Economics and Philosophy, **22** (2006) 231–241 Copyright © Cambridge University Press
doi:10.1017/S0266267106000885

ECONOMIC (IR)RATIONALITY IN RISK ANALYSIS

Sven Ove Hansson
Royal Institute of Technology, Stockholm

Mainstream risk analysis deviates in at least two important respects from the rationality ideal of mainstream economics. First, expected utility maximization is not applied in a consistent way. It is applied to endodoxastic uncertainty, i.e. the uncertainty (or risk) expressed in a risk assessment, but in many cases not to metadoxastic uncertainty, i.e. uncertainty about which of several competing assessments is correct. Instead, a common approach to metadoxastic uncertainty is to only take the most plausible assessment into account. This will typically lead to risk-prone deviations from risk-neutrality. Secondly, risks and benefits for different persons are added to form a total value of risk. Such calculations are used to support the view that one should accept being exposed to a risk if it brings greater benefits for others. This is in stark contrast to modern Paretian welfare economics, that refrains from interindividual comparisons and does not require people to accept a disadvantage because it brings a larger advantage for others.

1. INTRODUCTION

Some of the earliest studies in risk analysis took the form of attempts to determine a general limit between acceptable and non-acceptable levels of risk. Such a limit should be common to different types of risk factors and different types of situations.[1] However, it was soon realized that a risk has to be assessed in relation to the benefits with which it is associated. Today, the standard approach is to weigh risks against benefits.[2] Risk

[1] These proposals have survived in the form of *de minimis* risk (Rechard 1999: 783–5; Peterson 2002).

[2] The terminology is somewhat misleading, since "risk" denotes disadvantages that may or may not materialize, whereas "benefit" denotes advantages of which one is sure. It would

https://doi.org/10.1017/S0266267106000885 Published online by Cambridge University Press

analysis has developed into a discipline that strives for rational solutions to decision problems, in much the same way as mainstream economics. In the subdiscipline of risk-benefit analysis, the overarching goal is to achieve standard economic rationality in the management of risks.

However, risks to human health and the environment are very different entities from those commonly treated in economics. Two major conversions are needed to squeeze risks into an economic framework, so that they can be weighed against economic benefits. First, the various types of outcomes referred to in risk analysis – from human deaths to the extinction of a crustacean species – are valued in monetary terms. Secondly, uncertain outcomes are given weight in proportion to their probabilities, in accordance with expected utility theory. A proponent of this methodology recently motivated it with explicit reference to the criterion of rationality:

> The only meaningful way to evaluate the riskiness of a technology is through probabilistic risk analysis (PRA). A PRA gives an estimate of the number of expected health impacts – e.g., the number of induced deaths – of the technology, which then allows comparisons to be made with the health impacts of competing technologies so a rational judgment can be made of their relative acceptability. No only is that procedure attractive from the standpoint of scientific logic, but it is easily understood by the public. (Cohen 2003: 909)

Both these conversions conform with standard assumptions in mainstream economics, but they have nevertheless been subject to criticism. Critics claim that negative outcomes such as the loss of human lives cannot be measured in terms of money. (Ashby 1980; Baram 1981; Kelman 1981). Stuart Hampshire (1972: 9) has warned that the habits of mind engendered by this type of impersonal computation may lead to "a coarseness and grossness of moral feeling, a blunting of sensibility, and a suppression of individual discrimination and gentleness." Another common criticism is that serious events with low probabilities should be given a higher weight than what they are given in the expected utility model (Burgos and Defeo 2004). In the present contribution, neither of these lines of criticism will be followed. To the contrary, it will be assumed for the sake of argument that all outcomes can be adequately valued in monetary terms and that expected utility maximization is the rational approach to uncertain outcomes. The purpose of the article is to appraise to what extent common practices in risk analysis (including risk-benefit analysis) actually conform with standard economic rationality. The analysis will focus on two crucial issues, namely how risk analysis conforms with the ideal of consistently maximizing expected utility (Section 2) and how its

be better to compare risks vs. chances, or benefits vs. harms. However, the established terminology will be used here.

https://doi.org/10.1017/S0266267106000885 Published online by Cambridge University Press

approach to interindividual comparisons relates to the dominating view of such comparisons in modern economics (Section 3).

2. HOW RISK-NEUTRAL IS RISK ANALYSIS?

Expected utility maximization is generally taken to be the *risk-neutral* decision rule. Decision rules that put more emphasis on avoiding the worst possible outcomes are called *risk-averse*, and those deviating in the opposite direction may be called *risk-prone*. Proponents of standard risk analysis accuse alternative, precautionary approaches of being unduly risk-averse (Charnley 1999; 2000; Durodié 2003). It is generally assumed in these discussions that probabilistic risk (and risk-benefit) analysis in its common form is risk-neutral since it employs expected utility maximization. However, as we will soon see, it is in fact often risk-prone due to the way in which it treats underlying uncertainties that are not explicitly represented in the formal analysis.

In order to see this we need to make the distinction between *endodoxastic* and *metadoxastic* uncertainty. Endodoxastic uncertainty is the uncertainty expressed or inherent in an assessment.[3] It can be communicated by the assessor either in probabilistic or non-probabilistic terms. Metadoxastic uncertainty is uncertainty about whether or not an assessment is correct, or about the choice between alternative assessments. The distinction between endo- and metadoxastic uncertainty is related to that between first- and second-order probability, but it is more general since neither type of uncertainty needs to be expressed in probabilistic terms.[4]

A study of actual practices in risk analysis, for instance as they are reported in the leading journal *Risk Analysis*, will confirm that expected utility maximization is the standard approach to endodoxastic uncertainty. Since this decision rule requires exact probabilities, it provides a strong

[3] More generally speaking: in a state of belief.

[4] Ever since Hume ([1739] 1888: 182–3) the distinction between first- and second-order probabilities has been eschewed by many philosophers who have feared that it may give rise to infinite regress. The regress argument can easily be blocked if the probabilities on the two levels are assumed to be different in nature, hence first-order probabilities can be categorized as objective and second-order probabilities as subjective. However, this solution may be too restrictive. As was noted by Bryan Skyrms (1980: 109), it is "hardly in dispute that people have beliefs about their beliefs," and this is also applicable to those beliefs that can be expressed as subjective probabilities. Therefore a formal model that mirrors certain aspects of actual reasoning will have to contain subjective probability assignments that refer to subjective probability assignments. Cognitive limitations preclude infinite regress in actual subjects. Consequently a formal model that mirrors actual reasoning should not contain an infinite number of levels; it is in fact difficult to find cases in which more than two levels of subjective probabilities are needed. The same argument is applicable to non-probabilistic uncertainty.

https://doi.org/10.1017/S0266267106000885 Published online by Cambridge University Press

impetus to assign probabilities to potential negative outcomes. This is a major reason why expert judgments of risks are often expressed in exact probabilistic terms, even when it takes considerable effort to transform uncertain knowledge into that format. In contrast, metadoxastic uncertainty is often dealt with by selecting or developing the most credible assessment, and then relying on that assessment in the analysis.[5] As a consequence, other possible assessments with lower but non-negligible credibility will have no influence on the calculations.[6]

For a simple example of this, consider a decision whether or not to install exhaust ventilation in a factory in order to reduce exposure to a solvent that is widely used in the process. The cost of the ventilation system is €200 000. The only known adverse effects (risks) associated with the solvent are headaches and irritation of the mucous membranes that are experienced by some of the workers. The total disvalue of these effects is estimated at €100 000. (Remember that for the sake of the argument, we have accepted the common practice in risk-benefit analysis to value health effects in monetary terms.) According to standard risk-benefit analysis, the ventilation system should not be installed.

However, the health effects of this solvent are insufficiently investigated. If it has unknown serious effects, then the disvalue of unchanged exposure conditions is much higher than €100 000. For any positive number x, a probability of at least $1/x$ that the substance has an additional health effect with a disvalue at least x times that of the known effects (headache and irritation) is sufficient to motivate the ventilation system, according to expected utility maximization. Judging by the experience we have from other substances, this may very well be the case for instance for $x = 20$. This example shows that the standard approach in risk-benefit analysis, in which the application of expected utility maximization is restricted to endodoxastic uncertainty, may yield a recommendation that

[5] More sophisticated approaches, that take into account the uncertainty deriving from the choice between different models and assessments, can be found in risk assessments performed in the nuclear sector, but less often in other areas such as the risk assessment of chemicals (Rechard 1999: 766, 781).

[6] Another way to express this is that metadoxastic uncertainty is often *de facto* dealt with in accordance with the maxiprobability rule, which consists in disregarding all alternatives except the one that has the highest probability. This is not a decision rule that has been proposed by decision-theorists, but rather a method that decision-makers can be observed making use of (Hansson and Johannesson 1997; Hansson 1999). The epistemic requirements of the maxiprobability rule are less demanding than those of expected utility maximization. Thus, in order to apply the maxiprobability rule to a case of metadoxastic uncertainty between different sources of information we only need to know which is the most reliable source, and use the information obtained from it. In order to maximize expected utility in such a case, we will have to integrate the information received from the different sources, assigning to each of them weight according to the estimated probability that the information it provides is correct.

https://doi.org/10.1017/S0266267106000885 Published online by Cambridge University Press

differs from that which would follow from consistent application of this decision rule to both endo- and metadoxastic uncertainty. Obviously, in this and many other cases, metadoxastic uncertainties are extremely difficult to quantify in probabilistic terms, but so are some endodoxastic uncertainties. In an expected utility framework, the fact that a probability is difficult to estimate is not sufficient reason to treat it as equal to zero.

For another example, consider one of the most controversial environmental issues, namely global climate change. The near-consensus of climate scientists is represented by the reports of the IPCC, according to which the Earth's temperature is slowly increasing due to the emission of greenhouse gases. There is however a possibility that the IPCC's conclusions are too pessimistic. Possibly, the greenhouse effect is completely counterbalanced by other mechanisms, so that no global warming disaster is approaching. On the other hand, there is also a possibility that global warming will accelerate much faster than what the IPCC predicts. For concreteness, suppose that probabilities and utilities are as follows:

Risk assessment	Probability that the risk assessment is correct	Utility of climate effects if the risk assessment is correct
A	.01	−10000
B	.04	−1000
C	.90	−100
D	.05	0

Here, A and B represent appraisals that are much more pessimistic than those of the IPCC, C the IPCC appraisal, and D the most optimistic appraisal. Under the assumptions given in the table, reliance on the best available assessment (the standard method) will lead to an estimate of the utility of continued greenhouse emissions at −100, whereas consistent application of the expected utility method to both endo- and metadoxastic uncertainty will lead to an estimate of it at −230. In most controversies about the greenhouse effect, the former approach has implicitly been taken for granted. The contested issue is which is the most plausible assessment, to which this method should be applied (McCright and Dunlap 2000).

Metadoxastic uncertainty is almost always present in risk analysis, but often it is not explictly treated. As we have seen, it is a common practice to base decisions exclusively on the best available expert assessment. This approach is problematic from a decision-theoretical point of view. There is ample historical evidence of mistakes made by experts on risks. A rational

https://doi.org/10.1017/S0266267106000885 Published online by Cambridge University Press

decision-maker should take into account the possibility that such mistakes may happen again. In an expected utility framework, this means that scenarios deviating from the majority view of experts should be entered into the analysis, and that such scenarios should be given weight according to how plausible they are estimated to be.[7]

Metadoxastic uncertainty is of course prevalent in other types of decisions than those concerned with health-related and environmental risks. However, in many other applications there is symmetry between possible errors in different directions. Hence an estimate of the financial outcome for a company of a particular investment may turn out to be €10 million too high or it may turn out to be €10 million too low, and in many cases the two types of mistakes seem to be equally probable. If there is perfect symmetry, i.e. over- and underestimates of the same size have equal probabilities, then it makes no difference if we disregard them both and base our analysis exclusively on the central, most probable estimate.

In applications to health-related and environmental risks, such symmetry is rare. Consider, for instance, an estimate that the probability of an explosion in a certain pressure vessel in the next year is 10^{-5}. This probability may be 2×10^{-5} too low (i.e. the correct value may be 3×10^{-5}), but it cannot be 2×10^{-5} too high (since it cannot be negative). Due to this asymmetry, a risk-benefit analysis based exclusively on the central, most probable estimate can be expected to be more risk-prone than the "risk-neutral" ideal of consistently maximized expected utility.

Perhaps surprisingly, the risk-neutral ideal may be better approximated by decision-rules constructed to comply with the precautionary principle. According to that principle, unproven but scientifically credible dangers should be taken into account in decision-making (Sandin 1999; Hansson 1999). If such dangers are assigned weights in accordance with their estimated probabilities, then this variant of the precautionary principle seems to be closer than standard risk analysis to a risk-neutral treatment of metadoxastic uncertainty.

[7] The determination of such metadoxastic probability assignments is a difficult task, but as mentioned above it cannot be avoided in an expected utility framework. (Assigning maximal probability to one assessment is of course a way to do it, not a way to avoid doing it.) One possible source of metadoxastic probabilities is systematized experience from previous experts' assessments (Hansson 2001). Another possible source of metadoxastic probabilities is the distribution of assessments in a representative group of experts. It could for instance be hypothesized that the distribution of experts' point estimates over the set of possible assessments coincides with a reasonable distribution of probabilities that these estimates are correct. However, if the point estimates of the experts gravitate towards the middle, then the average of these point estimates will tend to be misleading through the same mechanism that operates when a single point estimate is relied upon. Other ways to aggregate assessments by different experts, that require each of them to report a probability distribution, rather than only a point estimate, may be more adequate (Hansson 1993).

https://doi.org/10.1017/S0266267106000885 Published online by Cambridge University Press

3. INTERPERSONAL AGGREGATION IN RISK ANALYSIS

With respect to interindividual comparisons, standard risk analysis follows the principles of classical utilitarianism. Thus, all risks are taken to be fully comparable and additively aggregable. In risk-benefit analysis, benefits are added in the same way, and finally the sum of benefits is compared to the sum of risks in order to determine whether the total effect is positive or negative. The practice of adding all risks and all benefits has immediate intuitive appeal. Of course we should weigh all benefits against all risks.

Or should we? Some reflection will show that this type of calculation is not as unproblematic as it may seem at first sight. Arguably, its most problematic feature is a far-reaching disregard for persons.[8] In the total calculations of risk analysis, a disadvantage to one person can always be compensated by an equally sized advantage to another person. Just as in classical utilitarianism, individuals have no other role than as carriers of utilities and disutilities, the values of which are independent of whom they are carried by. Such a framework excludes from consideration many types of moral criteria that have considerable intuitive appeal, such as the distinction between taking and being exposed to a risk, voluntary vs. involuntary risk exposure, and fairness in the distribution of risks and benefits (Hansson 2003).

The calculative framework of modern risk analysis is largely based on the methodology used in the influential *Reactor Safety Study* (WASH-1400, the Rasmussen report) from 1975. The definition of risk introduced in this study, namely as frequency × magnitude, is the origin of the common notion among risk analysts that risk equals probability times consequence (Rechard 1999: 776). The calculation of the total risk was based on the identification, through fault tree analysis, of various possible accidents. For each of these, a calculation was made of its "risk" in the technical sense just referred to. The total risk associated with reactor accidents was obtained by adding up all these different risks.

This methodology can also be exemplified by the several risk assessments that have been made of the transportation of nuclear material on roads or rails. In such studies, the radiological risks associated with normal handling and with various types of possible accidents are quantified, and so are non-radiological risks including fatalities caused by accidents and vehicle exhaust emissions. As noted by the authors of an unusually careful emissions study, "[t]he end result of all estimates are unit

[8] More precisely, the standard forms of utilitarianism and risk analysis are person-*regarding* in the (weak) sense that the advantages and disadvantages referred to in these theories are advantages and disadvantages to persons. These theories are person-*disregarding* in the sense that they disregard the identity of the person to which an advantage or disadvantage accrues, and consequently also how advantages and disadvantages are distributed among persons.

https://doi.org/10.1017/S0266267106000885 Published online by Cambridge University Press

risk factors for use in transportation risk assessment. The unit risk factors are in the form of latent cancer fatalities per kilometer (LCFs/km) or latent fatalities per kilometer or travel by truck or railcar" (Biwer and Butler 1999: 1158). The risk of a given shipment is then obtained by multiplying the distance travelled by this unit risk factor. Obviously, these calculations will provide us with the total number of (statistically expected) deaths. They do not make it possible to infer the risk for individual persons, such as the workers performing the transport, other road users, or residents in the areas close to the road.

An obvious alternative to this utilitarian approach is to treat each individual as a separate moral unit. Then risks and benefits pertaining to one and the same person can be weighed against each other, whereas risks and benefits for different persons are treated as incomparable. This type of risk-weighing can be found in certain applications that are strongly influenced by the individual-centered traditions of clinical medicine. Dietary advice is one example. Due to environmental contamination, health authorities recommend limits in the consumption of fish caught in certain waters. Such recommendations are based on endeavours to balance the negative health effects of the contaminants against the positive effects of fish as a constituent of the diet (Knuth *et al.* 2003). This balance is struck separately for each individual; hence positive effects for others (such as the fishing industry) are excluded from consideration.

The ethical appraisal of clinical trials is another example. It is an almost universally accepted principle in research ethics that a patient should not be included in a clinical trial unless there is genuine uncertainty on whether or not participation in the trial is better for her than the standard treatment that she would otherwise receive. That her participation is beneficial for others (such as future patients) cannot outweigh a negative net effect on her own health; in other words her participation has to be supported by an appraisal that is restricted to risks and benefits for herself (London 2001; Hansson 2004, 2006).

Hence, there are two traditions in the weighing of risks and benefits that differ in their approaches to interindividual comparisons. This issue is also the dividing line between the "old" and "new" schools of welfare economics. Mainstream risk analysis adheres to the same utilitarian approach, with full interindividual comparability and aggregation, as Arthur Pigou's so-called old welfare economics. The more individualist approach that was exemplified above with dietary advice and clinical trials shares the approach of the so-called new school in welfare economics. The latter school has dominated mainstream economics since the 1930s, when Lionel Robbins showed how economic analysis can dispense with interpersonal comparability.

Mainstream risk analysis and mainstream economics represent two extremes with respect to interindividual comparisons. The aggregations

https://doi.org/10.1017/S0266267106000885 Published online by Cambridge University Press

of total risk that are performed routinely in risk analysis stand in stark contrast to the consistent avoidance of interindividual comparisons that is a major guiding principle in modern economics. This difference also has repercussions in the ideological uses of the respective disciplines. It is an implicit message of risk-benefit analysis that a rational person should accept being exposed to a risk if this brings greater benefits for others. The implicit message of Paretian welfare economics is much more appreciative of self-interested behaviour.[9]

It could perhaps be argued that the different approaches in economics and risk analysis are motivated by differences in the interindividual comparability of the outcomes treated in the two disciplines. A reasonable argument could be made that the disvalue of disease or death differs less between persons than the marginal value of money. If this is true, then total aggregation in accordance with utilitarian principles can be performed with less uncertainty in risk analysis than in economics generally. However, even if this is granted, it does not follow that total aggregation is morally on safer ground in risk analysis than in economics. The crucial issue is not whether or not a benefit (or harm) for one person can be said to be greater than a benefit (or harm) for another person. The crucial issue is whether or not an action that brings about the greater benefit (smaller harm) compensates for the non-realization of a smaller benefit (greater harm) accruing to another person. Such interpersonal compensability does not follow automatically from interpersonal comparability (Hansson 2004). The fact that it is worse for you to lose your thumb than for me to lose my little finger does not in itself make it allowable for you or anybody else to chop off my little finger in order to save your thumb.

Mainstream economics and risk analysis thus produce two extreme answers to the question whether an advantage for one person can morally compensate for a disadvantage for another person, namely the answers "never," respectively "if and only if the advantage is larger than the disadvantage." Both these answers have implausible implications (irrespective of whether we are dealing with risks or with economic assets). Some sort of intermediate, but yet principled answer to this question seems to be needed.

As I have argued more extensively elsewhere (Hansson 2003), a plausible answer for the case of risks can be based on a system of prima facie rights. 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

[9] The choice of this (Paretian) approach to welfare economics is largely based on scepticism with regard to the meaningfulness of interindividual comparisons, rather than on considerations of fairness.

https://doi.org/10.1017/S0266267106000885 Published online by Cambridge University Press

right not to be damaged by others. This right can be extended to a prima facie right not to be put at risk by others: *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 right, it can be overridden. In practice, it has to be overridden in quite a few cases, since social life would be impossible if we were not allowed to expose each other to certain risks. To mention just one example: as car drivers we put each other's lives at risk. In order to make the prima facie right to risk avoidance workable, we need a normatively reasonable account of the overriding considerations in view of which these and similar risk impositions can be accepted. This can be obtained by 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 deal with the complexities of modern society, we also need to apply this principle to exchanges of different types of risks and benefits. 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 advantage.

This standpoint has the important advantage of recognizing each person's individual rights, but still making mutually beneficial adjustments possible. It does not allow you to stop your neighbor from driving a car, although this implies a (small) risk for you as a pedestrian, since this method of transportation is also accessible to you, and beneficial to you as well. On the other hand it would prohibit exploitative arrangements in which someone is exposed to risks in order to achieve benefits only for others. Clearly, this approach has its own problems, and is in need of detailed development. For our present purposes, the important conclusion is that a plausible solution will have to include distinctions and considerations that are absent both from standard risk analysis and from standard economic theory.

4. CONCLUSION

We have found that the common forms of risk analysis and risk-benefit analysis deviate in two important ways from the rationality assumptions of modern mainstream economics. First, although risk analysis applies expected utility maximization fairly consistently to endodoxastic uncertainties, its approach to metadoxastic uncertainty may often lead to risk-prone deviations from risk-neutrality. Secondly, risk (and risk-benefit) analysis assumes full comparability and aggregability of risks (and benefits), which contradicts the standard assumptions of modern welfare economics. Contrary to the latter discipline, mainstream risk

https://doi.org/10.1017/S0266267106000885 Published online by Cambridge University Press

analysis provides support for the idea that a rational person should accept a disadvantage for herself if it brings a greater advantage to others or to society as a whole. These are important differences that need to be carefully considered, in particular when economics and risk analysis are combined in studies intended to guide policy decisions.

REFERENCES

Ashby, E. 1980. What price the Furbish lousewort? *Environmental Science & Technology* 14:1176–81

Baram, M. S. 1981. The use of cost-benefit analysis in regulatory decision-making is proving harmful to public health. *Annals of the New York Academy of Sciences* 363:123–8

Biwer, B. M. and J. P. Butler 1999. Vehicle emission unit risk factors for transportation risk assessments. *Risk Analysis* 19:1157–71

Burgos, R. and O. Defeo 2004. Long-term population structure, mortality and modeling of a tropical multi-fleet fishery: the red grouper Epinephelus morio of the Campeche Bank, Gulf of Mexico. *Fisheries Research* 66:325–35

Charnley, G. 1999. President's Message, *RISK Newsletter* 19(2):2

Charnley, G. 2000. 1999 Annual Meeting: past president's message: risk analysis under fire. *RISK Newsletter* 20(1):3

Cohen, B. 2003. Probabilistic risk analysis for a high-level radioactive waste repository. *Risk Analysis* 23:909–15

Durodie, B. 2003. The true cost of precautionary chemicals regulation. *Risk Analysis* 23:389–98

Hampshire, S. 1972. *Morality and pessimism*. Cambridge University Press

Hansson, S. O. 1993. Entscheidungsfindung bei Uneinigkeit der Experten. In *Die Modernisierung der Demokratie, Internationale Ansätze*, ed. H. Zilleßen, P. C. Dienel, and W. Strubelt, 87–96. Westdeutscher Verlag

Hansson, S. O. 1999. Adjusting scientific practices to the precautionary principle. *Human and Ecological Risk Assessment* 5:909–21

Hansson, S. O. 2001. Choosing priority-setting criteria for carcinogens. *Human and Ecological Risk Assessment* 7:475–91

Hansson, S. O. 2003. Ethical criteria of risk acceptance. *Erkenntnis* 59:291–309

Hansson, S. O. 2004. Weighing risks and benefits. *Topoi* 23:145–52

Hansson, S. O. 2006. Uncertainty and the Ethics of Clinical Trials. *Theoretical Medicine and Bioethics*, in press.

Hansson, S. O. and M. Johannesson 1997. Decision-theoretic approaches to global climate change. In *International politics of climate change*, ed. Gunnar Fermann, 153–78. Scandinavian University Press

Hume, D. [1739] 1888. *A treatise of human nature*, ed. LA Selby-Bigge. Clarendon

Kelman, S. 1981. Cost-benefit analysis: an ethical critique. *Regulation* 5:33–40

Knuth, B. A., N. A. Connelly, J. Sheeshka, and J. Patterson 2003. Weighing health benefit and health risk information when consuming sport-caught fish. *Risk Analysis* 23:1185–97

London, A. J. 2001. Equipoise and international human-subjects research. *Bioethics* 15:312–32

McCright, A. M. and R. E. Dunlap, 2000. Challenging global warming as a social problem: an analysis of the conservative movement's counter-claims. *Social Problems* 47:499–522

Peterson, M. 2002. What is a de minimis risk? *Risk Management* 4:47–55

Rechard, R. P. 1999. Historical relationship between performance assessment for radioactive waste disposal and other types of risk assessment. *Risk Analysis* 19:763–807

Sandin, P. 1999. Dimensions of the precautionary principle, *Human and Ecological Risk Assessment* 5:889–907

Skyrms, B. 1980. Higher order degrees of belief. In *Prospects for pragmatism*, ed. D. H. Mellor. 109–37

https://doi.org/10.1017/S0266267106000885 Published online by Cambridge University Press