

Philosophical Perspectives on Risk and Uncertainty

Niels Gottschalk-Mazouz (University of Bayreuth)

Gottschalk-Mazouz, N. (2019): „Philosophical Perspectives on Safety and Risk” in: Raue, M.; Streicher, B.; Lerner, E. (Hg.) *Perceived Safety. A Multidisciplinary Perspective*, Cham: Springer, 27-41.

Introduction

Humans care about themselves and their future. If nothing bad can happen, they are safe. If they know something for certain, they have certainty. But, as finite beings, humans do not enjoy much certainty. And in a dynamic world, they are rarely safe. “Risk” is one of the concepts that may help to come to terms with this. It is, however, notoriously unclear. That is why this text is, in its first parts, dealing with terminology. The main objective of this text, however, is a discussion of rational and moral aspects of risk-taking, and of risk governance within what will be called “risk cultures”.

Risk: Four basic uses of the term

The different uses of this concept can be systemised as follows (Gottschalk-Mazouz 2011):

Risk means either (1) the *possibility* of an event, or (2) the *probability* of an event, or (3) the *value* of a (possible or probable) event, or (4) the *product* of probability and value of an event. All these attributions can be meant *objectively* or *subjectively*. E.g. with respect to the possibility of some bad thing that may happen, i.e. risk (1), objective risk means that bad things *can* happen, whereas subjective risk means that one *thinks* that bad things can happen.

Of course, not only bad things can happen. But, when speaking of risk, very often we seem to be focussing exclusively on bad things. So, it is presupposed, then, that either nothing happens (if we are lucky), or that something bad happens (if we are unlucky). In other words, the event under scrutiny is a bad thing, a loss, a damage. With respect to risk (3) and (4), we take the value of the event to be the extent of this damage then. If the extent of damage is measured in some monetary value, the symbolic expression of risk (4) as $R = P \times E$ amounts to what is sometimes called the “insurance formula”: Risk equals probability of damage times extent of damage. When we focus exclusively on bad things, we are framing risk in a negative way, because the value of the event can only be negative. The “greater” the risk the worse it is. That is why I call this the *negative* concept of risk.

But good things can happen as well. E.g. it is often said that we should not focus only on the risks but also on the chances, or opportunities. This talk of chances or opportunities can be analysed in completely symmetric way: With this it is presupposed that either nothing happens (if we are unlucky) or that something good happens (if we are lucky). This amounts to something that, if people were using the term ‘risk’ for it, should be called the *positive* concept of risk. But people rarely do.

Rather, what people do, is incorporate the chances into a broader concept of risk that is neither purely negative nor purely positive, and that I want to call *speculative* risk. With it, we allow for bad as well as good things to happen *within* the concept of risk. So what we presuppose then is that either a bad thing happens (if we are unlucky), or a good thing happens (if we are lucky), or nothing at all happens (which is taken to be neutral, i.e. is setting the baseline). We do so by allowing for negative or positive values of the events.

Now, the world is rarely such that events come alone. They come as clusters, or as consequences, i.e. as sets, which usually are of mixed value: some negative, some positive.

Now, we have two equivalent ways of expressing that. Following the concept of negative risk, we can say that a given set of events contains risks and opportunities. If we know enough about it to compare them, we can say that it contains more risk than opportunity (or the other way around, or that they are equal). Following the concept of speculative risk, we would rather say that the set of events contains negative and positive risk. And if we know enough about it, that the *overall speculative risk* or the *risk balance* is negative (or positive, or neither). What we are saying in any of these two ways can be given the same precise meaning depending on what concept of risk (1-4) is invoked. That there is more risk than opportunity, or that the risk balance is negative, would mean that (1) more of the events are negative than positive, or that (2) the negative events are more probable than the positive ones, or that (3) the value of the negative events is – in absolute terms – greater than the value of the positive events, or that (4) the product of probability and value of the negative events is – again in absolute terms – greater than the respective product for the positive valued events.

To be able to form a risk balance, you have to be able to compare risks (1-4) according to some apt metrics. The definitions of risk (1-4) point at respective natural metrics, however, which for risk (1) is a binary number, for risk (2) is a real number between 0 and 1 or any other probability measure, for risk (3) is a monetary value or any other cardinal (or, for that matter, at least ordinal) value, and for risk (4) is the product of probability and value which is typically also a value, namely a cardinal value.

However, events do not form clusters by nature, at least not in a sense that can be used in balancing risks: Causal chains (or causal networks) stretch out indefinitely in time and, in a connected world, get very complicated to trace even for relatively short timespans. So every question about risks (1-4) have to be constrained according to some background relevance

conditions and standards. If you ask for a possibility, as in risk (1), it has to be made clear what is taken to be constant (e.g. natural laws, social structures, individual behaviour etc.) and what is taken to be variable (e.g. environmental conditions). Moreover, of course, you need standards of demarcation that allow you to give a yes/no answer. Only then it makes sense to ask whether an event x is possible or not (e.g. that your house will suffer from flooding). If you ask for a probability, as in risk (2), further constraints have to be introduced: The standard way to understand probabilities is to see them as relative frequencies, i.e. ratios of numbers of events (e.g. the weeks where your house will suffer from flooding divided by the total number of weeks). Now we need a standard of division (e.g. the weeks). If the risk definition involves evaluations, as in risk (3) and (4), it is clear that an evaluative standard (as one of many possible standards) is presupposed as being adequate for evaluating these events (e.g. a monetary value of the damage due to flooding). All these standards might be controversial, which may be evident for evaluations (e.g. shall we monetarize when loss of lives is involved?), but is also the case for divisions (e.g. shall traffic accidents be counted per year or not better per 10.000 km?) or demarcations (e.g. from where on shall we speak of “flooding”, or of a “traffic accident”?). This is virulent in all the risk comparisons that we frequently make, e.g. travelling by plane might be riskier than by car if compared by the hour, but is less risky if compared by distance travelled.

Risk: Objective, subjective, and perceived

Subjective risk is not necessarily the same as perceived risk. As the terms have been introduced, objective and subjective denote kinds of properties of events. “Objective” means that something is attributed with respect to the object, i.e. that it lies in the nature of the object whether it is correctly attributed, and this is what makes it true or false.

“Subjective” means that now it lies in the nature of the subject, i.e. its beliefs, expectations or desires, whether it is correctly attributed, and this is what makes it true or false.

In the case of risk (1), objective risk would mean that something can happen. It is an objective question, independent of what we believe or want, whether this is the case or not. Subjective risk would mean that we expect that something can happen. In the case of risk (2), objective risk would mean either that we apply some concept of objective probability for single events, which I find rather unintuitive, or that we follow the concept of relative frequencies to come up with some ratio. Subjective risk would mean that we assign subjective probabilities, i.e. expectations that depend on what we believe to be the case or what will be the case. In the case of risk (3), an objective risk would mean that the evaluation is objectively true, i.e. that a given event would, objectively, cost a certain number of lives, if “lives lost” is our evaluative standard. Subjective value would mean that we expect it to cost those lives. Finally, risk (4) might be purely objective, subjective or hybrid.

“Objective risk” does not mean that no subject’s activity is involved in assigning it, however. To the contrary, because properties are not self-attributing, there would be no instance of any concept objective or subjective without a subject that performs attributions. It is just that the attributions are meant objectively, to be true or false in virtue of the things that there are. “Subjective risk”, on the other hand, does not mean that the standards of attribution become subjective, but whether that which is attributed is a property, ultimately, of the subject and not the object. That there is subjective risk (3) e.g. does not mean that one person measures value in dollars and the other one in lives lost. But it means that we talk now about *expecting* something to cost some amount of money if it happens, or some number of lives. So even though evaluation needs an evaluative framework which might be

man-made, like money, it might nevertheless be an objective fact that some damage costs us some amount of money to repair it – whether we know it or not.

Because we as human beings have no direct access to the world as such, in its pure objectivity, the distinction between objective and subjective risk might seem purely academic. But it is not. First, objectivity can be made sense of as to be the asymptotic limit of intersubjectivity. Second, subjectivity can be said to be a matter of degree: Combining both points, we can say that some risks may be more objective or more subjective in character depending on the level of intersubjective conformation of the supporting beliefs or values.

“Perceived risk” is ambivalent. It can either be understood as a demarcating label, such that there are perceived risks and other, maybe: real, or hidden, or whatever risks. But perceived risk can also be understood as an explanatory label, such that every risk is a percept, maybe even a perception of a perception (or of perceptions). In any case, speaking of perception is presupposing that there is something to perceive, and that we might perceive it (and perceive it correctly or adequately) or not. Understood as such, it draws from the same distinction as objective and subjective risk. There are three options to further specify what is perceived: It may be either objective risk, or any other objective entity, or a subjective entity. To go with the first option would preclude to hold that every risk is subjective, of course. The latter two options are compatible with both the explanatory and the demarcating understanding of “perceived risk”.

As for “perceived safety”, this expression usually does not mean that adverse events are believed to be impossible. If German politicians say “Our nuclear power plants are safe!” or “Our dams are safe!” it is clear that they cannot mean that literally; they can only mean that either the probability that things go wrong is close to zero, or that the possible damage is

close to zero, or both. And we may believe this or not. But they cannot mean that one (or both) are really zero. Perceived safety, in this sense, would always be a misperception. A better alternative would be to understand safety as practical safety: Something is practically safe when anything that might go wrong is either so improbable or causing so little damage that it makes no difference to our choices. Some people might live in perceived safety in this sense; they think that they need not worry. And while this is not always a misperception, it is well known that it sometimes is.

Risk and uncertainty, hypothetical and meta-risks

Depending on the chosen definition of risk, it might make sense to differentiate between risk and uncertainty. Certainty and uncertainty are subjective concepts, so uncertainty and *objective* risk are not of the same kind anyway. It makes more sense to compare subjective risk and uncertainty. We can be uncertain about many things, but for our purposes it might be useful to say that “uncertainty” is more radical than risk. Sometimes uncertainty is used to describe anything that is less determinate than risk (4). So, if we know that something might happen but do not know its probability or value, we have uncertainty and not risk (4). But what we have, then, can be easily spelled out taking the other definitions of risk (1-3) into account. Namely, if we know that something might happen, but not the probability or value, then this is the situation of risk (1). If we know its probability or value, but not both, this is risk (2) or (3) resp.

So, the demarcation between risk and uncertainty is relative to the involved definition of risk. For any risk (n), what is less certain turns out to be a risk (m) with $m < n$. Only for something that is less certain than risk (1) we need a new concept. This could mean that we do not expect that something bad or good might happen at all, or that something bad or

good of a certain kind might happen (so that we can at least qualitatively distinguish it from other possibilities, i.e. in its essence, but not in terms of probability or value). For the latter, I would thus suggest to use the term “essential uncertainty”, and for the former the term “complete uncertainty” or “ignorance”. Please note that ignorance does not mean that we do expect that nothing at all happens, it means that we expect something to happen for sure (business as usual, so to say) and thus we do not expect that something *might* happen (or not happen), or, in other words, we do not expect any surprises at all. Essential uncertainty means that, after something unexpected has happened, we would be saying: “I did not expect that it might go wrong in *that* way”. Ignorance, however, means that we would have to be saying: “I did not expect that anything *at all* might go wrong here”.

These extents of uncertainty have to be distinguished from orders of uncertainty. First-order uncertainty about something, say a damage extent, means that we do not know the exact value that quantifies it. Sometimes we can quantify this first-order uncertainty by using tools and concepts from statistics, e.g. by using error bars. Second-order uncertainty would mean, however, that we do not know how exact we know (or not know) the value in question.

A further distinction, that is cross-cutting to those explained so far, is that between real, hypothetical and meta-risks (cf. Hubig 1993: 75). A risk is real, in the sense of this distinction, if the domains of definition are well-known and the events under scrutiny uncontroversially fall in these domains. A risk is hypothetical, however, if this is not the case. The events under scrutiny are only theoretically described (and we lack the experience that would reassure us that we got the theory right), or that are only qualitatively observed such that quantities have to be estimated etc. – and hence we cannot be sure that the events will be confined to the domains of definition. That a risk is hypothetical, in the sense just introduced, does not mean that we should take it less seriously or so, it means that we do not really know what to

expect. In the case of meta-risks however, the event under scrutiny has the potential to change the domains of definition for other risk evaluations, such that many other risk assessments will have to be revised in a typically quite unclear way. It should be clear that these risks must not be neglected. An example for a hypothetical risk would be the release of genetically modified organisms into the wild when we know about their behaviour only from model calculations, theoretical estimates and analogies to other organisms' behaviour in the wild. Examples of meta-risks would be the triggering of massive climate change, the creation of a hybrid species between animal and human, or of some singularity-like self-enhancing artificial intelligence.

Risk and choice

It has been suggested to distinguish between risk (Risiko) and danger (Gefahr): While danger, it is said, concerns events that might happen independent of what we do, to speak of something as a risk means that our own decisions, our choices, are involved (cf. Evers and Nowotny 1987). In a social setting, for those that decide something may appear as a risk what for those affected appears as a danger (Luhmann 1991). The neat example, following Luhmann 1993, is that before there were umbrellas, given that I had to go out, there had been a certain danger to get wet. With an umbrella at hand, this danger turns into a risk. Now I can lower the risk of getting wet, by taking the umbrella with me, but I run other risks with that, e.g. the risk of losing the umbrella. While this story nicely illustrates one of the many repercussions and side-effects of technology, or more general, of the use of means, and the risk/danger distinction may be helpful for social theory, I do not think that ordinary language use warrants the use of the labels "danger" and "risk" for it. After all, we say that the risk of a fatal hit of the earth by a meteor is so and so (we mean: the probability, and

may presuppose that “fatal”, i.e. the damage under scrutiny, is the extinction of mankind). Now, one can say that we should not use the word “risk” here, but then I do not think that “danger” would be a viable alternative. Yes, we can say that this danger is greater than that danger. But nobody would say: “The danger of ... is 0.01 per year” or “The danger is 100 EURO”) or the like. Rather, the best alternative to risk jargon would be to say “The probability of ... is 0.01 per year” or “The possible damage is 100 EURO”. So, danger talk cannot fully substitute risk talk *regardless* of whether what may happen is dependent on our choices or not. Thus, I do not see that danger and risk are alternative concepts, but rather that for some negative risk, in the sense of risk (1), we also use the term “danger”, as well as for some undifferentiated risk (2-4)-talk. Hence, I will stick to the risk-jargon.

Independent of such labelling problems, the connection between risk and choice has been thoroughly discussed in economics and ethics in terms of rationality. By definition one should try to avoid bad events, so vis-à-vis *negative* risks the rational behaviour is aversion: After all, who would prefer a 50/50 chance of losing 1 Euro (i.e. a 50% chance of losing the Euro and a 50% chance of keeping it) over a 100% chance of not losing or gaining any money? When economists speak of “risk aversion” (cf. Kahneman and Tversky 1979), however, they refer to what I called *speculative* risk. Then aversion means to shy away from uncertain outcomes just because of their uncertainty. So if you are risk averse in the economists’ sense, you would prefer a 100% chance of not losing or gaining any money over a 50/50 chance of losing *or gaining* 1 Euro. If you are in a situation where you cannot afford to lose 1 Euro (e.g. because you need it to buy your next meal) but one more Euro would not make much of a difference to you, it would be rational to be risk averse. Whereas, when one Euro less would not harm you but 1 Euro more would make a big difference (e.g. with just

this Euro more in your pocket you could buy what you always wanted), it would be rational to be risk seeking.

This rational risk seeking (or risk aversion) behaviour is not to be confused with the pleasure (or disgust) that some people might find in gambling. Because if you love safety for its own sake and hate gambling, you might prefer not to gamble at all over a gamble that promises you only wins. Or, if you really like gambling as such, you would be willing to gamble in situations where, say, you have a 49/51 chance of winning or losing a Euro, or where every outcome makes you lose something (but the grief from losing some money is outweighed by the intrinsic pleasure from gambling). So, if one does not confine rationality to making or maximising profits (in some narrow sense, e.g. of winning or losing money), this latter behaviour is also not irrational. It is then guided by *intrinsic* risk aversion.

Rational choice under risk and uncertainty

In situations that can suitably be described as risk (4), i.e. where the values and probabilities of the options are known, and the according demarcations, divisions and evaluations are uncontroversially adequate and inclusive, i.e. that all relevant short- and long-term effects are accounted for, all ex-ante as well as ex-post costs, e.g. for precautionary measures or damage repair, are included in the evaluation, and also subjective costs due to persistent fear or dread etc. – in such an ideal case, risk (4) can be interpreted as the expected value of the event. Rational choice theory then tells us that it is rational to go with positive expected values, or in the case of multiple options, that it would then be most rational to always carry out what promises the highest expected value (sometimes dubbed the “Bernoulli-Principle”), cf. Mongin 1997.

Real-world risk calculations however, typically are considerably more narrow in scope. In them, one looks only at certain types of consequences (e.g. life, or property), and only the negative consequences (loss of life, or property), and aggregates along quite limited dimensions (typically: by counting loss of life, or by adding up the value of property loss, or by both combined with an exchange rate based on willingness-to-pay or welfare loss, cf. World Bank 2016, 47ff.). With respect to expected value, these calculations are necessarily incomplete, for above all they do not include the chances: They are only expressions of negative risk, not of speculative risk, and moreover only of certain aspects of it. Because of this, not accepting different but equally high risks *in any such calculatory framework* cannot normally be regarded as irrational, but may well be a very rational attitude of encompassing the aspects that are neglected by the framework.

Moreover, every interesting real-world calculation of a risk (4) will fall more or less short even of the narrow-scope ideal. Typically, it has to deal with imprecise numbers, estimates and guesses. Most challenging are events that are rare or that occur within complex systems (more on that below), because then the proper basis of experience is missing or the causal analysis does not work, resp. There are some strategies of dealing with such uncertainties that are worth explaining in more detail. *Probabilism* would attempt to assign probabilities and values despite of these uncertainties, and then decide for one of the options. It comprises a variety of substrategies that differ in how to do the assignment. One such substrategy would be *probabiliorism*, which consists in identifying the most probable consequences and act as if they were certain (and not only quite probable). Indicators of such a strategy are justifications that include phrases like “it is virtually certain that...” or “it is beyond reasonable doubt that...”. Another substrategy would be “Bernoulli-Maximisation”: Here we assign equal probabilities to all unknown consequences, deal

similarly with damage extents, and do then follow the Bernoulli-Principle (i.e. choose the option with the highest expected value). *Tutorism* is an alternative to probabilism, though. Here we do not assign any probabilities at all, but look only at the values of the consequences, i.e. the extent of damage. We then try to avoid the option (or those options) that would bring maximal damage, or would choose those options that would bring the maximal relief from suffering. The maxim recommended by the German philosopher Hans Jonas to give priority to the worst outlook (“Vorrang der ungünstigsten Prognose”, Jonas 1979) can be seen as an example of this strategy. To judge a measure only with regards of how it harms or benefits those that are worst off (as the Rawlsian “difference principle” suggests, cf. Rawls 1999), or most vulnerable, in a society would also be a substrategy of *tutorism*.

As soon as hypothetical risks or meta-risks are involved, all risk quantifications are becoming more and more insignificant and inadequate. In these cases, qualitative rules or scenario-based methods might provide better guidance. Ideally, we can then characterize the risk and identify best ways of dealing with them vis-à-vis the full spectrum of uncertainties. One proposal of a set of qualitative rules comes from the WBGU 1999, which takes into account second- and third-order uncertainties and uses criteria such as ubiquity, persistence, irreversibility, delayed effect and mobilisation potential for the characterisation of risks and the recommendation of adequate strategies. Scenarios, finally, can help to structure the possibility space and thus help to identify or construct meaningful options that can then be qualitatively or (sometimes) quantitatively be compared. Especially for meta-risks this may sometimes be the only viable option.

Risk acceptance and risk acceptability

Risk acceptance is a descriptive term that refers to observed behaviour (which risks people take) or judgements (which risks people are willing to take). In social settings, however, my choices affect others and I am affected by theirs – and I know that. This creates some complexity in understanding risk behaviour/judgements. Very often risk decisions are to some degree *allocentric*, “other-directed”, i.e. those that decide about an action are not identical with those that are affected by the possible costs or benefits of this action. The risk-taker is not identical to the risk-sufferer, or so one may say. And, to a certain degree, each side takes into account the perspective of the other. In attributing risk perception, the question would be then whose perceptions are involved (or, more generally, to which subject the subjective components shall be related). Another way in which a risk may be “social” is when the possible benefits and the possible harms do not affect the same persons. The risk may be *asymmetric*, i.e. individual risk-balances may be diverging. And again we have to ask to which subjects any subjective elements shall relate.

Every actor knows that he or she sometimes decides for others (e.g. children) which risks to take, and that those that may benefit are sometimes not those that may suffer. Thus, in our judgements about which risks we want to accept as well as in our behaviour in dealing with risks, we have taken the social dimensions of risk somehow into account, i.e. we have taken some stance towards the other. Empirical research in psychology and sociology may help to describe risk behaviour and judgements, but it not easy to interpret these findings, because of the complexity of decisions under risk and uncertainty.

Acceptability of risk, on the other hand, is a normative term. It means possible acceptance, but not in a prognostic sense, but in the sense of meeting considered judgements about legitimate acceptance. These judgements are normative judgements because they involve either prudential evaluations (about what is good for me, or for us as a group) or moral

evaluations (about what is good, or even equally good, for all of us). Prudential evaluations are, in its general form, linked to conceptions of a good life – and it is in the light of these conceptions that strategies of dealing with risk and uncertainty (including the strategies explained above) are chosen. Psychology and sociology can help to understand these links; philosophy should also be able to contribute but is traditionally more occupied with moral evaluations. These evaluations are mainly discussed within the frameworks of the grand moral theories, utilitarianism, deontology and contractualism. Utilitarianism typically see the maximization of the sum over everybody's expected value as the right thing to do. It is then sensitive only to the sum, not to the distribution (between individuals) of such value, and it may well be that risk decisions make some people worse off as long this is compensated by making some other people better off. This approach connects to economic theory.

Deontology, however, operates with indexical rights of individuals that do not allow for such calculus. The infringement of rights has to be justified by the exertion of some other, equally high or higher-indexed right, and that there are probabilities involved does not change this principle (Nida-Rümelin 1996). So e.g. if the right to live is ranked higher than the right to property, it is justified to impose risks of property loss on some people if this increases the chances of other people to survive. While in some constellations it is still allowed to come to risk decisions that impose risks on some to the potential or immediate benefit of some others, equity considerations now require that those that are being made potentially worse off have to be compensated by those that benefit, or at least those that are being made actually worse off (if things go wrong and damage occurs). This approach connects to legal theory. Contractualism relies not on calculus or on antecedent rights to determine what is allowed or required to do, but on the agreement of those that are affected. The individuals shall be put in a position where they can make an informed decision which risks they are

willing to take (vis-à-vis certain compensations etc.), and where this is not individually possible, where they can control the institutions that govern the risk that do affect them (Renn et al. 2007). This approach connects to the theory of democracy (Shrader-Frechette 1991) and has led to a variety of models of citizen participation in shaping collective risk decisions (Gottschalk and Elstner 1997). It can be grounded in “discourse ethics” (ibid.; Skorupinski and Ott 2000) that contains deontological and if not utilitarian then at least consequentialist elements.

Risk governance, systemic risks and the precautionary principle

“Risk governance” has more recently been introduced as a label for the wide-scope attempt to deal with risks in a rational and equitable way (Renn 2008) following the lead of the US National Research Council (NRC 1996). It consists of four consecutive phases: Pre-assessment (framing, early warning, screening), appraisal (estimation of hazards and exposures, assessment of risks), characterization & evaluation (see the last sections of this text), management (selection of options). Communication, in this framework, is a cross-cutting task in all four phases.

Risk governance in modern societies has become quite challenging: Modern societies are organized as a network of very sophisticated institutions, are relying in many ways advanced technology, and are long-term interfering with the natural environment. Risk decisions often have consequences that run through all three of these social, technological and environmental systems, which are coupled in many ways. Accordingly, concepts of “systemic risk” have been introduced. That risks occur (and have to be governed) in these systemic structures has been pointed out by the International Risk Governance Council: “Systemic risks are at the crossroads between natural events (partially altered and amplified by human

action such as the emission of greenhouse gases), economic, social and technological developments and policy-driven actions, both at the domestic and the international level.” (IRGC 2005, 81). On the other hand, some of these systems are itself essential for the fulfilment of basic human needs. The OECD has pointed out this aspect: “A systemic risk ... is one that affects the systems on which society depends – health, transport, environment, telecommunications, etc.” (OECD 2003, 30). So, “systemic risk” may mean that risks occur due to systems (they are creating risks) as in the IRGC definition, or that systems are in danger due to risks (they are affected by risks) as in the OECD definition. In fact, both may be the case, due to the multitude of systems involved.

As the major challenges of modern societies in governing risks, it has been suggested to distinguish complexity, uncertainty and ambiguity (Klinke and Renn 2002; IRGC 2005). Complexity means that system behaviour involves feedbacks and triggers effects that makes it hard to predict medium- and long-term behaviour and to identify causal relations, i.e. describe them in terms of if-then-statements (which would be needed to discuss interventions). Uncertainty, used in contrast to complexity, points out not the objective complexities of these systems, but the subjective ignorance in determining the exact states of these systems, and their change with time. The interplay of complexity and uncertainty is well known from physics, where simple models already show deterministic chaos: They render very complex dynamic results that change qualitatively due to only minor changes in the control parameters, such that that any uncertainties about these parameters make these systems then practically unpredictable.

Ambiguity, finally, shall be an umbrella term for all evaluative and normative degrees of freedom. Modern societies are pluralistic, so common values for collective decisions are hard to find. On top of that, risk-specific ambiguities occur, i.e. due to diverging levels of risk-

adversity. On the level of moral theory, a similar pluralism occurs (that of utilitarianism, deontology and contractualism, e.g.). A major problem is the time horizon of many risk decisions (among the most prominent here are those concerning nuclear energy). It is disputed whether future events shall be discounted and if so with which rate. Moreover, evaluations may change with economic and scientific progress. What we see as worthless today may become a precious resource in the future, and what we see as harmless may be considered dangerous. Finally, it would be parochial to assume that values and moral ideals will not change with time, but we have no idea how they will change. However, we cannot simply take the preferences of future generations to be identical to our own preferences – we should allow them to form their own preferences. This makes it very hard to meaningfully evaluate risks - even if all natural consequences were known.

Suitable reactions to this threefold challenge include openness and explicitness in risk-governance. For risk identification and characterisation, this means that the underlying uncertainties are not ignored (e.g. when some parameter is not known exactly then a plausible distribution of parameter values should be used and not just the best estimate) and that they are made explicit (e.g. by depicting error bars for any quantity, or by adding a verbal commentary that points out known unknowns, simplifications, second-order uncertainties etc.). Underlying ambiguities should be handled in a similar way: By making divergent evaluations explicit, by laying open diverging or converging normative standards due to competing moral theories, etc. Openness and explicitness are essential components of the ideal of a “discursive” risk governance (Renn et al. 2007, 234).

Moreover, it is often said that we should follow a precautionary principle. This principle can be interpreted in different ways, though. In its most simple form, it says that we should try to avoid taking risks when the potential damage is catastrophic. These risks are, in the sense

outlined above, meta-risks and the principle recommends a tutorist strategy. Such catastrophes threaten our agency, i.e. narrow down our future options and our abilities to identify and evaluate them – we would merely be able to react to external circumstances in a daily struggle for survival and can no longer properly plan ahead (or, for that matter, assess risks) at all. In a more complex form, the precautionary principle says that we should also try to avoid taking severe risks, i.e. those with high damage potential, when the potential gains are in the same order of magnitude. This can be seen as an answer to complexity, uncertainty and ambiguity, where we better be careful if we do not really know what we are doing. A reason for this may be that what we see as potential gains (e.g. what we find pleasure in) is usually more variable than what we see as potential loss (e.g. life and limb, avoiding pain). And while we do not know what future generations will find enjoyable, we know well enough what will make them suffer. Both versions of the precautionary principle can be justified within any of the discussed moral theories, not only within deontology, although different reasons would have to be given and different concrete advice of levels of extra care would result.

Sometimes the term “precautionary principle” is used in a less abstract sense, however, namely as a label for regulative strategies that allow any innovation only if proven to be (relatively) harmless. This is combined with a tort law that allows only for minor compensation claims in case that (by definition) unforeseen risks materialize. As opposed to that, a “risk principle” would allow any innovation if not proven to be (relatively) harmful. This is combined with hefty compensation claims, then. Both models allow to deal with an open future and set checks and balances against irresponsible or egoistic risk behaviour. The former is more in line with deontology and the latter with utilitarianism, and it is sometimes said that Germany (or the EU) would follow the former and the US the latter principle. What

makes no sense, however, is to combine the liberal admission policy of the risk principle with restrictive tort law, i.e. only minor compensation claims. On the other hand, restrictive admission plus high compensation claims may mean forfeiting too many opportunities. Thus it might be difficult to harmonize regulative frameworks across cultural borders.

Risk Cultures

Ultimately, this is what any society and any of its individuals has to acknowledge: That *nolens volens* something has been established that I would like to call, in a broad sense, a “risk culture”. Culture, in this analysis, has three dimensions, a physical, an informational, and a social. In dealing with risks such as flooding (cf. Gottschalk-Mazouz 2006), the physical dimension includes dams, water buffers etc., the informational dimension includes public broadcasting and education, and the social dimension includes civil protection organisations, insurances, donations etc. Typically, only the adequate interplay of all the components in the three dimensions allows for a society, and an individual as its part, to adequately deal with risks. But for every culture, it is not only the static aspect (what the resp. risk culture is like at the moment) but also the dynamic aspect (how it develops or is being developed) that deserves attention when dealing with risks: Because risks change over time, new risks emerge and old ones disappear or become less relevant, and because social, technological and environmental conditions change as well, risk cultures are (or should be) in continuous transformation.

This “risk culture” is not the same as the famous “risk society” hypothesis that Ulrich Beck has taken our modern societies to be. He has been suggesting (1992, 21) that risks are a “systematic way of dealing with hazards and insecurities induced and introduced by modernisation itself” and that we are thus on a way to a “New Modernity”. That risks are

induced and introduced (but also mitigated and absorbed) by modernisation may well be. But ever since man has been keeping stock, has been fabricating its own tools, has been communicating about adverse events and has been organised such as to collectively absorb them, there have been risk cultures. And in those cultures, risks have been transformed. Naturally, people would not have used these terms. And maybe they would have described what we take to be a risk as something else, some divine trial maybe or something that could not have been altered (and has been predetermined) anyway. It is true that to develop *deliberate* risk cultures, modern concepts of human agency and choice, of nature and society are required. But as soon as these concepts were at hand, people have also using the concept of risk (and, in fact, also the term “risk”; cf. Rammstedt 1992) to develop such deliberate risk cultures.

References

- Beck, Ulrich (1992): *The Risk Society. Toward a New Modernity*. London. (Orig.: *Risikogesellschaft. Auf dem Weg in eine andere Moderne*. Frankfurt/M. 1986).
- Evers, Adalbert and Nowotny, Helga (1987): *Über den Umgang mit Unsicherheit: Die Entdeckung der Gestaltbarkeit von Gesellschaft*. Frankfurt/M.
- Gottschalk-Mazouz, Niels (2011): Risiko. In: M. Düwell; C. Hübenthal; M. H. Werner (Eds.), *Handbuch Ethik*, 3rd ed. Stuttgart, 502-508.
- Gottschalk-Mazouz, Niels (2008): Risikokulturen, in: J. Köngeter (Ed.), *Sicherheit und Risiko wasserbaulicher Anlagen*. Aachen, A1-A22.

Gottschalk, Niels and Elstner, Marcus (1997): Politik und Technik: Überlegungen zu einer innovativen Technikgestaltung in partizipativen Verfahren. In: M. Elstner (Ed.), Gentechnik, Ethik und Gesellschaft. Heidelberg, 143-179.

Hubig, Christoph (1993): Technik und Wissenschaftsethik: ein Leitfaden. Berlin/New York.

International Risk Governance Council IRGC (2005): Risk Governance: Towards an Integrative Approach. White Paper #1. Geneve.

Jonas, Hans (1979): Das Prinzip Verantwortung. Versuch einer Ethik für die technologische Zivilisation. Frankfurt/M.

Kahneman Daniel and Tversky, Amos (1979): Prospect theory: an analysis of decision under risk. In: *Econometrica* 47, 263–91

Klinke, Andreas and Renn, Ortwin (2002): A New Approach to Risk Evaluation and Management: Risk-Based, Precaution-Based and Discourse-Based Management, In: *Risk Analysis* 22(6), 1071–1094

Luhmann, Niklas (1991): *Soziologie des Risikos*. Berlin/New York.

Luhmann, Niklas (1993): Die Moral des Risikos und das Risiko der Moral. In: G. Bechmann (Ed.), *Risiko und Gesellschaft. Grundlagen und Ergebnisse interdisziplinärer Risikoforschung*. Opladen, 327-338.

Mongin, Philippe (1997): Expected Utility Theory. In: J. Davis; W. Hands; U. Mäki (Eds.), *Handbook of Economic Methodology*. London, 342-350.

Nida-Rümelin, Julian (1996): Ethik des Risikos. In: J. Nida-Rümelin (Ed.), *Angewandte Ethik*. Stuttgart, 806-831.

National Research Council NRC (1996): Understanding risk. Informing decisions in a democratic society. Washington DC.

OECD (2003): Emerging Systemic Risks in the 21st Century: An Agenda for Action. Paris.

Rawls, J. (1999): A Theory of Justice (Revised Edition). Oxford.

Renn, Ortwin; Schweizer, Pia-Johanna; Dreyer, Marion; Klinke, Andreas (2007): Risiko. München.

Renn, Ortwin (2008): Risk Governance. Coping with uncertainty in a complex world. London.

Shrader-Frechette, Kristin (1991): Risk and Rationality. Berkeley CA.

Skorupinski, Barbara and Ott, Konrad (2000): Technikfolgenabschätzung und Ethik. Zürich.

Rammstedt, Ottheim (1992): Risiko. In: J.Ritter; K.Gründer (Eds.), Historisches Wörterbuch der Philosophie. Basel/Stuttgart, Vol. 8, 1045-1051.