smoke from your chimney contributes to my risk of respiratory disease, etc. In order to make the prima facie no-risk principle workable, we need a normatively reasonable account of the overriding considerations in view of which these and similar risk impositions can be accepted.

In risk analysis, the acceptability of these risk impositions is accounted for in impersonal terms. The two most common ways to approach this issue is to argue that these risks are too small to be worried about and that they are outweighed by greater social benefits.²⁹ Both these approaches have the disadvantage of depersonalizing risks, i.e. treating them as artificially detached from the persons exposed to them. In the present framework we need a solution that takes persons seriously.

One approach with some intuitive appeal is the "single owner heuristics" that is used in some legal contexts. Its basic idea is that in negligence cases, "due care is the care an average, reasonable person takes in his or her own person and property." Therefore, "the inquiry reduces to whether the average person would take the precaution if he or she bore

²⁹ On the former approach, see Bicevskis (1982) and Peterson (2002).

both the costs and benefits in full.” (Gilles 1994, pp 1020 and 1035) This criterion has the advantage of providing a simple litmus test for risk impositions: If someone exposes someone else to a risk to which most people would not expose themselves, then a warning flag should be raised. However, this seems to be a too lax criterion of acceptability. Individual persons may be more risk-averse than the average person, either in general or with respect to some specific type of risks. In a society that values autonomy, such preferences should not be infringed without good reason. The fact that most people are willing to accept a certain risk does not entitle us to expose persons to this risk who are not willing to accept it.

Fortunately, a more promising solution to the problem is available. To explain why most everyday risks are accepted, we may appeal to reciprocal exchanges of risks and benefits. Each of us takes risks in order to obtain benefits for ourselves. It is beneficial for all of us to extend this practice to mutual exchanges of risks and benefits. Hence, if others are allowed to drive a car, exposing me to certain risks, then in exchange I am allowed to drive a car and expose them to the corresponding risks. This (we may suppose) is to the benefit of all of us. In order to reap the advantages of modern society with its division of labour and its complex production chains, we also need to apply this principle to exchanges of different types of risks and benefits. (A person who does not travel by car may then have to accept the negative effects of other people’s car-driving in exchange for instance for his use of a wood stove that gives rise to much more unhealthy emissions than the stoves that others use.) We can then regard exposure of a person to a risk as acceptable if it is part of a social system of risk-taking that works to her own advantage.³⁰ In hypothetical retrospection, this criterion emerges rather directly since we do not compare risks per se but different courses of action that give rise to different patterns of risks. The socially realistic alternatives in which others are not allowed to drive a car in my vicinity are all alternatives in which my own car-driving is equally restricted. Hence the mutual benefits and the mutual risks that are associated with each other in the social system will be treated together, which goes a long way to explain why such risks are accepted.

References

- Bell DE (1982) Regret in decision making under uncertainty. *Oper Res* 30:961–981
 Bicevskis A (1982) Unacceptability of acceptable risk. *Search* 13(1–2):31–34
 Dickenson D (1991) Moral luck in medical ethics and practical politics. Avebury, Aldershot
 Gilles SG (1994) The invisible hand formula. *Va Law Rev* 80:1015–1054
 Hansson SO (1993) The false promises of risk analysis. *Ratio* 6:16–26
 Hansson SO (1999) But what should I do? *Philosophia* 27:433–440
 Hansson SO (2001) The modes of value. *Philos Stud* 104:33–46
 Hansson SO (2003) Ethical criteria of risk acceptance. *Erkenntnis* 59:291–309
 Hansson SO (2004a) Weighing risks and benefits. *Topoi* 23:145–152
 Hansson SO (2004b) Great uncertainty about small things. *Techne* 8(2)
 Hansson SO (2005a) Seven myths of risk. *Risk Manage* 7(2):7–17
 Hansson SO (2005b) Extended antipaternalism. *J Med Ethics* 31:97–100
 Hansson SO (2006a) Economic (ir)rationality in risk analysis. *Econ Philos* 22:231–241
 Hansson SO (2006b) Uncertainty and the ethics of clinical trials. *Theor Med Bioethics* 27:149–167
 Humberstone IL (1980) You’ll regret it. *Analysis* 40:175–176
 Jeffrey RC (1983) *The logic of decision*, 2nd edn. University of Chicago Press, Chicago, IL
 Kahneman D, Tversky A ([1979] 1988) Prospect theory: an analysis of decision under risk. In: Gärdenfors P,

³⁰ For details, see Hansson (2003).

- Sahlin N-E (eds) *Decision, probability, and utility: selected readings*. Cambridge University Press, Cambridge, UK, pp 183–214
- Lackey D (1976) Empirical disconfirmation and ethical counter-example. *J Value Inq* 10:30–34
- Ladmores C (1999) The idea of a life plan. *Soc Philos Policy* 16:96–112
- Loomes G, Sugden R (1982) Regret theory: an alternative theory of rational choice under uncertainty. *Econ J* 92:805–824
- Luce RD, Raiffa H (1957) *Games and decisions: introduction and critical survey*. Wiley, New York
- Lucey KG (1976) Counter-examples and borderline cases. *Personalist* 57:351–355
- Mulvaney RJ (1992) Wisdom, time, and avarice in St Thomas Aquinas's treatise on prudence. *Mod Schman* 69:443–462
- Nagel T (1986) *The view from nowhere*. Oxford University Press, New York
- Peterson M (2002) What is a de minimis risk? *Risk Manage* 4:47–55
- Rawls J (1972) *A theory of justice*. Oxford University Press, Oxford
- Sugden R (1985) Regret, recrimination and rationality. *Theory Decis* 19:77–99
- Vanden Houten A (2002) Prudence in Hobbes's political philosophy. *Hist Polit Thought* 23:266–287
- Ward DE (1995) Imaginary scenarios, black boxes and philosophical method. *Erkenntnis* 43:181–198
- Weber M (1998) The resilience of the Allais paradox. *Ethics* 109:94–118
- Weirich P (1981) A bias of rationality. *Australas J Philos* 59:31–37
- Williams BAO (1976) Moral luck. *Proc Aristot Soc, Suppl Vol* 50:115–135
- Wit E-JC (1997) *The ethics of chance*. PhD thesis, Pennsylvania State University