

ROESER, S.

Philosophy Department, Faculty of Technology, Policy and Management,
Delft University of Technology, Jaffalaan 5, 2628 BX Delft, The Netherlands,
T: +31-15-2788779, F: +31-15-2786233, E: S.Roeser@tudelft.nl, <http://www.tbm.tudelft.nl/webstaf/sabiner/>

ETHICAL INTUITIONS ABOUT RISKS

Abstract - Psychologists, sociologists and empirical decision researchers agree that concerning risks, intuitive judgments are a bad guide. However, there is an important ambiguity inherent to judgments about risks. Such judgments might concern the purely quantitative risk, but they might also concern the ethical acceptability of a risk. Whereas intuitions are a bad guide concerning the quantitative aspect of risk judgments, intuitions might be inevitable and legitimate concerning ethical risk judgments. This is indeed the position defended in this paper. To this purpose empirical findings from Paul Slovic concerning intuitions of laypeople about risks are studied closely. I will place those in the context of the philosophical theory of ethical intuitionism. This framework shows that intuitive risk judgments of laypeople express legitimate ethical concerns that should be included in methodologies for risk policy.

1. INTRODUCTION

Perhaps the most important message from the research done to date is that there is wisdom as well as error in public attitudes and perceptions. Laypeople sometimes lack certain information about hazards. However, their basic conceptualization of risk is much richer than that of experts and reflects legitimate concerns that are typically omitted from expert risk assessments (Slovic 2000, 191).

Psychologists, sociologists and empirical decision researchers agree that concerning risks, intuitive judgments are a bad guide. However, there is an important ambiguity inherent to judgments about risks. Such judgments might concern the purely quantitative risk, but they might also concern the moral acceptability of a risk. Whereas intuitions are a bad guide concerning the quantitative aspect of risk judgments, intuitions might be inevitable and legitimate concerning normative, ethical¹ risk judgments. This is indeed the position I will defend in this paper.

2. INTUITIONS CONCERNING RISKS

There is a vast amount of literature on the fallibility of intuitive judgments concerning risks. Most famously, Tversky and Kahnemann (1974) have shown how not only lay people but also experts make judgments concerning risks which divert significantly from judgments based on normative decision theories such as rational choice theory. Psychologists, sociologists and empirical decision researchers are in agreement that if it comes to decisions concerning risk, intuitive judgments are an unreliable guide. Although it is admitted that such intuitive judgments might have *pragmatic* value in allowing us to navigate through a complex world (in that sense they are *heuristics*), we should not expect them to *correctly* represent probabilities or relations between probabilities (in that sense they are distortions or *biases*; cf. the title of Gilovich et al. (2002): *Intuitive Judgment: Heuristics and Biases*).

Examples of how intuitive judgments concerning probabilities can lead us astray are the following:

1. 'the law of small numbers': overgeneralization on basis of small samples;

¹ I will use the notions 'ethical' and 'moral' interchangeably.

2. judgments of causality and correlation: perception of relationships that one expects to see, even if they are not there;
3. availability: influence of imaginability or memorability on perception of hazards;
4. anchoring and insufficient adjustment: making use of simple starting points and adjustment mechanisms;
5. information processing shortcuts: using simple decision strategies that avoid the weighing of multiple considerations;
6. probability neglect: focusing on the worst case even if it is highly improbable;
7. loss aversion: dislike of loss from the status quo;
8. ‘benevolence’ of nature versus ‘suspect’ man-made processes;
9. system neglect: inability to see that risks are part of systems, where intervention into a system can create risks;
10. ‘framing’: the way information is presented influences the choices people make.²

There are large amounts of studies by psychologists that confirm that not only laypeople but even experts tend to misjudge probabilities, based on the above mentioned fallacies or biases.

However, there is an important ambiguity when it comes to judgments about risks. Such judgments might concern the purely quantitative risk (in terms of the probability of a given effect), but they might also concern the moral acceptability of a risk. The two are not identical (cf. for example, Fischhoff et al. (1981) and Shrader-Frechette (1991)). It is possible that in judging risks, people do not only assess the quantitative aspect of risk but at the same time make an assessment of the desirability of a risk, which is a normative issue.³ If this is true, then the fact that intuitive judgments diverge from probabilistic judgments does not need to mean that such intuitive judgments are flawed. It might mean that with intuitive judgments we try to capture more than mere probabilities can tell us.

There is evidence that risk judgments by laypeople indeed incorporate both aspects (Slovic 2000, 116). When asked to judge risk of death, the judgments of laypeople diverged significantly from those of experts (whose judgments do come close to factual annual fatalities, Slovic 2000, 113, 114). At first sight this might indicate that intuitive risk judgments of laypeople are significantly less reliable than those of experts. However, when laypeople were asked to judge the amount of annual fatalities, those judgments differed from their risk judgments and came closer to the judgments of experts and to factual annual fatalities. Apparently, for laypeople the concept of risk includes more than the amount of deaths but also qualitative concerns:

These results lead us to reject the idea that laypeople wanted to equate risk with annual fatality estimates but were inaccurate in doing so. Instead, we are led to believe that laypeople incorporate other considerations besides annual fatalities into their concept of risk (Slovic 2000, 116).

While the first (quantitative) kind of risk judgments should be made according to scientific methods, the second kind of judgment by definition requires normative considerations. The empirical data from the studies by Paul Slovic are consistent with this analysis. It is far from clear that formal decision theory provides for an ultimate standard with which to assess the latter kind of risk judgments. In the standard approach to risk assessment and risk management risk is understood as a function of probabilities and unwanted outcomes. Cost-benefit analysis is used to determine which hazardous activity is preferable to an alternative hazardous activity. As numerous sociologists and psychologists have emphasized, this methodology is inherently normative as well, for example by determining what counts as an unwanted outcome (cf. for example, Fischhoff et al. 1981, Jasanoff 1993, Slovic 1999).

Further on in this paper, the additional considerations of laypeople will be critically examined from a normative, philosophical point of view, in order to determine whether they might be justified concerns. It will be

² For discussions of and references for the various items: concerning items 1-5, cf. Slovic (2000), 21; concerning items 6-9, cf. Sunstein (2005), 35; concerning item 10, cf., for example, Tversky and Kahneman (1981).

³ Slovic refers to this ambiguity with the formulation: ‘perception of actual or acceptable risk’ (Slovic 2000, 86). ‘Actual risk’ denotes the quantitative aspect of risk, ‘acceptable risk’ denotes the normative aspect of risk.

argued that this is indeed the case. The risk intuitions of laypeople incorporate important ethical considerations that are left out in cost-benefit analysis. But first let us take a closer look at which concerns Slovic has found to play a role in laypeople's risk judgments.

3. LAYPEOPLE'S CONCERNS IN JUDGING RISKS

Paul Slovic (often together with others) has executed extensive psychological studies which spell out which kinds of concerns laypeople take into account in making risk judgments (Slovic 2000). This section discusses some interesting results of Slovic's studies which can serve as direct input for a normative analysis, which will follow further on.

Slovic and his colleagues have studied the intercorrelations between the following risk characteristics in judgments by laypeople:

1. Severity not controllable
2. Dread
3. Globally catastrophic
4. Little preventive control
5. Certain to be fatal
6. Risks and benefits inequitable
7. Catastrophic
8. Threatens future generations
9. Not easily reduced
10. Risks increasing
11. Involuntary
12. Affects me personally
13. Not observable
14. Unknown to those exposed
15. Effects immediate
16. New (unfamiliar)
17. Unknown to science
18. Many people exposed (Slovic 2000, 140).⁴

These characteristics were hypothesized by various researchers to play an important role in laypeople's risk perceptions (cf. Slovic 2000, 86). This means that the items were not just construed by Slovic and his colleagues but they were found to be important features in laypeople's risk judgments, based on previous studies by Slovic and other researchers. The reason why I pick this article by Slovic to apply a normative analysis to is that from all his many publications on this subject, this one presents the most extensive list of concerns that are important in the risk judgments of laypeople.

⁴ The table of intercorrelations between these risk characteristics is reprinted in appendix 1 of my paper. For the complete questions related to these characteristics which the subjects in this study had to answer, cf. Slovic (2000), 86, 87 and 138 and appendix 2 of my paper. Note that the numbering and labels of the characteristics in the table of intercorrelations (Slovic 2000, 140) does not correspond with the numbers and not always with the labels of the two initial studies. I have tried to track back these correspondences, but I might be mistaken at times where it was not clear which characteristic of the initial studies was referred to. This only affects the places in this article where I refer to the full text of questions related to specific characteristics.

Subjects⁵ were presented a list of 90 hazards which they had to rate with the risk characteristics given above on a scale from 1 through 7. Slovic et al. examined through factor analysis in how far items are correlated. According to Slovic et al., based on correlations between these characteristics, two clusters can be formed: 1-12 and 13-17. Slovic et al. call the clusters 'dread' and 'familiarity' respectively (cf. Slovic 2000, 141). They take characteristic 18 as separate, since it is rather independent of the others. Slovic et al. are mainly concerned with the statistical relationships between these characteristics. I would like to take a closer look at the contents of the various characteristics. A following step will then be to assess in how far these characteristics are normatively justified (section 6).

Many philosophers distinguish between 'ontology' and 'epistemology'. 'Ontology' refers to the way the world is in itself, independently of our judgments and beliefs, 'epistemology' refers to the way we believe the world to be like, i.e. to our judgments and knowledge about the world. It is interesting that many of the features in the second cluster (13-17) concern epistemic issues, i.e. lack of knowledge: characteristics 13, 14 and 17 are definitely epistemic, whereas characteristic 16 is ambiguous. The original question runs as follows: 'Are these risks new, novel ones or old, familiar ones?' 'New' sounds like an ontological notion, 'unfamiliar' could be either ontological or epistemological. 'Effects immediate' (15.) in the second cluster is highly *negatively* correlated with all the other elements in that cluster. Apparently, only remote effects qualify for lack of knowledge.

Characteristics 1-12 sound mainly ontological. However, characteristic 2, 'dread', is interesting because it introduces a further aspect, i.e. people's responses. The full question runs as follows: 'Is this a risk that people have learned to live with and can think about reasonably calmly, or is it one that people have great dread for – on the level of a gut reaction...?' (Slovic 2000, 87). Characteristics 1, 3, 4, 5, 7 and 9 clearly concern the issue of controllability and they are all rather highly correlated (ranging from .63 to .84). Characteristics 8, 10 and 11 also are related to controllability. Dread (characteristic 2) can be seen as a response to the perception of an uncontrollable risk. Indeed, characteristics 1, 3, 4 and 5 are highly intercorrelated with 2 (values at around .80), whereas 7 and 2 have a lower correlation of .66. Rather surprisingly, the correlation of characteristic 9, 'Not easily reduced', with those characteristics is only around .60.

Let us take a closer look at 'dread' (2.) and see with which other concerns it is correlated, since this might point into the direction of which concerns are taken to be especially dreadful. One might expect that 10., 'risks increasing', should be highly correlated to 'dread', but the correlation is only .63. This might be explained by the idea that also 'stable' (non-increasing) risks might be considered to be dreadful. 'Affects me personally' (12.) and 'dread' are also not very highly correlated (.64), whereas 'threatens future generations' (8.) and 'dread' have a much higher correlation, .76. This is an interesting finding, because it might indicate that people are more concerned with those who will live later on than just with their own personal interests. Alternatively, it might mean that people expect a given risk to get worse in the future and that this is what they find dreadful. Indeed, the correlation between 8 and 10 ('risks increasing') is rather high: .75. The correlation between 'effects immediate' (15.) and 'dread' is very low: .15; and the correlation between 15 and 8 is negative: -.29, which is consistent. Apparently a risk is more dreadful if it is expected to increase and to show its effects later on.

Let me note that I think that the full formulation of characteristic 2 (given before) is problematic for two reasons: firstly, what is dreadful is contrasted with what people 'have learned to live with (and can think about reasonably calmly)'. This seems to exclude the possibility that people can feel dread for risks that are common, but why should that be? Maybe this is true, but this should be something which is open for debate and research, it is surely not given. Indeed, further on in the same article we see a diagram in which risks are plotted against the axes 'dread' and 'familiarity' (which are the two main clusters of this study). Apparently, the relation between dread and familiarity is variable. Hence, it seems even stranger to presuppose that they are exclusionary. This formulation might account for some of the correlations I have discussed in the preceding paragraph. The second problem with the full formulation of characteristic 2 is that it asks the subjects to assess the reactions of others: 'is it [a risk] that people have great dread for?' Rating this item highly does not need to say anything about the degree of dread that the subject might feel. Worse, subjects might be mistaken about the amount of dread that others feel. This is an empirical, quantitative issue. In order to address this issue it would have been much more straightforward to measure the subjects' attitudes by asking them directly about their own feeling of dread rather than taking the detour of asking them about the feelings of dread of 'the public'. These features of the full

⁵ The subjects were college students, the article does not mention how large the sample was. The beginning of the article contains the following statement: 'Although the data have thus far been collected only from college students, the results appear to provide further insights into the nature of risk perception' (Slovic et al., 2000, 137). One might question in how far the results of this study can be taken to be representative. It goes beyond the scope of my paper to examine this question in any detail. In my paper I will treat the results as being representative.

formulation of characteristic 2 are at least ambiguous or even misleading and might distort the message we can get out of Slovic et al.'s study.

Remarkably, 'many people exposed' (18.) and dread only have a correlation of .04, whereas 'catastrophic' (7.) and dread have a relatively much higher correlation of .66. Apparently, it is not so much the amount of people who are exposed to a risk which laypeople find dreadful but rather the severity of the risk. Indeed, 'severity not controllable' (1., full text of the question: 'If a mishap occurs, can the damage be controlled?' Slovic 2000, 138) and dread have one of the highest correlations, .82, and 'certain to be fatal' (5.) and dread score equally high, .82. Note that also 'risks and benefits inequitable' (6.) and dread are rather highly correlated: .76.

Note that dread correlates low with the elements of the 'epistemological' second cluster. Apparently, risks about which we lack knowledge are not considered to be specifically dreadful, although one might expect the opposite. This is interesting because experts often charge laypeople to be afraid of certain technologies just because these technologies are new and laypeople lack information. This study falsifies this claim, at least concerning risks about which people are aware that they (and experts) lack information ('unknown to science', 17., and 'dread' even have a negative correlation of -.10).

Whereas in the present section I have looked at the content of the characteristics, in section 6 I will discuss in how far these characteristics can be considered as legitimate ethical concerns about risks. In order to do so, I will need a theoretical framework which I will develop in sections 4 and 5.

4. THE LEGITIMACY OF ETHICAL RISK CONSIDERATIONS

Slovic provides us with valuable empirical material. However, this material gives rise to the normative, philosophical question in how far the concerns of laypeople are justified or reasonable moral concerns. Slovic's work can be seen as an attempt to provide for a basis to include laypeople's judgments in decision procedures about risk. Initially, Slovic seemed to take scientific risk assessment as a standard against which to evaluate the assessment of laypeople. The main argument seemed to be instrumental: we have to include the concerns of citizens even if they are mistaken, if only to avoid public unrest. However, in his later work he emphasized the importance and legitimacy of the considerations of laypeople (cf. the passage at the beginning of this paper). Rather than taking the judgments of experts and scientific approaches to risk as a standard, he emphasizes that all approaches to risk involve normative considerations, such as which effects to take into account and which methods to use to assess them. Slovic conflates this with the claim that all risk judgments are inherently subjective. However, this seems to be throwing the baby away with the bathwater. Stating that all risk judgments are inherently *normative* is not the same as stating that all risk judgments are inherently *subjective*. It is a philosophically controversial claim to understand normativity as a form of subjectivity, which should be supported by philosophical arguments which Slovic does not provide. However, I don't think that it is wise to make this claim given the project that Slovic is working on, and there are good philosophical reasons for not understanding normativity as a form of subjectivity.

Understanding normativity as a form of subjective preferences (be it of individuals or of groups) boils down to normative relativism. Normative relativism states that there are no universal, objective normative standards. This is problematic because it makes criticism of other normative points of view impossible. Rather than being a liberating approach as it is seen and endorsed as by many social scientists (cf. for example Jaeger et al. 2001), it means that it leaves one normatively empty-handed. Taken literally, normative relativism implies that it does not matter what we do. All normative judgments are merely projections on a normatively blank world. However, this means that we might just as well throw dice to determine what to do or assign a dictator. Normative or ethical deliberation would not be anything beyond giving people the illusion of doing something meaningful, in the end, any answer would do. However, this seems highly counterintuitive. In normative or ethical deliberation, we feel that a lot is at stake and we try to find a correct answer. If ethics would be a form of subjectivism, whatever a person would prefer would be morally right. However, that would be a very arrogant view and it would make criticism of others impossible. Or if one thinks that normativity is constituted by cultural practices (this is also called a form of intersubjectivity, in order to distinguish it from individual preferences), then the practice of one culture could never be really, i.e. objectively, better or worse than that of another culture. In order to form ethical judgments we would then only need to ask 'what does my group think'. However, this excludes the possibility that we could criticize the predominant views in our own culture, which would lead to conservatism. Furthermore, we feel that we actually can criticize the dominant views of our culture. If normativity and ethics would be relative (either to individuals or to groups), trying to find correct moral answers would be a futile endeavor, and the critical stance inherent to ethical thinking would lose its bite. In moral philosophy, these are very common objections against subjectivism and relativism, but unfortunately, these insights have not made it into many other academic disciplines yet.

People might feel attracted to relativism because of its supposed liberating power, which is a misunderstanding as I have just argued. Another reason to endorse relativism is that objective values do not seem to fit into a scientific world view. However, this presupposes strong and controversial philosophical and metaphysical assumptions about what could be a correct world view. The idea that there are objective moral norms and values does not need to be at odds with scientific insights. Science is there to tell us how the world is in a descriptive, empirical sense, but science cannot make any claims about how the world is in a metaphysical or normative sense, even though scientific findings can be relevant for metaphysical and normative issues. Metaphysical and normative issues are by definition the domain of philosophy.

To connect this discussion back to the issue we are dealing with in this paper, judgments about risks, this means the following: the first kind of risk judgments concerns scientific aspects of the world, the second kind of risk judgments concerns normative aspects of the world (which are of course related to scientific and other factual aspects of the world). None of these domains of the world can be captured by the other kind of judgment, and none of these judgments is inferior to the other. They both have their own domain in which they are irreplaceable (for more on this issue, cf. Shrader-Frechette 1991).

5. ETHICAL INTUITIONS CONCERNING RISKS

In this section I will argue that ethical intuitionism such as defended by, for example, the philosophers Thomas Reid (1969), G.E. Moore (1988), W.D. Ross (1967, 1968), A.C. Ewing (1929) and more recently, Jonathan Dancy (1993, 2004), provides us with a theoretical framework based on which the considerations of laypeople can be understood as justified, reasonable moral intuitions. This way we can provide for a theoretical framework for Slovic's normative claims.

Ethical intuitionism states that our ethical intuitions can be justified even if we are not able to base them on further arguments. Ethical intuitions can be understood as perceptions of moral reality, i.e. the moral aspects of the world. Ethical intuitions cannot be reduced to or replaced by other kinds of considerations such as scientific truths, be it from the natural sciences or from the social sciences. Trying to replace ethics by one of the sciences means replacing normative statements by descriptive statements. However, the whole point of normative thinking is to have a critical attitude towards what is descriptively the case. Maybe it is part of human nature to be cruel, but is cruelty a morally good thing? No. Such a moral insight cannot be replaced by a descriptive way of thinking, since it is always possible to ask about a descriptive fact of the world whether it is morally good or right. This is G.E. Moore's so-called 'open question argument'. Trying to replace ethics by a descriptive discipline amounts to what Moore calls the 'naturalistic fallacy' (cf. Moore 1988). Ethical intuitionists argue that we have to take our ethical intuitions at face value. This does not mean that ethical intuitions cannot be mistaken, but they are 'innocent until proven guilty', and even if we cast certain ethical intuitions into doubt, we cannot avoid using other ethical intuitions to do so. Ethical reasoning always involves ethical intuitions, i.e. basic moral insights which we cannot replace by non-ethical insights or by further arguments.

According to ethical intuitionists, ethical intuitions are the foundations of more complex moral reasoning. Every kind of reasoning has to be based on basic beliefs or intuitions, otherwise our reasoning would be circular, or it would lead us to an infinite regress or we would need to make arbitrary assumptions (cf. Reid 1969, Alston 1993). Ethical intuitionists draw an analogy with mathematics. In mathematics we start from axioms from which we can build more complex theorems. Ethical intuitionists think that ethical intuitions function in a similar way as axioms in mathematics. They also invoke the analogy with sense perception. Just like our beliefs about the factual, empirical world are based on our sense perception, our complex ethical beliefs are based on basic moral beliefs or intuitions. Basic moral beliefs are 'self-evident': this means that they are evident in themselves, not by the provision of further arguments; they are 'non-inferential' (cf. Ewing 1929). Note that the claim that a belief is 'self-evident' does not mean to suggest that a belief is immediate in time or that it is infallible (cf. Moore 1988). Just like the insight in mathematical axioms may take time, the same is the case with insight in basic moral considerations. And just as we can err in our basic mathematical and perceptual beliefs, we can err in our basic moral beliefs (cf. Reid 1969; Roeser 2005). In addition, Reid emphasizes that we can also directly understand some moral truths that as a matter of fact do allow for further justification (Reid 1969).

Furthermore, most ethical intuitionists think that there is an irreducible plurality of morally relevant features.⁶ We cannot reduce justice, benevolence, happiness, honesty, gratefulness, promise keeping etc to each other. All these notions are self-evidently morally relevant; we cannot derive them from more fundamental ethical

⁶ With the exception of the utilitarian and intuitionist Henry Sidgwick (1901).

considerations or principles. This distinguishes intuitionism from Kantianism and utilitarianism, which are monistic theories. Kantianism and utilitarianism both think that there is one fundamental ethical principle and that all other ethical principles or considerations can be reduced to this. Kantians state that all ethical considerations have to be in accordance with the categorical imperative.⁷ Utilitarians state that the most basic moral principle is that we have to maximize happiness or utility. Monistic ethical theories cannot account for the fact that sometimes there can be conflicting moral demands on us which we cannot solve by trying to find one ‘master-principle’. In one situation, consideration A might be more important than consideration B; in another situation, this might be the other way around. This is W.D. Ross’s famous account of *prima facie*-duties: general moral principles only hold *prima facie*. In concrete cases one duty can overrule another one. There is no pre-established method to judge which duty is more important in which situation, even less is there a method providing us with a general serial ordering between moral principles (Ross 1967; also cf. Ewing 1929). Furthermore, there can be genuine dilemmas in which there are equally good reasons not to do A and not to do B, but in which there are no alternatives. Monistic ethical theories such as Kantian ethics or utilitarianism ignore the real existence of ethical dilemmas.

Utilitarians reduce all ethical considerations to maximizing happiness or utility. However, this gives rise to many counterexamples, for example if we intuitively feel a moral obligation to do something, for example because we promised it, even though it does not maximize happiness or utility. Cost-benefit analysis, which is the standard methodology in risk assessment, is modeled on utilitarianism, or more general, on consequentialism. According to consequentialism, the end justifies the means. The only morally important consideration in consequentialism is which option maximizes outcomes. This means that issues such as whether costs and benefits are fairly distributed and whether an action is performed voluntarily or not are ignored. Accordingly, consequentialist approaches such as utilitarianism and cost-benefit analysis require us to revise our ethical thinking. Instead, ethical intuitionists take side of our ethical intuitions. Intuitionists state that there is an irreducible plurality of morally relevant considerations that have to be balanced on a case by case basis.

Jonathan Dancy even goes further than that: he thinks that there are no necessary moral principles. According to him there are only context-specific moral truths, given by the concrete circumstances in a specific situation. Moral rules should be compared with inductive rules in the empirical sciences; they are merely generalizations which might have a heuristic value, but they are not a guarantee that something which is morally significant in a certain way in most situations will be so in all situations. In some situations it might not be important at all or its significance might even be reversed. For example honesty might generally be a good thing, but it can be very bad if it merely hurts somebody. Dancy’s opponents respond to this by suggesting to invoke more complex moral principles, but Dancy’s reply is that we can always think of new situations which undermine such a more complex rule. Furthermore, if we are able to make up candidates for complex general moral rules, why not use this ability to make particular moral judgments without ultimate reference to general rules? Apparently we are able to form complex moral judgments anyway (for all this, cf. Dancy 2004).

Let us look how we can apply these ideas to ethical judgments about risks. As indicated, ethical intuitionists think that we should start out with taking our ethical judgments at face value. In that sense, ethical intuitionism is a common sense approach. Most people know that it is (generally) wrong to steal, kill, lie etc. Moral philosophers might be able to reason more explicitly about ethics, but their intuitive moral judgments are not by definition better than those of people who have never read any scholarly literature on ethics (cf. Reid 1969). Furthermore, moral philosophers are far from being saints who never do things that they should not do. Ethical intuitionism is a theory that allows for the idea that the moral judgments of laypeople can be justified even if laypeople are not able to articulate a theoretical framework for their moral views. This can shed some light on the ethical risk judgments by laypeople. These judgments can be justified moral concerns, whether or not it is possible to come up with a further ethical justification for them. Maybe the only ‘justification’ that we can come up with is ‘this seems self-evident’. Ethical intuitionism is the only ethical theory which explicitly endorses this possibility. Other ethical theories try to avoid appeal to self-evident beliefs, but in practice, they cannot get around basic normative assumptions. Alternatively, they collapse into relativism (cf. Roeser forthcoming).

After having provided for a rough sketch of ethical intuitionism, it is time to take a closer look at the ethical risk judgments of laypeople and to examine whether they can be justified by further ethical arguments or whether we should consider them to be ‘self-evident’.

⁷ ‘Act only according to that maxim by which you can at the same time will that it would become a universal law’ (Kant 1964).

6. A PLURALIST ACCOUNT OF ETHICAL RISK CONSIDERATIONS

In this section I will examine the most important concerns of laypeople that Slovic has identified and I will discuss in how far they can be taken as reasonable moral concerns by using the framework of ethical intuitionism as sketched in the previous section. This means that ethical risk considerations can be understood as judgments concerning objective moral *prima facie* considerations or morally relevant features.⁸

Let us look at the risk characteristics which I have analyzed in section 3. The question I will ask in this section is whether these risk characteristics can be seen as justified moral concerns, either based on further going ethical argumentation or because they strike us as ‘self-evident’.⁹

The first cluster that Slovic et al. identified was ‘dread’. The question now is whether the elements of this cluster can reasonably be said to be of moral importance. I think that the answer is positive. Let us start with ‘dread’ (2): it is an open question in how far something that is frightening should be taken seriously. Does fear in itself provide for justification or only for so far as it points to something that deserves this response? I would say that it depends on whether one is justified in perceiving something as frightening. The other items in this cluster might provide justification for feelings of dread. Let us take a closer look at these items.

‘Severity not controllable’ (1): this is clearly morally relevant. Allowing an activity which might have negative effects the severity of which are beyond our control is like opening Pandora’s box. This is a risk which we reasonably only want to take if a lot can be gained which cannot be gained in a better way.

‘Globally catastrophic’ (3): The full text of the question refers to ‘catastrophic death and destruction across the whole world’. It is clearly self-evident that it is a legitimate moral concern to take into account in how far a risk might cause a global catastrophe.

‘Little preventive control’ (4): obviously it is relevant for the moral acceptability of a technology or activity whether we can prevent its negative effects or not, and obviously it is morally better if we can avoid negative side-effects.

‘Certain to be fatal’ (5): this is again a rather obvious concern. Surely it is a morally relevant feature of a risk whether it is certain to be fatal when it materializes.

It is not clear whether or how any of these considerations are or can be incorporated in cost-benefit analysis. The next considerations by definition are not part of cost-benefit analysis.

‘Risks and benefits inequitable’ (6): most philosophers consider fairness and equity as inherently morally relevant. The issue of fairness is not at all recognized in cost-benefit analysis. Many moral philosophers have criticized utilitarianism and consequentialism in general for this shortcoming, and cost-benefit analysis inherits this. Cost-benefit analysis could justify the exploitation of one group by another, no matter how small or large each of the groups, as long as the benefits of the one group outweigh the costs of the other group. This is not to say that such considerations may never be made, they are often unavoidable. But the point is that we cannot simply add and subtract costs and benefits without looking in more detail into issues about fair distribution and possible compensations. Certain costs might be unacceptable under (almost) all circumstances, such as intentionally killing or enslaving innocent people. To determine when unfair distributions might be unavoidable we need to go beyond the data that a cost-benefit analysis provides us with.

‘Catastrophic’ (7): In absolute numbers, two risks might be identical concerning how many victims they cause, but one hazard kills people one by one whereas another hazard kills many people at once. Does this mean that these hazards are morally on a par or might a chronic risk be worse than a catastrophic one or vice versa? This is a more complicated issue. From a consequentialist perspective, such as in cost-benefit analysis, we should be morally indifferent concerning which of these risks materializes. However, imagine two hazards. While the one

⁸ Ross (1967) talks about *prima facie*-duties, but we can also extend his account to other ethical notions such as values and virtues, all of which I mean to capture with the notion ‘consideration’. Note that Shrader-Frechette (1991) also refers to Ross’s *prima facie*-duties in this context, but she uses the phrase ‘subjective value judgments’, which completely goes against the spirit of Ross’s theory and that of all other intuitionists. Far from being subjectivists, all intuitionists, including Ross, are moral realists or objectivists, and all of them have argued extensively against subjectivism and other forms of relativism. This is not merely a verbal matter but a substantial philosophical issue, cf. my discussion of this issue in section 4.

⁹ For full questions, cf. appendix 2. A complicating factor might be that we do not know the full questions for all items and that in some cases where full questions are available, it is not entirely clear which question refers to which item, cf. note 3.

hazard can kill 1.000 people all at once, the other hazard kills 1.000 people distributed over 1.000 years. In that case, the second hazard seems to be more acceptable because it might be that in the future we can reduce the lethal power of this hazard. Hence, I think that it is not trivial to treat both hazards equally as consequentialism would demand. On a consequentialist account, this would be an irrelevant concern, whereas I think that it can be reasonable to take it into account.

This problem is related to the following considerations. ‘Not easily reduced’ (9): it is clear that it is morally relevant whether we can reduce the riskiness of a hazard or not. If the negative effects of a hazard can be easily reduced it can be a reason to accept this hazard and to try to reduce these effects.

‘Risks increasing’ (10): This is obviously important since it affects the acceptability of a risk through time, and it might also change our cost-benefit analysis. Cost-benefit analysis should take into account expected future developments. To my knowledge, this is not a standard requirement for cost-benefit analyses.

This is directly related to ‘Threatens future generations’ (8). Many philosophers and also other people think that we have a moral obligation to leave a world behind that is worthwhile living in for people who will live long after we have died. It is not clear in how far cost-benefit analysis takes into account effects that only manifest in a remote future. If it does not, this is morally reprehensible, if it does, this gives rise to further methodological problems concerning the predictability of future events. These methodological problems might induce a precautionary approach rather than trying to make a cost-benefit analysis based on shaky predictions.

‘Involuntary’ (11): voluntariness is a central moral concept which is directly related to the principle of autonomy which is endorsed by many ethical theories (cf. e.g. Mill ...). However, consequentialist approaches such as cost-benefit analysis can allow inflictions of autonomy since they favor actions which maximize outcomes, no matter how these outcomes are achieved. Hence, this is a concern that has a lot to go for it and yet is at odds with the standard approach in risk assessment.

‘Affects me personally’ (12): This is a problem for consequentialism and hence also for cost-benefit analysis, since those approaches do not attach different weights to different people. This can be morally right since it reflects the principle of equality. However, critics of consequentialism point out that personal considerations are not per se morally wrong. In the context of risks this is especially important concerning individual risks. For example, if a person employing the hazardous activity is the only one who might undergo the risks, and he or she is informed about the risks, he or she might be justified in undergoing the activity even if a cost-benefit analysis would tell that it is not worth it. This becomes more problematic concerning collective risks. But even in this context it is interesting that some moral philosophers argue that consequentialism unduly denies any form of self-interest and personal attachments. All these are debatable issues, but it is not unreasonable to at least take these considerations into account, which cost-benefit analysis does not allow.

Most or all items in the first cluster do not play a role in cost-benefit analysis. If I am right in my assessment that these are all in fact legitimate concerns, then it is shown that cost-benefit analysis misses out on numerous important issues which do figure in laypeople’s judgments.

Let us now look at the second cluster, ‘familiarity’. Most of the characteristics in this cluster are related to a lack of knowledge. In general we can state that a lack of knowledge about possible negative effects of a hazard is directly morally relevant. It might force us to be cautious. How would a lack of knowledge affect a cost-benefit analysis? Cost-benefit analysis incorporates probabilistic knowledge, but if the basis for statistical analysis is insufficient due to a substantial lack of knowledge, a cost-benefit analysis becomes highly problematic, unreliable or even impossible because of missing information. Hence, a lack of knowledge is considered to be of *methodological* importance for a cost-benefit analysis. However, the fact that there is a lack of knowledge is in itself not a concern of explicit *moral* importance in cost-benefit analysis. This is a shortcoming. This becomes more clear if we look at the individual features that this cluster consists of.

‘Not observable’ (13): this means that people might have no idea that an activity they employ might give rise to negative side effects. Our conventional alarm systems might not work in such cases. It seems reasonable that we should be extra careful with such risks.

‘Unknown to those exposed’ (14): this is highly relevant for the same reason as the previous item. In addition, both these items might undermine autonomous decision making, since autonomous choice requires adequate information (‘informed consent’; this is an important moral principle in medical ethics). These items state that information is lacking, which is directly relevant.

‘Effects immediate’ (15): it is clear that if a hazard can kill a person ten years after the exposure instead of immediately, this is still morally relevant. This is a possible problem for consequentialism and cost-benefit analysis, namely how to take into account risks that only manifest in the future, and until how far into the future we should take them into account (cf. my comments concerning characteristics 7-10). It should be incorporated into a cost-benefit analysis and it is probably possible to some degree, but the question is whether it always *is* incorporated.

‘New (unfamiliar)’ (16): newness does not need to be in itself a reason not to do something, however, it might be a reason to be careful and use a precautionary approach because we might not yet know what the possible negative consequences of this hazard are and how likely they are to occur.

‘Unknown to science’ (17): Here we can make the same point as concerning the previous consideration. A lack of scientific knowledge about risks can be a reason to apply the precautionary principle in some form or another. If we lack scientific evidence concerning possible negative consequences of a hazard, this is an important concern and can be a reason to be careful and demand further research before people can become exposed to that hazard. Of course there is a huge debate about the precautionary principle, whether it is applicable at all and if so, in which form, but this mainly concerns the implementation of the principle as a concrete methodology. That being cautious can be a morally wise thing should not be ruled out at forehand, but this is exactly what cost-benefit analysis does by not morally problematizing a lack of scientific evidence.

Hence, all the items in the cluster ‘familiarity’ point to important ethical considerations and principles such as ‘informed consent’ and the ‘precautionary principle’, considerations that are left out in cost-benefit analysis, which points to further shortcomings of that methodology.

Let us now look at the statistically separate feature, ‘Many people exposed’ (18): this obviously makes a difference, and it also plays a role in standard cost-benefit analysis. However, in cost-benefit analysis the exposure of many people to a risk can be justified if the benefits to a group of people are large enough, even if this only concerns a small group of people. It is doubtful whether this is morally justified (cf. characteristic 6.).

Hence, after having discussed all the risk considerations that Slovic et al. found to play a role in laypeople’s risk considerations, we can conclude that they are all reasonable or legitimate to a lesser or higher degree. We also saw that whereas some considerations can be substantiated with additional arguments, others simply strike us as self-evident. Most of these considerations are also central in various ethical theories. Furthermore, most of these considerations cannot be reduced to one or a few fundamental principles. For example, equity (6.) cannot be reduced to concerns about controllability. All these considerations form an irreducible plurality of morally relevant concerns about risks. Now the point might be made that all this is not very illuminating. Surely, the fact that something is catastrophic, lethal or unfair is morally relevant, but is this not a trivial or analytic claim? What does the appeal to intuitionism and self-evidence add? Well, this would be exactly the kind of response that an intuitionist would like to hear, since it would only confirm that we are obviously dealing with morally relevant features. So much the worse for traditional risk analysis to not take these considerations into account! We saw that all these considerations are insufficiently or even not at all addressed in the standard methodology for risk assessment, which is cost-benefit analysis. Cost-benefit analysis might often be an unavoidable basis for judgments about whether a risk is morally acceptable or not, but it is surely not always sufficient. The risk judgments of laypeople provide us with a substantial addition to the standard methodology if we want to achieve more adequate judgments about the moral acceptability of risks.

7. CONCLUSION

Let us return to the beginning of this paper where I mentioned that intuitions about risks are seen as very unreliable. We can now conclude that concerning intuitions about risks, these are indeed very unreliable in so far as they concern the quantitative aspect of risks, but they are necessary and can be legitimate concerning the ethical aspects of risks. Whereas the quantitative part of risk assessment should not be based on intuitions but should be done according to scientific methods, evaluations whether a risk is morally acceptable have to be made based on ethical intuitions. The philosophical theory of ethical intuitionism can provide us with a framework in which the considerations of laypeople concerning risks which Slovic et al. have found through surveys can be understood as legitimate moral concerns. Of course we could also reject cost-benefit analysis on philosophical grounds only, without taking the detour through the judgments of laypeople. However, the claim that Slovic has made is that laypeople have legitimate concerns about risks, and the discussion in section 6 was meant to substantiate this claim. This discussion showed that the concerns of laypeople are legitimate, and these concerns cannot be reduced to one simple methodology such as cost-benefit analysis.

However, this approach comes at a prize, which is that it does not give us any guidelines as to how to make overall evaluations about the moral acceptability of risks. A hope might be that we can establish a serial ordering of ethical concerns, rated by importance. However, it is doubtful whether this is possible. It might not do justice to each possible case and hence lead to oversimplification again. The intuitionist framework might do more justice to the complexity of the ethical aspects of risks than cost-benefit analysis, but this means that discussions about acceptable risks will have to be made on a case by case basis, involving thorough ethical reflection instead of applying a clear cut methodology. This is what the intuitionist Ross has to say about this:

Loyalty to the facts is worth more than a symmetrical architectonic or a hastily reached simplicity. If further reflection discovers a perfect logical basis for this or a better classification, so much the better (Ross 1967, 23).

From the previous discussion of legitimate concerns of laypeople it is clear that cost-benefit analysis does not provide for a perfect logical basis for moral judgments about risks. This means that we should find different ways of doing risk analysis. According to Ross, judgments about morally complex issues are merely ‘fallible opinions’ where many good judges might disagree (Ross 1968, 189). In such cases, the best approach might be to hear many different voices and let various people exchange considerations and arguments. This will be no guarantee to come up with infallible solutions, but it might be the best we can do in the face of morally complex issues. In any case, this means that in risk assessment we need more than technological and policy expertise, we also need ethical ‘expertise’, which can for example be provided by concerned citizens who are willing to contribute to ethics panels and other alternative public decision making bodies.¹⁰

REFERENCES

- Alston, W. P., 1993. *The Reliability of Sense Perception*, Cornell University Press, Ithaca and London.
- Dancy, J., 2004. *Ethics without Principles*, Oxford University Press, Oxford.
- Dancy, J., 1993. *Moral Reasons*, Blackwell, Oxford.
- Ewing, A.C., 1929. *The Morality of Punishment*, Kegan Paul, London.
- Fischhoff, B., Lichtenstein, S., Slovic, P., Derby, S.L., Keeney, R., 1981. *Acceptable Risk*, Cambridge University Press, Cambridge.
- Gilovich, T., Griffin, D., Kahnemann, D., (Eds.), 2002. *Intuitive Judgment: Heuristics and Biases*, Cambridge University Press, Cambridge.
- Jaeger, C.J., Renn, O., Rosa, E.A., Webler, T., 2001. *Risk, Uncertainty, and Rational Action*, Earthscan, London.
- Jasanoff, S., 1993. Bridging the Two Cultures of Risk Analysis. *Risk Analysis* 13, 123-129.
- Kant, I., 1964 [1786]. *Groundwork of the Metaphysic of Morals*, Harper and Row, New York.
- Mill, J.S., 1985 [1859]. *On Liberty*, Penguin, Harmondsworth.
- Moore, G.E., 1988 [1903]. *Principia Ethica*, Prometheus Books, Buffalo, NY.
- Reid, T., 1969 [1788]. *Essays on the Active Powers of the Human Mind*, Introduction by Baruch Brody, The M.I.T. Press, Cambridge, Massachusetts, and London, England.
- Roeser, S., *Intuitionism, Moral Truth, and Tolerance*. *The Journal of Value Inquiry* (forthcoming)
- Roeser, S., *Defending Moral Intuition*. In: Van Woudenberg, R., Roeser, S., Rood, R., (Eds.) 2005, *Basic Belief and Basic Knowledge: Papers in Epistemology*, Ontos Verlag, Frankfurt, pp. 231-250
- Ross, W.D., 1967 [1930]. *The Right and the Good*, Clarendon Press, Oxford.
- Ross, W.D., 1968 [1939]. *Foundations of Ethics*. The Gifford Lectures, Clarendon Press, Oxford.
- Shrader-Frechette, K., 1991. *Risk and Rationality*, University of California Press, Berkeley.
- Sidgwick, H., 1901 [1874]. *The Methods of Ethics*, Macmillan, London, New York.
- Slovic, P., 2000. *The Perception of Risk*, Earthscan, London.
- Sunstein, C.R., 2005. *Laws of Fear*, Cambridge University Press, Cambridge.
- Tversky, A., & Kahneman, D., 1974. Judgment under Uncertainty: Heuristics and Biases. *Science* 185, 1124-1131

¹⁰ Work on this paper was supported by the Netherlands Organisation for Scientific Research (NWO) under grant number 275-20-007.

APPENDIX 1

Appendix 1 to Roeser, S., 'Ethical Intuitions about Risks': Table 'Incorrelations among 18 Risk Characteristics in the Extended Stur

	Severity not controllable	Dread	Globally catastrophic	Little preventive control	Certain to be fatal	Risks & benefits inequitable	Catastrophic	Threatens future generations	Not easily reduced	Risks increasing	Involuntary	Affects me personally	Not observable	Unknown to those exposed	Effects immediate	New (unfamiliar)	Unknown to science	Many people exposed	
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	
1. Severity not controllable																			
2. Dread	0.82																		
3. Globally catastrophic	0.78	0.83																	
4. Little preventive control	0.86	0.72	0.71																
5. Certain to be fatal	0.8	0.82	0.73	0.77															
6. Risks & benefits inequitable	0.75	0.76	0.84	0.63	0.65														
7. Catastrophic	0.77	0.66	0.77	0.74	0.67	0.77													
8. Threatens future generations	0.62	0.76	0.86	0.52	0.64	0.81	0.63												
9. Not easily reduced	0.59	0.67	0.64	0.67	0.66	0.59	0.56	0.59											
10. Risks increasing	0.51	0.63	0.76	0.48	0.58	0.76	0.57	0.75	0.6										
11. Involuntary	0.68	0.58	0.69	0.52	0.42	0.77	0.74	0.62	0.33	0.44									
12. Affects me personally	0.57	0.64	0.77	0.43	0.5	0.71	0.64	0.78	0.35	0.67	0.61								
13. Not observable	-0.04	-0.14	0.04	-0.19	-0.28	0.08	0	0.24	-0.15	-0.1	0.33	0.1							
14. Unknown to those exposed	0.14	0.05	0.22	-0.05	-0.12	0.28	0.24	0.35	-0.2	0.05	0.63	0.31	0.79						
15. Effects immediate	0.21	0.15	0	0.34	0.36	-0.08	0.11	-0.29	0.22	0.01	-0.25	-0.14	-0.87	-0.77					
16. New (unfamiliar)	0.32	0.29	0.32	0.22	0.2	0.25	0.18	0.44	0.13	0.06	0.36	0.2	0.63	0.56	-0.48				
17. Unknown to science	-0.1	-0.1	-0.04	-0.07	-0.15	-0.02	-0.04	0.02	-0.23	0.02	0.13	0	0.43	0.5	-0.44	0.32			
18. Many people exposed	-0.04	0.04	0.23	-0.14	-0.11	0.32	0.23	0.47	0.03	0.26	0.34	0.56	0.37	0.46	-0.52	0.07	-0.01		

Permission republication appendix 1 - Ethical Intuitions about Risks (S. Roeser):

Paul Slovic, Baruch Fischhoff and Sarah Lichtenstein, Facts and Fears: Understanding Perceived Risk, in R C Schwing and W A Albers, Jr (eds), Societal Risk Assessment: How Safe is Safe Enough? New York, Plenum, 1980, pp 181-214. This book is currently published by Springer. The paper has been republished in Paul Slovic (2000), The Perception of Risk, Earthscan, chapter 8, table 8.2. With kind permission of Springer Science and Business Media, Earthscan and Paul Slovic.

APPENDIX 2

Full questions related to the various characteristics (but cf. note 3):

1. 'If a mishap occurs, can the damage be controlled?', (Slovic 2000, 138; question 10).
2. 'Is this a risk that people have learned to live with and can think about reasonably calmly, or is it one that people have great dread for – on the level of a gut reaction...?' (Slovic 2000, 87, question 8).
3. Formulation 1: 'Does the hazard threaten global catastrophe?' (Slovic 2000, 138, question 15), formulation 2: 'To what extent does pursuit of this activity, substance or technology have the potential to cause catastrophic death and destruction across the whole world?' (Slovic 2000, 139).
4. 'Can mishaps be prevented?' (Slovic 2000, 138, question 9) or 'If you are exposed to the risk of each activity or technology, to what extent can you, by personal skill or diligence, avoid death while engaging in the activity?' (Slovic 2000, 86, question 5). This is a characteristic where it is unclear which of the questions corresponds to it, cf. note 3. While the second formulation refers to the ability of the subject answering the survey if he or she can prevent a risk to die by being him or herself engaged in the activity, the first formulation does not specify who will be at risk and who will be able to prevent the risk from manifesting, and it does not specify what kind of risk might be at stake ('mishap'). Since the first question is more general I have this in mind in my discussion. But it might obviously make a significant difference whether a subject is to rate a hazard according to the first or the second formulation.
5. 'Severity of consequences: When the risk from the activity is realized in the form of a mishap or illness, how likely is it that the consequence will be fatal?' (Slovic 2000, 87, question 9).
6. 'Are the benefits equitably distributed among those at risk?' (Slovic 2000, 138, question 14).

7. 'Is this a risk that kills people one at a time (chronic risk) or a risk that kills large numbers of people at once (catastrophic risk)?' (Slovic 2000, 87, question 7).
8. 'Does the hazard threaten future generations?' (Slovic 2000, 138, question 12).
9. 'Can the risks be reduced easily?' (Slovic 138, question 18).
10. 'Are the risks increasing or decreasing?' (Slovic 2000, 138, question 17).
11. 'Do people get into these situations voluntarily?' (Slovic 86, question 1).
12. 'Are you personally at risk from this hazard?' (Slovic 2000, 138, question 138).
13. 'Are the damage-producing processes observable as they occur?' (Slovic 2000, 138, question 16).
14. 'To what extent are the risks known precisely by the persons who are exposed to those risks?' (Slovic 2000, 86, question 3).
15. 'To what extent is the risk of death immediate - or is death likely to occur at some later time?' (Slovic 2000, 86, question 86).
16. 'Are these risks new, novel ones or old, familiar ones?' (Slovic 2000, 87, question 6).
17. 'To what extent are the risks known to science?' (Slovic 2000, 86, question 4).
18. 'How many people are exposed to this hazard?' (Slovic 2000, 138, question 11).