

Seven Myths of Risk

Talk at the conference *Stockholm thirty years on. Progress achieved and challenges ahead in international environmental co-operation*. Swedish Ministry of the Environment, June 17-18, 2000

Sven Ove Hansson, Royal Institute of Technology, Stockholm

soh@infra.kth.se

The purpose of this presentation is to introduce both the concept of risk and the precautionary principle, that is a major policy principle in present-day risk management. Since risk has been the subject of many misconceptions I will do this in large part by criticizing seven views on risk that I believe to have caused considerable confusion both among scientists and policy-makers. But before looking at the seven myths of risk, let us begin with the basic issue of defining “risk”.

The word “risk” often refers, rather vaguely, to situations in which it is possible but not certain that some undesirable event will occur. In addition, the word has several more specialized meanings. Let me illustrate this by making a few statements about the single most important preventable health hazard in non-starving countries. First: “Lung cancer is one of the major risks that affect smokers.” Here, we use “risk” in the following sense:

(1) risk = an *unwanted event* which may or may not occur.

Next, let me repeat what I just said about the severity of this hazard: “Smoking is by far the most important health risk in industrialized countries.” Here we use “risk” in another sense:

(2) risk = the *cause* of an unwanted event which may or may not occur.

Next, we can quantify this risk in the following way: “There is evidence that the risk of having one’s life shortened by smoking is about 50%.” This is yet another concept of risk:

(3) risk = the *probability* of an unwanted event which may or may not occur.

In risk analysis, smoking is often compared to other risk factors in terms of the statistically expected number of victims. We may say, for instance: “The total risk from smoking is higher than that from any other cause that has been analyzed by risk analysts.” I will return shortly to what this means more precisely; for the moment let us just record the usage:

(4) risk = the statistical *expectation value* of unwanted events which may or may not occur.

In decision theory, an essential distinction is that between decisions “under risk” and “under uncertainty”. The difference is that in the former case, but not the latter, probabilities are assumed to be known. Hence we may say: “The probabilities of

various smoking-related diseases are so well-known that a decision whether or not to smoke can be classified as a decision under risk”.

(5) risk = the fact that a decision is made under conditions of *known probabilities*
 (“decision under risk”)

With this we have identified five common meanings or “risk”. There are also several other more technical meanings, in particular in economic theory. In spite of this, however, one often hears claims that “risk” should be defined once and for all.

The first myth of risk:

“Risk” must have a single, well-defined meaning.

In practice, the project of defining one single meaning of “risk” and excluding all others is a futile form of linguistic imperialism, as impossible as the corresponding task for most other words with several entrenched uses in both everyday and technical language. In spite of this, attempts at such linguistic imperialism are not uncommon. In most cases, they take place in favour of the fourth meaning of “risk” on my list, namely risk as statistical expectation value. This is the standard meaning of “risk” in professional risk analysis. In that discipline, “risk” often denotes a numerical representation of severity, that is obtained by multiplying the probability of an unwanted event with a measure of its disvalue (negative value). When, for instance, the risks associated with nuclear energy are compared in numerical terms to those of fossil fuels, “risk” is usually taken in this sense. All the major variants of technological risk analysis are based on the identification of risk with expectation value. In addition, this concept of

risk is used in several related fields. In risk-benefit analysis, the risks that are weighed against benefits are expectation values. In studies of “risk perception”, the “subjective risk” reported by the subjects is compared to the “objective risk”, which is identified with the value obtained in this way.

It could be argued that it does not matter much how we use the word “risk” – or any other term – if only the usage is well-defined. However, another more serious and much more substantial matter lies behind this, namely the issue of how to judge and assess the severity or the acceptability of risks.

The second myth of risk:

The severity of risks should be judged according to probability-weighted averages of the severity of their outcomes.

Although this approach has the advantage of being simple, operative, and mathematizable, it often goes severely wrong when applied to real life problems. The reason for this is that it does not take all the relevant factors into account. In real life, there are always other factors in addition to probabilities and utilities that can – and should – influence appraisals of risk. Risks are inextricably connected with interpersonal relationships. They do not just “exist”; they are taken, run, or imposed. To take just one example, it makes a big difference if it is my own life or that of somebody else that I risk in order to earn a fortune for myself. Therefore, person-related aspects such as agency, intentionality, consent, equity etc. will have to be taken seriously in any reasonably accurate general format for the assessment of risk. An analysis of risk that includes considerations of agency and responsibility will be an analysis more in terms of the verb (to) ‘risk’ than of the noun (a) ‘risk’. Major policy debates on risks

have in part been clashes between the “noun” and the “verb” approach to risk. Risk analysts and other experts tend to emphasize the *size* of risks, hence treating risks as isolated objects. Members of the public often question the very act of *risking* improbable but potentially calamitous accidents. Risk analysts’ one-sided focus on probabilities and outcomes, to the exclusion of other important factors that could legitimately influence decisions, is a major reason why risk analysis has had such great difficulties in communicating with the public. Instead of blaming the public for not understanding probabilistic reasoning, risk analysts should learn to deal with the moral and social issues that the public – rightly – put on the agenda.

The third myth of risk permeates most of modern risk analysis:

The third myth of risk:

Decisions on risk should be made by weighing total risks against total benefits.

This “weighing principle” is seductive; at first glance it seems reasonable, perhaps even obviously true. But it is the source of much trouble, since it requires a farreaching disregard of persons.

It follows from this weighing principle that a risk imposition is acceptable if the total benefits that it gives rise to outweigh the total risks, irrespective of who is affected by the risks and the benefits. In particular, since risks and benefits are treated separately, this model prevents us from attaching importance to whether the risks and the benefits associated with a particular technology accrue to the same persons or to different persons. In this sense, mainstream risk analysis does not take persons seriously.

This is a feature that risk analysis has in common with utilitarian moral theory. In that theory, utilities and disutilities that pertain to different individuals are added, with no respect being paid to the fact that they are bound to different persons. Persons have no role in the ethical calculus other than as bearers of risks and benefits whose value is independent of whom they are carried by. Therefore, a disadvantage affecting one person can always be justified by a sufficiently large advantage to some other person. Similarly, in mainstream risk analysis, benefits for one person may easily outweigh risk-exposure affecting other persons. Consider a polluting industry somewhere in Sweden. The total economic advantages to the Swedish population of this industry may outweigh the total health risks that the pollution gives rise to. However, for those who live in the neighbourhood the situation is radically different. The whole health risk burden that the pollution from the plant gives rise to falls on them. Nevertheless, they receive a much smaller share of the economic advantages. In risk-benefit analysis, performed in the standard way as expected utility maximization, such distributional issues are disregarded – or at best treated as an addition to the main analysis. To the common moral intuition, this is an implausible way of thinking.

In its impersonal mode of analysis, risk analysis represents an extreme central planning paradigm, in sharp contrast to mainstream normative economics. The dominant approach in modern economics does not condone reallocations without mutual consent. The reason why Pareto optimality has a central role in modern normative economics is precisely that it respects the person, i.e. it does not allow the comparison of advantages and disadvantages that accrue to different persons. Indeed, in standard normative economics, an advantage to one person cannot outweigh a disadvantage to another person. Standard risk analysis represents the contrary approach.

Hence, we have presently a strange combination of dominant ideologies: a form of risk analysis that requires interpersonal aggregation and a form of economic analysis that prohibits it. This combination is particularly unfortunate from the perspective of the poor, risk-imposed person. The economic criterion of Pareto-optimality precludes the transfer of economic goods to her from rich persons if they object to such a transfer, whereas risk analysis allows others to expose her to risks, against her own will, if only someone else – perhaps the same rich persons – have sufficient benefits from this exposure.

My proposal is that we give up the impersonality of traditional risk analysis, and develop tools for *ethical risk analysis* that respect the person. This requires that we accept every person's 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, but barring exceptional cases this should only happen when it is in the individual's own interest. In other words, we should treat each risk-exposed person as a sovereign individual who has a right to fair treatment, rather than as a carrier of utilities and disutilities that would have the same worth if they were carried by someone else. With this approach follows a new agenda. According to traditional risk analysis, in order to show that it is acceptable to impose a risk on Ms Smith, one has to give sufficient reasons for accepting the risk as such, as an impersonal entity. According to the ethical risk analysis that I propose, one instead has to give sufficient reasons for accepting that Ms Smith is exposed to the risk.

One sometimes encounters the view that issues of risk are "experts' issues", and that non-experts such as our elected representatives, should not have much of a say.

The fourth myth of risk:

Decisions on risk should be taken by experts rather than by laymen.

In practice, it is not possible to dissect out decisions on risk from the political decision-making process as a whole, and it is not either possible to dissect out the criteria of risk assessment from social values in general. Almost all social decisions can be described in terms of avoiding undesirable events. Some of these events can be quantified in a meaningful way, for example health risks and economic risks. Others are virtually impossible to quantify, like risks of cultural impoverishment, social isolation, and increased tensions between social strata. The various models advocated for risk evaluation involve selecting certain of the more quantifiable factors, for example expected number of deaths and expected economic losses. A comparison is then made between the effects on these factors of different alternatives, and the effects are weighed against each other. Some social values necessarily fall outside of such comparisons. The more general social and political decision-making process is suited to make decisions on a more comprehensive basis, including the non-quantifiable factors that are excluded from risk analysis. Furthermore, successful social and political decision-making processes in controversial issues typically involve bargaining and compromises, that are not part of the expert-driven risk analysis process.

It is often claimed that policy-makers are irrational when they do not follow the risk assessments made by their experts. I have already mentioned one reason why this is not in general a fair accusation, namely that decision-makers have to take into account many aspects not covered by the experts in their risk assessments. An additional reason is that experts are known to have made mistakes. A rational decision-maker should take into account the possibility that this may happen again. Experts often do not realize that for the non-expert, the possibility of the experts being wrong may very well be a

dominant part of the risk involved e.g. in the use of a complex technology. When there is a wide divergence between the views of experts and those of the public, this is certainly a sign of failure in the social system for division of intellectual labour, but it does not necessarily follow that this failure is located entirely within the minds of the non-experts who distrust the experts. It cannot be a criterion of rationality that one takes experts for infallible.

In issues of risk, like all other social issues, the role of the expert is to investigate facts and options, not to make decisions or to misrepresent facts in a unidimensional way that gives the decision-maker no choice.

The fifth myth is an extreme form of the fourth one. I call it the “technocratic dream”.

The fifth myth of risk:

Risk-reducing measures in all different sectors of society should be decided according to the same standards.

According to this idea, risk analysts should perform analysis with uniform methods for all parts of society: be it mammography, workplace safety, railway safety, or chemicals in the environment. The idea is to calculate the risks in all these different places, and then allocate resources for abatement in a way that minimizes the total risks.

What the technocratic dreamers do not realize is that risk issues are parts of larger, more complex issues. Traffic safety is closely connected to issues of traffic and community planning. Mammography has important social and psychological aspects, and must be seen in the context of the general organization of preventive medicine. Workplace safety issues are related to issues of workplace organisation, labour law etc.

In short, the risk issues of different social sectors always have important aspects that connect them to other issues in these respective sectors. The technocratic dream, with its unified calculation for all social sectors, is insensitive to the concerns and the decision procedures of the various social sectors. It is in fact not compatible with democratic decision-making as we know it. The technocrats' dream is a nightmare for the rest of us.

I will now turn to the precautionary principle, and with that to the sixth myth

The sixth myth of risk:

Risk assessments should be based only on well-established scientific facts.

How could anyone object to that? Why not base our decisions only on good, "sound" science? To explain this, let me first say a few words about scientific knowledge in general.

Scientific knowledge begins with data that originate in experiments and other observations. Through a process of critical assessment, these data give rise to the scientific corpus. Roughly speaking, the corpus consists of those statements that could, at the time being, legitimately be made, without reservation, in a (sufficiently detailed) textbook. When determining whether or not a scientific hypothesis should be accepted for the time being as part of the corpus, the onus of proof falls squarely to its adherents. Similarly, those who claim the existence of an as yet unproven phenomenon have the burden of proof. These proof standards are essential for the integrity of science.

The obvious way to use scientific information for policy purpose is to use information from the corpus. For many purposes, this is the only sensible thing to do. However, in the context of risk it may have unwanted consequences to rely exclusively

on the corpus. Suppose that there are suspicions, based on relevant but insufficient scientific evidence, that a certain chemical substance is dangerous to human health. Since the evidence is not sufficient to warrant an addition to the scientific corpus, this information cannot influence policies in the “standard” way. However, the evidence may nevertheless be sufficient to warrant changes in technologies in which that chemical is being used. We want, in cases like this, to have a direct way from data to policies.

In order to avoid unwarranted action due to misinterpreted scientific data, it is essential that this direct road from data to policy be guided by scientific judgement in essentially the same way as the road from data to corpus. The major difference is that in the former case, the level of required proof is adjusted to policy purposes.

We have therefore, two different decision processes. One consists in determining which scientific hypothesis should be included into the scientific corpus. The other consists in determining which data and hypotheses should influence practical measures to protect health and the environment. It would be a strange coincidence if these two criteria always coincided. Strong reasons can be given for strict standards of proof for scientific purposes. At the same time, according to the dictum “better safe than sorry”, there are strong reasons to allow risk management decisions to be influenced by sound scientific indications of danger that are yet not sufficiently well-established to qualify for inclusion into the scientific corpus. This difference in criteria between the two decisions is the essential insight on which the precautionary principle is based.

As should be obvious from this, the precautionary principle does not require less science, but rather more science than non-precautionary decision-making. We need to learn how to use science to distinguish in a consistent way between serious and less serious indications of danger. In the NewS program – New Strategy for the Risk

Management of Chemicals – a group of Swedish scientists are developing scientific methodology for this purpose.

According to some critics, the precautionary principle fails to pay enough respect to science, since it requires that precautionary measures be taken also against threats for which full scientific evidence has not been established. For instance, the use of a chemical substance can be prohibited by the precautionary principle even if we do not *know* whether it threatens the environment or not. In spite of its convincing first appearances, this argument breaks down as soon as sufficient attention is paid to its key term ‘unscientific’. There are two meanings of this word. A statement is unscientific in what we may call the weak sense if it is *not based on* science. It is unscientific in what we may call the strong sense if it *contradicts* science. Creationism is unscientific in the strong sense. Your aesthetic judgments are unscientific in the weak but presumably not in the strong sense.

The precautionary principle is certainly unscientific in the weak sense, but then so are all decision rules – including the rule that equates the evidence required for practical measures against a possible hazard with the evidence required for scientific proof that the hazard exists.

On the other hand, the precautionary principle is *not* unscientific in the strong sense. A decision-maker who applies the precautionary principle will use the same type of scientific evidence, and assign the same relative weights to different kinds of evidence, as a decision-maker who requires full scientific evidence before actions are taken. The difference lies, as I have already said, in the amount of scientific evidence that is required for a decision.

Finally, I will turn to the most dangerous of the seven myths

The seventh myth or risk:

If there is a serious risk, then scientists will find it if they look for it.

It is very often implicitly assumed that what cannot be detected cannot be a matter of concern. Occasionally, this has also been explicitly stated. Hence, the Health Physics Society wrote in a position statement:

“...[E]stimate of risk should be limited to individuals receiving a dose of 5 rem in one year or a lifetime dose of 10 rem in addition to natural background.

Below these doses, risk estimates should not be used; expressions of risk should only be qualitative emphasizing the inability to detect any increased health detriment (i.e., zero health effects is the most likely outcome).”¹

The reason why this is an untenable standpoint is that many risks are in fact undetectable. The following hypothetical example can be used to explain why.

There are three chemical substances A, B, and C, and 1000 persons exposed to each of them. Exposure to A gives rise to hepatic angiosarcoma among 0.5 % of the exposed. Among unexposed individuals, the frequency of this disease is very close to 0. Therefore, the individual victims can be identified. This effect is detectable on the individual level.

Exposure to B causes a rise in the incidence of leukemia from 1.0 to 1.5 %. Hence, the number of victims will be the same as for A, but although we know that about 10 of the about 15 leukemia patients would also have contracted the disease in the absence of exposure to the substance, we cannot find out who these ten patients are.

¹ Health Physics Society (1996), *Radiation Risk in Perspective*. Position statement of the Health Physics Society, adopted January 1996. Downloaded in December 1998 from <http://www2.org/hps/rad.htm>.

The victims cannot be identified. On the other hand, the increased incidence is clearly distinguishable from random variations (given the usual criteria for statistical significance). Therefore, the effect of substance B is detectable on the collective (statistical) but not on the individual level.

Exposure to C leads to a rise in the incidence of lung cancer from 10.0 to 10.5 %. Again, the number of additional cancer cases is the same as for the other two substances. Just as in the previous case, individual victims cannot be identified. In addition, since the difference between 10.0 and 10.5 % is indistinguishable from random variations, the effects of this substance are undetectable even on the collective level.

We can therefore distinguish between effects that are completely undetectable, like the effects of substance C, and effects that are only individually undetectable, like those of substance B. To see why this is not a “merely academic” issue, but a central concern in practical risk management, let us focus on lifetime risks of lethal effects. As a rough rule of thumb, epidemiological studies can reliably detect excess relative risks only if they are about 10 % or greater. For the more common types of lethal diseases, such as coronary disease and lung cancer, lifetime risks are of the order of magnitude of about 10 %. Therefore, even in the most sensitive studies, an increase in lifetime risk of the size 10^{-2} (10 % of 10 %) or smaller may be undetectable (i.e. indistinguishable from random variations). In animal experiments we have similar experimental problems, and in addition problems of extrapolation from one species to another.

How small health effects should be of concern to us? Many attempts have been made to set a limit of concern, expressed either as “acceptable risk” or “de minimis risk”. Most of us would agree that if a human population is exposed to a risk factor that will, statistically, kill one person out of 10^9 , then that risk will not be an issue of high

priority. Arguably, it would be no disaster if our risk assessment methods are insufficient to discover risks of that order of magnitude. On the other hand, most of us would consider it a serious problem if a risk factor kills one person out of 100 or 1000. The most common proposals for limits of concern for lethal risks are 1 in 100 000 and 1 in 1 000 000. It is difficult to find proposals above 1 in 10 000. These values are of course not objective or scientific limits; I just report what seems to be levels at which lethal risks are often *de facto* accepted (as distinguished from *acceptable*).

We therefore have a wide gap between those (probabilistic) risk levels that are scientifically detectable and those that are commonly regarded to be of minor concern. This little-known gap has a breadth of 2–4 orders of magnitude. It is a major reason why the inference from “no known risk” to “no risk” is a dangerous one.

There are two ways to bridge this gap. Scientists can develop mechanistic knowledge, for instance of toxicity, that allows us to infer effects through other means than direct detection. Risk managers can develop cautious strategies, such as uncertainty factors (safety factors) through which the gap can be bridged. It is my personal conviction that what we need to solve the serious risk problems that are ahead of us is the combination of these two approaches. Science without precaution means acting too little and too late against environmental hazards. Precaution without science means acting with the wrong priorities. What we need is science-based precaution.