

# Chapter 2

## Levels of Uncertainty

Hauke Riesch

**Abstract** There exist a variety of different understandings, definitions, and classifications of risk, which can make the resulting landscape of academic literature on the topic seem somewhat disjointed and often confusing. In this chapter, I will introduce a map on how to think about risks, and in particular uncertainty, which is arranged along the different questions of what the different academic disciplines find interesting about risk. This aims to give a more integrated idea of where the different literatures intersect and thus provide some order in our understanding of what risk is and what is interesting about it. One particular dimension will be presented in more detail, answering the question of what exactly we are uncertain about and distinguishing between five different levels of uncertainty. I will argue, through some concrete examples, that concentrating on the objects of uncertainty can give us an appreciation on how different perspectives on a given risk scenario are formed and will use the more general map to show how this perspective intersects with other classifications and analyses of risk.

I beseech you, in the bowels of Christ, think it possible you may be mistaken (Oliver Cromwell, addressing the Church of Scotland, 1650) (From Carlyle [1871](#)).

### Introduction

What we mean by risk is not a clear issue because many writers use the word with slightly different meanings and definitions, even beyond the more vague everyday usage of the term. Aven and Renn ([2009](#)), for example, have found ten different

---

H. Riesch (✉)

School of Social Sciences, Brunel University, London, UK  
e-mail: [hauke.riesch@brunel.ac.uk](mailto:hauke.riesch@brunel.ac.uk)

definitions they gathered from the wider risk literature. The problem of clear terminology continues if we go into the various classifications and clarifications of risk and uncertainty, with scholars distinguishing between risks, uncertainties, indeterminacies, ambiguities, and levels, objects or locations of risk and/or uncertainty. With this in mind, I feel slightly apologetic about writing about another scheme devised by myself and David Spiegelhalter (Spiegelhalter 2010; Spiegelhalter et al. 2011), where we use, again, our own terminology, this time in trying to distinguish between different things we can be uncertain about. In this chapter, I will try to explain our distinctions and where they correlate and/or fit in with other classifications of risk and uncertainty, as well as provide an argument on why we feel this particular classification adds to the literature on risk theory by going through a couple of real-world examples.

As Norton et al. (2006) note in their reply to the paper by Walker et al. (2003) discussed below, “an important barrier to achieving a common understanding or interdisciplinary framework is the diversity of meanings associated with terms such as ‘uncertainty’ and ‘ignorance,’ both within and between disciplines” (Norton et al. 2006, p. 84). The proliferation of what we mean by risk and how we categorize it within the literature is partly due to the different agendas the different disciplines have with regard to the topic. The question “what do we want to know about risk?” will be answered differently by scholars, for example, interested in risk perception and those interested in the “risk society.” Asking this question explicitly may help us in finding out where the different disciplinary approaches to risk intersect. Our classification is partly intended to do just that, mostly because I (as a sociologist) and David Spiegelhalter (as a statistician) have always had slightly different conceptions of what is academically interesting about risk, and our collaboration was partly an attempt to build a conception of risk which is useful for both social and scientific/technical disciplines and will be useful for communicating across this divide by giving a clear account of how and why, in Funtowicz and Ravetz’s phrase, there is “a plurality of legitimate perspectives” on risk (Funtowicz and Ravetz 1993, p. 739). At the same time, I hope it will provide a useful and relatively simple map through which the different academic disciplines’ interests in risk can be compared and connections seen more easily.

In this chapter, I want to advance the idea that one can confront the different meanings which risk is given and offer an idea of how they are related. There exist many other schemes that try to categorize risks such as Renn and Klinke (2004), Stirling (2007), or van Asselt and Rotmans (2002), and I will try to show how they fit into our overall picture in the following section. As a departing point, I take risk to mean roughly a function of the uncertainty of an outcome and its impact. This definition leaves room for plenty of uncertainties itself, especially since there is no agreement on how to measure impact, or how to compare impacts of completely different categorization, and there is plenty of literature in risk studies devoted to this problem.

The uncertainty part of risk however is itself very problematic. There are some uncertainties we can put a number on, some where we can only evaluate qualitatively and some we have absolutely no idea on how to even start evaluating

them. The classification I will propose here is meant to bring some order into the way we think about uncertainty and provide a way in which different types of uncertainty and its classifications can be, if not directly compared, at least brought under the same scheme. Comprehensive surveys of what uncertainties and risks really are and how they should be classified can easily lead to a rather complicated structure that becomes less useful as a heuristic tool for people working within risk. This is more so on the social, policy, and communication aspects than in the technical risk assessment areas, for whom such schemes will be more useful, and I will concentrate on the former in this chapter. Out of the many different dimensions in which uncertainty can be categorized, we chose one in particular which we believe is most helpful when we seek to understand how different people and groups conceptualize and react toward risks. It is meant to analyze risks according to the following question: What kind of thing exactly are we uncertain about?

## Background

Philosophical classifications of probability have traditionally focused on questioning where our uncertainty derives from, with the two main choices being uncertainty inherent in the system, and uncertainty arising from our incomplete knowledge. These two interpretations of probability are named by philosophers (Hacking 1975, see also Gillies 2000) *epistemic* probability and *aleatoric* (also often called *ontological* or *ontic*) probability. This basic distinction still underlies modern philosophical theories of probability and can be seen, for example, in the philosophical split between Bayesian (subjective) and frequentist (objective) interpretations of probability in statistics (see also Gillies 2000).

Uncertainty in a larger sense, as opposed to the mathematically defined concept of probability, has also seen attempts at classification. An early and very influential distinction came from Frank Knight, who distinguished uncertainties which are quantifiable which he called risks, and those that are not quantifiable, which he called uncertainties:

The essential fact is that 'risk' means in some cases a quantity susceptible of measurement, while at other times it is something distinctly not of this character; and there are far-reaching and crucial differences in the bearings of the phenomena depending on which of the two is really present and operating. [...] It will appear that a measurable uncertainty, or 'risk' proper, as we shall use the term, is so far different from an unmeasurable one that it is not in effect an uncertainty at all (Knight 1971 [1921]).

This classification has proved to be very influential especially among sociologists, but is in my opinion slightly unfortunate as it propagates confusion with the traditionally defined concept of risk equaling probability times outcome (or, in the more modern sense focusing on negative outcomes, probability times harm). Although I recognize the usefulness of Knight's distinction for this work, to avoid confusion I prefer to work with the conception that risk refers to a measure of uncertainty combined with the potential outcome.

Combining these two perspectives in a sense, Stirling (2007) recently proposed to divide both the uncertainty as well as the outcome aspects of risk into “problematic” versus “unproblematic” in a similar way to which Knight distinguished between quantifiable and unquantifiable uncertainty. This results in a two by two matrix: at the corner where the probabilities as well as (our knowledge of) the outcomes are unproblematic there are risks associated with the typical statistical risk analyses such as Monte Carlo simulations or costbenefit analyses—these scenarios he terms “risks” in the traditional sense used by most scientists and risk analysts. Scenarios where the probability is knowable, but we are more unsure about the outcomes, he terms “ambiguities”; risks where conversely the outcomes are unproblematic but the probabilities are, he calls “uncertainties.” When neither are unproblematic, he talks about conditions of “ignorance.” It is worth also pointing out that the term “ambiguity” is used in other disciplines, for example, behavioral economics, to mean unknown probabilities, which is almost precisely the opposite to Stirling’s sense—this demonstrates, again, the problems of terminology within the wider risk literature. Technology assessment on the other hand traditionally uses similar terminology but without taking Stirling’s ambiguity into account.

Stirling argues that dividing risks into these categories can give us guidance on the circumstances when the precautionary principle could be a valid rule: by dividing risks into qualitatively distinct groups, he argues that the principle can be an important rule for helping with decision making in those circumstances where the outcomes or probabilities are not well understood, and no other type of decision rule would otherwise be helpful.

Another influential attempt at classifying risk elaborated to inform risk assessment policy eventually evolved to inform Funtowicz and Ravetz’s very influential concept of postnormal science (Funtowicz and Ravetz 1990, 1993). Funtowicz and Ravetz proposed to map risks as a measure of uncertainty and impact (“decision stakes”) and claimed that risks with low uncertainty and impact are the ones familiar from applied science for which traditional mathematical tools of risk analysis are most appropriate. Risks with medium but not high uncertainty and/or impact are in the domain of “professional consultancy,” which “uses science; but its problems and hence its solutions and methods, are radically different” (Ravetz 2006, p. 276). The label “postnormal science” applies to situations characterized by high uncertainties and/or high stakes.

Renn and Klinke (2004) similarly use this map with axes denoting uncertainty and impact and identify several areas on that map that delineate qualitatively different risk situations, though these depart from Funtowicz and Ravetz’s three areas on the map by being more fine grained: For example, the points in the map where the probability is low but the potential harm is great, they call “Damocles” risks, named after the Greek king who according to the legend had a sword suspended above him by a thin piece of string (the analogy being that the probability of the string breaking at any one point in time is low, but when it happens, the outcome is rather dramatic, at least for Damocles). Points with high probability and high harm they call “Cassandra” risks, after the Trojan prophet who knew

about the fate of the city but whose warnings were ignored. Hovering more in the background is a larger area of the map, where we are not very knowledgeable about the event's probabilities or its outcomes ("Pandora" risks).

Brian Wynne introduced his classification of risks as an improvement on the Funtowicz and Ravetz (Wynne 1992) classification which defines postnormal science. Like Stirling, Wynne sees "risks" as situations where the outcomes and the probabilities are well known and quantifiable. Uncertainties are present when "we know the important system parameters, but not the probability distributions" (p. 114). By contrast, the next level, "ignorance," is more difficult to define: "This is not so much a characteristic of knowledge itself as of the linkages between knowledge and commitments based on it" (p. 114). It is "endemic to scientific knowledge" (p. 115), because science has to simplify what it knows in order to work within its own methods. Finally, "indeterminacy" is seen as largely perpendicular to risks and uncertainties, because it questions the assumption on the causal chains and networks themselves. Thus, indeterminacy can be a huge factor in a particular situation even when the risks and uncertainties are judged to be small.

I am sympathetic to Wynne's classification because it recognizes that both quantifiable types of uncertainties as well as the less tangible deeper uncertainties are present at the same time in some situations and thus not mutually independent, which is a necessary realization away from other schemes such as Funtowicz and Ravetz's map. According to Wynne,

Ravetz et al. imply that uncertainty exists on an objective scale from small (risk) to large (ignorance), whereas I would see risk, uncertainty, ignorance and indeterminacy as overlaid one on the other, being expressed depending on the scale of the social commitments ('decision stakes') which are bet on the knowledge being correct (Wynne 1992, p. 116).

However, there are for me still some problems with it. First, and more trivially, is the question of terminology. Like almost every other theorist of risk that comes from the social science side, Wynne and Stirling take "risk" itself to be one of their categories, and then proceed to label the other categories somewhat arbitrarily—this results in a mess of technical definitions that leave no special terminology for the overall thing they intend to classify. We cannot call them classifications of risks (or uncertainties) because risk and uncertainty are already part of the classification system. Moreover, this use of the term risk clashes somewhat with the common definition of risk as a measure that combines uncertainty and outcome. This has not helped that another influential tradition of risk theory embodied by Beck (1992) and Giddens' (1999) work takes risk to mean something altogether more nebulous.

Another concern over Wynne's classification, though, is that the categories seem somewhat hard to pin down, in the sense that indeterminacy, for example, includes the various social contingencies that are not usually captured in conventional risk assessments, but what these social contingencies are, and how they relate to the other types of uncertainties is not categorically stated. It is not

entirely clear, at least to me, where the boundaries lie, or even if there are supposed to be any precise boundaries. Ignorance, he writes, is “conceptually more elusive” and best explained through a lengthy example. All this in effect makes Wynne’s conceptualization hard to explain and therefore possibly ineffective as a tool for bridging the divide between the social and the technical aspects of risks. The inclusion of broad concepts such as social contingencies as well as quantifiability leaves the feeling that Wynne’s categories slice through several useful other distinctions on risk (such as those introduced below, in particular that of Walker et al. 2003). While Wynne’s categories are helpful as a conceptual tool to analyze reactions to risk and identifying shortcomings in conventional scientific approaches to risk that need to be addressed, it remains unclear exactly how they intersect and relate to each other. In a way, our own classification presented below is an attempt to reformulate Wynne’s insights in a way that makes more intuitive sense and which hopefully helps in addressing the question of how Wynne’s categories relate to each other.

Van Asselt and Rotmans (2002) classify risks according to the source of our uncertainty, distinguishing primarily between the two major sources introduced above of epistemic and aleatoric uncertainties (or, in their terminology, uncertainties due to lack of knowledge and uncertainties due to the variability of nature). Uncertainties due to lack of knowledge include, for example, lack of observations/measurements, inexactness or conflicting evidence, while uncertainty due to the variability of nature includes variability in human behavior, value diversity, and the inherent randomness of nature. Aiming to go further than this, Walker et al. (2003) include more dimensions in the classification than merely the source of uncertainty. Thus, they distinguish between location, level, and nature of uncertainty: the location uncertainty can be subdivided between context, model, input and parameter uncertainties, and the final outcome uncertainty. Location uncertainty therefore roughly describes what we are uncertain about, i.e., “where uncertainty manifests itself within the whole model complex” (p. 9). The levels of uncertainty describe the “progression between determinism and total ignorance” and include, in order, statistical uncertainty, scenario uncertainty, recognized uncertainty, and total ignorance. Finally, the nature of uncertainty is, like in van Asselt and Rotman’s classification, mainly about the source of uncertainty, and can roughly be divided into epistemic and ontological uncertainties and subclassified as done by van Asselt and Rotmans.

In this chapter, I hope to be able to add a more inclusive categorization that stays within the spirit of Wynne’s as well as Walker et al. (2003) ideas but revolves more centrally around the question of what exactly it is that we are uncertain about, which roughly translates to the “location of uncertainty” dimension in Walker et al. This I will try to use to find interconnections between different literatures on risk. I will argue also that it is useful to apply the scheme to a selection of real-life uncertainties and use it to delineate and make sense of different groups’ varying assessments of a situation because they place different importance on the different objects of uncertainty that are all present to various degrees in all of the cases. I will start by making some preliminary distinctions

about risk and uncertainty which will enable us to see where this fits into the various other definitions and classifications of risk. I will borrow Walker et al. (2003) idea of different dimensions here, but add that, in our context, these dimensions can best be thought of as different answers to the question on what we want to know about risk.

Firstly, we conceptualize risk as a measure of uncertainty of an event happening times the severity of the outcome. As argued above, this is the usual definition of risk, though it is not used like this by all commentators, some of whom depart more from Knight's (1971) famous distinction between risks as quantifiable uncertainties versus uncertainties that are not quantifiable, which explains Wynne and Stirling's decisions to put "risk" as one of the categories within their overall schemes. Other writers such as those from the "risk society" tradition (Beck 1992 and Giddens 1999) use risk in a much more vague way which is not so much interested in quantifiable or nonquantifiable or even in the separation of uncertainty and severity of the outcome, but sees it more as the vague possibility that things can go wrong. This is again due to the fact that risk sociologists are interested in different aspects of risk (for example, how increasing awareness and preoccupation of risk affects late modern society). There is therefore not much point in criticizing some work for using vague definitions of risk because, from their point of view, there is simply not that much value added to having a precise working definition of what risk is. However, I hope to be able to show how our distinctions can contribute nevertheless to a better understanding of how the conception of risk that is seen as interesting to sociological and cultural approaches can be compared to other conceptions of risk.

Starting from the definition of risk being a measure of uncertainty and severity of outcome, it is secondly to be noted that neither uncertainty nor severity of outcome are in most cases easily measurable or even definable. Our scheme will leave the very interesting problem of severity of outcome for others to work out and concentrate specifically on the uncertainty aspect of risk.

Starting from the question of "what do we want to know about risk?" we can produce a table of different classifications of risk which are designed to answer that question in different ways. We may, for example, be interested in why we are uncertain, we may be interested in who is uncertain, how it affects individuals or society at large, how is risk represented and how should it be represented, and what is it exactly that we are uncertain about? These are the categories I use below, though there will possibly be more dimensions than those, and other authors may want to divide them differently (Walker et al. (2003), for example, distinguish between levels of uncertainty (whether we take a deterministic position or not) and nature of uncertainty (i.e., uncertainty seen as either aleatoric or epistemological), which I would both see as different sources of uncertainty (we can be uncertain *because* we take an epistemological stance and *because* we have an idealized deterministic situation).

*Why are we uncertain?* Here we can list classifications that have been made regarding the sources of uncertainty, such as in the scheme of van Asselt and Rotmans, which also relies on the philosophical distinction between

epistemological and aleatoric (or ontological) uncertainty described above: we can be uncertain either because of our lack of knowledge or because there is an inherent variability in nature. These two fundamental positions are often seen within the more sociological literature on risk as aspects that different situations of uncertainty can take on, so that, depending on the context, an uncertainty can be either epistemological or aleatoric: “it often remains a matter of convenience and judgment linked up to features of the problem under study as well as to the current state of knowledge or ignorance” (Walker et al. 2003, p. 13). In the philosophical literature, by contrast, it is more often assumed that the distinction is a result of different worldviews: we can, for example, be determinists in our general philosophical outlook, in which case, strictly speaking, all uncertainties are epistemic. In most everyday examples, the boundaries of whether an uncertainty should be considered epistemic or aleatoric seems to be a result of the setup, but the precise boundaries or even existence of the boundary to a large extent also depends on our philosophical stances and background assumptions and knowledge. We can, for example, see the probability of winning the lottery jackpot with a given set of numbers as purely aleatoric, because even with the most sophisticated current scientific methods, we are some way away from predicting the numbers drawn even if, philosophically, we are strictly speaking determinists who believe that an all knowing demon could calculate the final result from the initial state. In this example, the existence of probability that is for practical purposes aleatoric even for strict determinists is fairly obvious, though this is not necessarily the case in others. As I will argue below, there are other, epistemic, considerations to be made when we assess the likelihood of winning the lottery.

*Who is uncertain?* The question of the subject of the uncertainty is interesting from the point of view of psychologists or sociologists, who want to know what effect uncertainty has on people or on society at large. Different people respond to uncertainty differently, as shown, for example, in the well-known “white male effect” and similar phenomena discovered by risk psychology research (Slovic 2000). The subject of the uncertainty is also important for policy making since we would need to know how different groups and individuals respond to risks and representations of risk. For example, my current project investigates local opinions on energy infrastructure: to understand the dynamics of risk opinions within the area the infrastructure is being planned, we need to have a more detailed understanding of who the local actors and groups of actors are and how they interact with respect to interpretations of risk. Whether “the public” consents to the infrastructure being build in their back-yard ultimately depends on a complex interplay between local and national politicians, civil servants, the project developers, media representations, local and national NGOs and residents’ interest groups, as well as the individual resident’s understanding which is strongly influenced by, and in turn influences, the other stakeholders. Putting “the public” in scare quotes above is meant to signal that there is no monolithic public, with similar agendas, identities, or worldviews. Understanding who the relevant actors are and how they arrive at their conceptions of risks and how they influence and



are influenced by other groups of actors is vital for the analysis of what role risk plays in planning decisions.

*How is uncertainty represented?* Representations of uncertainty can take on different forms, which is again related to where the risk stands and is perceived along the other dimensions. We can, for example, simply deny that there is any uncertainty or risk at all or just concede that there is some, but more or less, undefined uncertainty. If we want (and know more about the situation), we can give a list of possible outcomes, either on their own or with some indication, qualitative or quantitative, on how likely each outcome would be. Should we have chosen a model we think is appropriate, we can give the result of the risk assessment as say a probability, with or without error bars or other representations of uncertainty on that final number.

How we represent risks depends very much on our knowledge of the situation, denying risk is a valid action when we do not know of any, and a simple list of possible outcomes is useful when we lack knowledge of how likely each outcome would be. However, which representation people chose in practice often depends also on what message they want to get across, or even reflects philosophical stances or implicit assumptions made. For example, if we want to make the risk of taking a particular medication look high, we can choose to represent it in relative rather than absolute terms. Similarly, we can give a positive or negative frame: for example, there is technically no difference between saying that “your chance of experiencing a heart attack or stroke in 10 years without statins is 10 %, which is reduced to 8 % with statins” and “your chance of avoiding a heart attack or stroke in 10 years without statins is 90 %, which is increased to 92 % with statins”—yet these two formulations have different connotations for the reader (example taken from Spiegelhalter and Pearson 2008). We can express probabilities in percentages or “natural frequencies,” where research has shown that people are intuitively better able to understand natural frequencies (Gigerenzer 2002). We can produce bar charts, pie charts, “smiley charts” on top of the verbal expressions, and these again convey different impressions of how risky something is. Finally, we can express uncertainties according to our philosophical understanding—if we say that I have a 10 % chance of having a heart attack within the next 10 years, that can either mean “10 % of people with test results like me will have a heart attack,” or “10 % of alternative future worlds will include me having a heart attack.” Again these scenarios while both expressions of the same amount of uncertainty will qualitatively feel different to people, with the second usually seen as the more persuasive way to get people taking their medicines, because it is more personalized (see also Edwards et al. 2001 on the effects of framing risk to patients).

*Responses to uncertainty:* How do people react to uncertainty? Do we, or should we, respond rationally to risk, for example, by doing a cost-benefit analysis to evaluate risks (Sunstein 2005)? Slovic et al. (2004) argue that while analytical and affective are two distinctive ways of reacting to risk, they interplay to produce rational behavior. But maybe even this distinction between affective and analytical needs to be challenged (Roeser 2009, 2010).

On a larger societal level, the risk society literature concerns itself, among of course other things, with how a society responds to risks (specifically our own, late modern—i.e., contemporary Western-society). Here, the issue is not so much about the nature of the risk as such (though it plays a role as I will outline below), or even whether the risks are real or not, but with the role that an increasing awareness of risk plays within late modern society. In particular, they describe the intuitive pessimistic induction through which people have come to realize (or at least believe) that there are always unexpected uncertainties and the possibility of things going horribly wrong with any possible new technological invention (the “unintended consequences of modernity”). Thus, as society has become more reflexive about its own technological achievements, the awareness of risk has become a more powerful driver of social forces than it was previously when risks were more perceived as due to intangible forces of nature rather than consequences of our own society, and therefore modern Western society’s response to risk has become qualitatively different to what it was before.

*Understanding uncertainty:* How people understand uncertainty is a related but somewhat orthogonal issue to the above—this may relate, for example, to the literature of social representations of risk (Joffe 1999; Washer 2004), which uses the social psychological literature of social representations (Moscovici 2000) to characterize how risk issues are perceived and made sense of through associated reasoning—new abstract and intangible concepts as are usually found in topics surrounding risk are conceptually anchored to concepts that are already understood, and thus new concepts are better assimilated into a group’s already held worldviews. Washer, for example, describes through the analysis of newspaper reports of recent new infectious disease outbreaks like SARS or avian flu and how these unfamiliar diseases (and the risks they represent) are being commonly anchored to already understood and familiar diseases (erg. the Spanish influenza outbreak of 1918), or to other aspects, such as vaguely xenophobic expectations of lax health and hygiene practices of the countries of origin; these mechanisms thus place the new disease into different categories of risk than they might otherwise have been perceived if anchored differently.

Hogg (2007) similarly uses a social psychological perspective, social identity theory (Tajfel 1981) to describe issues of intergroup and in group trust, arguing that our social identities about which groups we belong to effect how we trust the risk statements of others—in-group members are trusted more than out-group members, and even within groups, individuals who are more prototypical in that their characteristics conform well to group norms and values, are trusted more than more marginal members.

Here, we can also list other approaches that are interested in the social construction of risk. Cultural theories of risk such as the influential approach of Mary Douglas (1992) and the risk society argument are relevant here as well, because it is concerned with how societies construct (and thus understand) risks.

*What exactly are we uncertain about?* I left this category until last because this will be my focus in the following section. The object of our uncertainty has been the concern of several classification systems described above when, for example,

Wynne talks about uncertainty over causal chains or networks. Similarly, Walker et al. (2003) talk about “locations” of uncertainty, defining that as “where uncertainty manifests itself within the whole model complex” (p. 9), and distinguishing between uncertainty about the context (uncertainties outside of the model), models, inputs, parameters, and the final model outcome. In the following section, I will propose our slightly similar scheme which aims more at a rather general classification of the main types of objects we can be uncertain about which translates somewhat into Walker et al. locations of uncertainty, but is aimed at dispensing a too fine-grained classification in favor of one that we feel makes intuitive sense and can help explain different groups’ reactions toward the same risk scenarios.

In slicing the risk literature into these different categories of what they find interesting about risk, I recognize that a lot of work on risk looks at interactions between these different categories: for example, we can be interested in how different representations of risk and different aspects of risk can affect different people or groups of people. But I hope that this way of presenting the risk literature helps make sense of these interactions, and can therefore provide an interesting look into how different aspects of research on risk interlock. Our specific distinction between different objects of risk is itself designed in part to explain different outlooks on risk. In the following sections, I will present the different objects of uncertainty and, following that, explain through a few examples of how objects of uncertainty interconnect with some of the other dimensions of risk in a way which will hopefully give us a fuller description of the different risk scenarios.

## Objects of Uncertainty

Our classification (Spiegelhalter et al. 2011), somewhat unwisely in retrospect, divides the objects of uncertainty into different “levels.” I am calling our decision unwise because this suggests a particular linear hierarchy which may be misleading, but also because other commentators have attached the label “level of uncertainty” to some of the other dimensions of uncertainty outlined above. Specifically, Walker et al. use the term “levels of uncertainty” to describe the spectrum from determinism to “total ignorance.”

We distinguish between three types of uncertainty within the modeling process, and two without. Our use of the term model here is meant to be rather generic. Philosophical and social studies of scientists have shown that the term “model” can be used in varying ways in science (Bailer-Jones 2003), and, thus, generally it does not have the precise definition that it would have in mathematics or statistics. For example, the everyday constructions through which we as laypeople make sense of risk situations is taken here to be a kind of model as well, since we take our own incomplete information of the world and how we understand things to work and thus gain an understanding of what might happen. The difference between the nonexpert modeling we do in our everyday life and the expert risk

assessments is at the end merely a matter of background technical knowledge and competence and levels of commitment, rather than a huge qualitative difference. There is of course much more to be said about lay understanding and construction of risk perspectives, but, for my purposes, it should be enough to use the term “modeling” in an inclusive way that encompasses both expert and lay processing of risk.

By using the term model in this broad sense, I can apply it to different and varying real-life uncertainties and can include the formal mathematical application of the term as well as the more vague, everyday usage of model, in order to achieve applicability of our scheme across a wide variety of real-life risk situations, where precise mathematical or statistical modeling is impossible, impractical, or simply overlooked.

The categories I will present here are not meant to be mutually inclusive, and they will overlap. On the contrary, as I will argue with a couple of examples, in most risk situations, various levels of uncertainty are present at the same time, and our differences of opinion about risks may be due to us giving different importance to different levels.

*Level 1.* Uncertainty about the outcome. The model is known, the parameters are known, and it predicts a certain outcome with a probability  $p$ . An example here is the throw of a pair of dice: Our model is in this case the fundamental laws of classical probability, the parameters are the assumption that the dice are fair and unloaded, and the predicted outcome of, say, two sixes is  $(1/6)(1/6) = 1/36$ .

This is comparable to the “final model outcome” in Walker et al. On its own, this level of uncertainty exists only in rather idealized situations, as in arguably the example above of the dice. However, this is the level at which we as members of the public are most likely to encounter risk, for example, when we read in a newspaper that “the chances of developing bowel cancer is heightened by 20 % if we eat a bacon sandwich every day” (which is a real example taken from the case study further elaborated below). Such clear numbers, in the vast majority of cases, hide the fact that there are additional uncertainties related to the process in which experts arrived at it.

*Level 2.* Uncertainty about the parameters: The model is known, but its parameters are not known (Once the parameters are fixed, then the model predicts an outcome with probability  $p$ ).

This may simply be a lack of empirical information: If only we knew more, we could fix the parameters.

Our concept of uncertainty about the parameters itself hides a variety of different ways in which we can be uncertain about them: We can have fairly good, quantified probabilities about what the parameters should be as they might simply be a matter of getting better information about the system that is being modeled, but, more problematically, we could also be uncertain about how better measurements themselves are achieved, and/or our uncertainty about the parameters can itself only be expressible as a probability distribution, or even only a qualitative list of possibilities, or lastly we might simply have no idea of what the possibilities could be in the first place. Thus, here some of the different dimensions

as outlined above intersect with the object of uncertainty: our uncertainty about the parameters can be due to epistemic or aleatoric sources, and it can be represented in different ways.

Unlike Walker et al. we make no distinction between parameter and input uncertainty here, firstly for reasons of simplicity, but also because this more fine-grained distinction is not all that useful when we try to apply our scheme to real-life examples. Similarly, we would also class uncertainties over boundary conditions and initial values into this category as well, all of which may strain the term parameter uncertainty into categories not strictly speaking considered parameters as such—at the end, however, we decided to balance usefulness and simplicity with detail.

*Level 3. Uncertainty about the model:* There are several models to choose from, and we have an idea of how likely each competing model is to reflect reality. Models are usually simplifications about how the world works, and there are often several ways of modeling any given situation.

This is analogous with Walker’s model uncertainty, and again, this uncertainty itself can be presented in different ways and may be due to different sources. The way we should represent the uncertainties over model choice is more of a contentious issue and, of course, depends on the precise source of that uncertainty itself. In Spiegelhalter et al. (2011), we advocate a Bayesian approach to compare competing models (after Hoeting et al. 1999). This though will not be everyone’s favored approach, which means that, in most situations, we will encounter varying approaches to representations of model uncertainty.

Here, there can be, and frequently are, disagreements between the experts themselves, which means that to the nonexpert public or other consumers of a risk assessment, the uncertainty over the model choice is often related to other factors, such as how much trust they place in the experts to evaluate their model choices honestly or competently, and involves furthermore making a judgment between different experts’ assessment when faced with disagreement—however competence and honesty are assumptions that are made only implicitly (only rarely will experts be honest enough to consider their own competence as part of the overall risk assessment—building in an estimation of your own honesty into a risk assessment poses even more problems) and not strictly speaking part of the modeling process. There is therefore is a qualitatively different uncertainty for consumers and for producers of risk assessment, which will be the next level.

*Level 4. Uncertainty about acknowledged inadequacies and our implicitly made assumptions.* Every model is only a model of the real world and never completely represents the real world as such. There are therefore inevitable limitations to even the best models. These limitations could arise because some aspects that we know of have been omitted, or because of extrapolations from data or limitations in the computations, or a host of other possible reasons. Similar in a way to Wynne’s concept of “indeterminacy,” this is about questioning the assumptions we make, for example, about the validity of the science itself, and thus goes slightly perpendicular to the problems of choosing the models and parameters. These include the “imaginable surprises” (Schneider et al. 1998), that is, things we

suspect could occur but about which we do not know enough to be able to include them in the model.

As outlined above, this is where the question of trust comes into force as these are factors that are implicitly not assumed to matter in the risk assessments but not (or rarely) part of it. Similarly, there are always many assumptions about the world that have to be made and that are not part of the modeling process because they are assumed for one reason or another. For example, the risk referred to above of eating too many bacon sandwiches relies not only on the empirical and theoretical studies performed in the analysis, but also on the accumulated medical knowledge about cancer that was taken as given within the risk analysis. Any error within the fundamental scientific background assumed in a model that is supposed to reflect the real world albeit simplified will make that model less reliable. Therefore, uncertainties in our assumptions and scientific background knowledge are also inadequacies that are acknowledged but not usually part of the modeling process itself. At the same time, these are inadequacies in the process that are at least acknowledged in some way even if not particularly acted upon.

Dealing with acknowledged inadequacies can be done through informal, qualitatively formulated acknowledgment or listing the factors that have been left out of the model, or of course simple denial that there are any in the first place.

*Level 5. Uncertainty about unknown inadequacies:* We do not even know what we don't know. This particular type of uncertainty was made notorious through Donald Rumsfeld's famous speech on "unknown unknowns":

There are known knowns. These are things we know that we know. There are known unknowns. That is to say, there are things that we now know we don't know. But there are also unknown unknowns. These are things we do not know we don't know (Rumsfeld 2002).

There are as yet not very many formal approaches to unknown unknowns literature, though the concept has been well known for a while—for example, Keynes wrote that about some uncertainties, "there is no scientific basis on which to form any calculable probability whatever. We simply do not know" (Keynes 1937). Long before Keynes and Rumsfeld, however, a concept similar to unknown unknowns was introduced by Plato through the famous "Meno's paradox": How can we get to know about something when we are ignorant of what it is in the first place? (Sorensen 2009). It is also related to Taleb's concept of "black swan events" in economics (Taleb 2007) which are events that were not even considered but which, due to their high impact, have a tendency to completely change the playing field, and which is one of the concepts he used to warn about (what turned out to be) the 2008 world financial crisis.

These inadequacies are difficult to deal with formally or informally because we don't really know what they may be, and we are constrained in a way by the limits of our imagination of what could possibly go wrong—Jasanoff (2003), for example, identifies lack of imagination as one of the factors limiting our knowledge for proper risk assessments in postnormal science (p. 234).

Responding to unknown unknowns is naturally very difficult because by definition we do not know what they are. We can however acknowledge them

through simple humility that it is always possible that we are mistaken, as demonstrated by Cromwell's quote in the epigraph. Another way is to brainstorm every possibility we can think of and letting our imaginations go wild. This approach is of course never going to be able to cover everything that could go wrong and will therefore not eliminate unknown unknowns.

A slightly more formal way of responding to unforeseen events is the introduction of "fudge factors," for example, in bridge or airplane design, where we design the structure to be a bit stronger than even the worst case scenarios that we could think of require—though even then there is always the conceivable possibility that something worse may happen.

These levels in a way relate to different concerns of different disciplines—who are after all interested in different aspects of risk. For example, the traditional mathematical and philosophical problems of probability theory are mostly concerned with level 1 uncertainty. Statisticians are mostly concerned about level 2 and 3 uncertainty, that is, finding the right model and, within that model, adjusting the parameters appropriately. It seems unfortunately that uncertainties on which we cannot have a particular mathematical handle on are so often ignored by statisticians and risk modelers—often probably for the pragmatic reason that there simply is not much they can say about the higher levels with the mathematical tools of their trade. Shackley and Wynne (1996), for example, write that in their study of policy discourse on climate change, policy makers were concerned about the validity of the models, while the scientists themselves never even considered that to be an issue, but were instead more concerned about measurement errors within their models. This is to an extent an unfair generalization. An informal survey of technical abstracts from a recent Carbon Capture and Storage conference (Riesch and Reiner 2010) has shown that while model uncertainty is not generally discussed, it does occasionally get mentioned, alongside even an occasional awareness that there are uncertainties associated with unmodeled or unmodelable inadequacies. Nevertheless, worries about model inadequacies were certainly not a prevalent concern among the scientists and risk modelers.

This expert discourse unfortunately distorts the way we perceive particular risks because higher level uncertainties still exist. This may lead to situations like the ones described by Taleb (2007) when he writes about economists having forgotten that unforeseen out-of-the-blue events can occasionally happen and completely mess up our predictions—the sort of events he calls "black swans" if they also have a high potential impact.

The risk society approach of Beck and Giddens is talking mostly about levels 4 and 5, where it is hypothesized that late modern society is living with the increased realization that unmodeled and unmodelable risks are pervasive, and that even if we had some kind of handle on them, there is always the possibility of completely unforeseen events, what Beck calls the "unintended consequences" that he mostly associates with new technology, but which need not necessarily be tied in with it. In Beck's characterization of late modern society, we have now become accustomed to the realization that despite the best risk modeling of science and engineering experts, technological innovations and advances always have

unforeseen consequences, completely left-field occurrences that the original evaluations failed to take into account—in other words, we now know that we live with level 4 and 5 uncertainties all around us. In a way, it matters less to the sociological literatures whether these risks are real or not, but the mere realization that they do happen affects the way late modern society evaluates technological progress and ultimately, itself. Beck's work has been criticized for ostensibly being about risk, but not quite understanding the concepts of risk analysis and probability (Campbell and Currie 2006), though this slightly misses the point because, within this scheme, it is not really the nature of risk that is important, but responses to it.

As I have tried to argue above, the different disciplinary approaches to risk intersect in different ways—not only do they find different objects of uncertainty important, but they are also interested in different topics among the other dimensions. However, I have not yet found a comprehensive way of translating between the different approaches, and my categorization between different dimensions of risk is meant to solve this. In particular, I feel that the objects of uncertainty dimension which I presented here in more detail can be an important perspective with which to analyze different risk situations in a way that makes sense to the different disciplinary approaches. In the following section, I will go through several examples to illustrate what this perspective can show how all levels of uncertainty are present in most situations involving risks or uncertainties. In particular, I am interested (as a sociologist) to explain how different groups' perceptions of essentially the same scenario can differ so dramatically: because through their background experience, assumptions, and worldviews, they will attach different importance to the different levels described above. One important departing point therefore is my assertion that all the levels are present in every risk situation, and that the relative importance that is attached to them depends on who is mulling over it, and I will argue for that below. This seems to be more important on the objects of uncertainty dimensions more than on some of the others, and therefore I feel concentrating on these will help us bring about a more comprehensive way of translating between the various risk literatures, as these will be interested in different objects of uncertainty within each situation.

## *Examples*

In this section, I will explore how these five levels of uncertainty can help explain what happens in various real-life cases in which risk, perceived or real, is a factor, and how our concept of the levels explain different perceptions and how this can lead to the communication difficulties between groups with different perspectives (say between to proponents and opponents of carbon capture in the third example). In each of these cases, all five levels of uncertainty are present, though they are differently important and relevant depending on the example.



## The Lottery

I will start with a situation which traditionally is seen as less problematic because it seems to rely only on outcome uncertainty. In a typical national lottery, such as in the UK, there are 49 balls, and each week, six balls are drawn; people who have chosen all six correctly win the jackpot. While the exact rules of how much you win are more complicated (depending on the lottery), the case is at least on the surface clearly of level 1: The model is known, the parameters are known, there are no known inadequacies in the model; all the uncertainty that remains is the probability predicted by the model.

This however does not mean that there are no uncertainties present of the other levels, they are simply more hidden and seem less relevant. Level 2 uncertainty concerns the uncertainty of the parameters. In this case, one of the parameters that we have assumed were fixed concerned the individual probability for each ball to come up, thus the question essentially revolves around whether the lottery machine is fair. This is of course a question that we should be asking ourselves when we play the lottery, though we rarely do because we trust the authorities that set up the game. As soon as that trust is lost however, level 2 uncertainty comes to the foreground in our evaluations of how likely a jackpot win is. But this is also an empirical question—for the regulator to make sure that the parameters are what we assume them to be, the equipment is regularly checked, and therefore even if we trust the operators to run a fair game, there is still residual empirical uncertainty over the measurements performed during the equipment checks.

Level 3 and 4 uncertainties, in this case, are less likely to bother us because the situation is relatively simple. We, thus, do not really have competing models with which to describe the lottery: unlike in the examples below, where we have situations for which we need a model to describe it, in this case, we start with the model and set up the reality to fit it—that is, after all how the game was constructed. Therefore, in this case, we have a lot of confidence that the mathematical model we use to describe the game is accurate and not likely to be replaced by one that reflects the situation better.

This however does not necessarily reflect the situation from the point of view of the consumer of the lottery—given that the rules of the game are published but the precise probabilities for a given type of win are not necessarily, we have to make our own calculations, and, for the mathematically less able among us (such as myself), there remains the very real possibility that I have made a mistake in estimating my chances of winning. Again the lottery is a pretty simple situation where even I will not have many difficulties; however, the same cannot be said about other games of chance such as blackjack where there is no real model uncertainty from the point of view of an able calculator, but where for the average player, the probabilities are very much subject to uncertainties over mathematical ability. The trust that we have in the operator mentioned above to demonstrate why our parameter assumptions may be wrong is itself a frequently *unacknowledged* inadequacy: The probability that the operator is cheating is, even if somehow quantifiable, rarely part of the model (which in turn makes out the parameter

uncertainty to be only dependent on empirical questions) used to estimate the probabilities of winning or the expected pay-out. Yet again, for the consumer who may have a different estimation of the trustworthiness of the operator, level 4 introduces an unmodeled uncertainty, and their estimations of this uncertainty will be different according to background knowledge and assumptions.

Considering completely unexpected scenarios now, maybe the machine could blow up during the draw, invoking maybe the need to refund punters—again this would affect the probability of winning overall in a slight way. Or the operator could be declared bankrupt, in which case, it is not clear whether there would be refunds at all, and the issue would probably only be solved on a case by case basis depending on the whims or political pressure of put upon the government as is the case when other companies fold (even though costumers will usually not get refunds if a company goes bankrupt, there are often cases, such as tour operators, where political intervention may make an exception). The possibilities here are of course only restrained by my imagination and as soon as I formulate them they are not strictly speaking unknown unknowns. However, the relative ease with which we can conjure up scenarios which are not foreseen at all points toward a large background level 5 uncertainty which cannot be eliminated completely or even adequately estimated through better modeling.

These considerations I hope demonstrate that even in seemingly very clear situations that are not usually assumed to be subject to other than level 1 uncertainties, our estimation of the uncertainties rely to a large extent on our trust in the operator, our background assumptions and mathematical abilities, and these differ from person to person.

## **Saving Our Bacon**

What exactly does it mean when we are informed that we are facing an increased risk (by 20 %) of bowel cancer if we eat more than 500 g of processed meat a day (WCRF 2007a)? Again, I will hope to demonstrate here that in this claim, there are several levels of uncertainty interwoven because, depending on which perspective we take, we can evaluate the uncertainties of different objects in varying ways. Therefore, making sense of that claim will involve untangling them. (Incidentally, the lifetime risk of bowel cancer is estimated in the report as 5 %, which raises to 6 % when we eat a lot of red meat a day. In relative terms, the increase of risk is 20 %, while in absolute terms it is 1 %. The fact that the WCRF chose to present the more scary relative increase in their press strategy, rather than the more informative absolute increase, tells us a lot about their communication priorities, see also Riesch and Spiegelhalter 2011).

The claim above is based on a meta-analysis performed by the World Cancer Research Foundation of various published trials that investigated the incidence of bowel cancer among people who consume a lot of red meat versus those that do not (WCRF 2007a). One level of uncertainty therefore involves what the studies, as aggregated by the accepted rules of how researchers should do meta-analyses,

tell us about eating processed meat: the model is known (in this case, the rules involved of doing the analysis, as well as the rules of the individual studies aggregated in the meta-analysis), the parameters are fixed (in this case the empirical evidence), and together they predict the outcome, bowel cancer, as 20 % higher than without the consumption of processed meat. This is the level of uncertainty at which the WCRF communication strategy operated: Our science has found that the risk is  $p$ , and that is what the public should know about red meat (as suggested by the WCRF press strategy; WCRF 2007b; 2007c).

However, especially when looked at from the perspective of the reader of the report, the other levels of uncertainty are there in the background as well and have been emphasized by some of the other actors in the debate: Level 2 uncertainty is, to a certain extent found in the report itself, as this represents the empirical uncertainties surrounding each individual study in the meta-analysis (i.e., fixing the parameters through empirical data): These empirical errors have been aggregated, and since the meta-analysis involved lots of different studies, the overall error has been reduced, and this level of uncertainty is represented through the use of error bars in their charts. Error bars of course did not make it into the verbal communication that accompanies the study's conclusion; instead, the information about uncertainty here is formulated qualitatively: The report distinguishes between the evidence being "convincing," "probable," and "limited." In the final communication of the report, the inherent experimental error, the level 2 uncertainty, was not quantitatively included and could possibly be said to be relatively low. There were though a small qualitative indicators in the wording of the press release:

*There is strong evidence* that red and processed meats are causes of bowel cancer, and that there is no amount of processed meat that can be confidently shown not to increase risk (WCRF 2007c, my emphasis).

While level 2 uncertainty has been addressed in the report, if only qualitatively as a way of showing some caution in interpreting the results, level 3 uncertainty posed more problems and was the sort of uncertainty that the expert critics of the report have focused on: This is uncertainty surrounding choosing the model itself. In this case, that translates to the controversy of how the meta-analysis was done, and specifically which studies were included in the analysis. Critics of the report have pointed out that the meta-analysis has left out many individual studies that, if included, would have given the whole analysis a different result. Whether or not there was much merit in these criticisms, they at least demonstrated that no amount of certainty in the analysis itself can remove the uncertainty inherent in choosing the model. The methodology of meta-analysis in general, while an established tool within medical research, is nonetheless not without its critics, and again, though I will not comment on whether these criticisms have much merit, they demonstrate that even within the expert community there are differences of opinion, and therefore, especially for nonexpert bystanders like me, there is an additional uncertainty over whether the whole methodology used by the report is sound in the first place.

This is compounded by level 4 uncertainty because even if we did have certainty over which studies should have been included, and whether meta-analysis in general is the best way to pool the results from these studies, there is still residual uncertainty about the scientific background assumptions underlying the study, which relies to a large part on previous medical knowledge on for example cancer which is seen as well-established and therefore not considered a factor to be included in the model at all. This is to an extent not too much of an inadequacy because the study being an empirical evaluation of several selected trials and observational studies does not rely much on previous medical knowledge; however, it does rely on previously established medical and scientific knowledge and assumptions that these methodologies are a valid way of establishing knowledge. Again though it is not my place to comment on whether there is any merit to these criticisms, that is a criticism that certainly has been made, not so much by medical experts, but by alternative health practitioners who reject a large amount of otherwise established medical knowledge and methods. For the nonexpert bystander, again, the situation is that of competing groups who both claim to have expert status and who have different opinions about what the study shows and can even in principal be expected to show. This is then a different level of uncertainty altogether for the consumer of the report.

Added to that, there are other implicitly made assumptions in the report which relate to the honesty and competence of the researchers themselves. Of course, we can't expect them to take these concerns seriously as an additional uncertainty in their own science, but these are not assumptions that can automatically be assumed by the reader. Both of these two levels of uncertainty (3 and 4) were emphatically not voiced in the official communication by the WCRF, which is understandable because they would have cast doubt on their own experts' judgment. However, they were certainly voiced by the critics: Level 3 uncertainty, as shown above, was the expert critic's response, of fellow medical researchers who accept the general methodology but object to the way it was performed in this instance, while level 4 uncertainty is more usually the response of critics who disagree with meta-analyses generally or who distrust or disbelieve some of the assumptions that medical research takes as established (these are not very influential among medical researchers, but have some influence among alternative medicine campaigners).

Finally, there is level 5 uncertainty: We may be completely wrong footed about the risks of processed meat—maybe the results of all the studies were a systematic error in the design of contemporary medical studies that we do not know about? Maybe something even more exotic has gone wrong? Admittedly, in this particular case, it is quite hard to imagine possible level 5 uncertainties, but this is of course the nature of this level of uncertainty as I have defined it by us not having any handle on it, and not even ever having thought of the possibility. Accordingly, giving a numerical estimation of this level of uncertainty is impossible. Beyond gut instincts, we cannot even tell if it is likely or not likely that something is fundamentally wrong with our conceptions of the problem. In the next examples, I will show that while in this case level 5 uncertainty is not much incorporated in the

current thinking about the subject, in many other situations involving risk, level 5 uncertainty can be central.

The complications and disputes involved in this case are connected with the protagonists talking about different levels of uncertainty: The WCRF experts talked about level 1 uncertainties in their take-home message to the general public, while at least acknowledging level 2 uncertainties when talking among themselves and in communication with other experts. Level 3 uncertainty is the level at which the expert critics attack the report, while the less listened to nonmedical critics attacked it at level 4. Meanwhile, level 5 uncertainty looms menacingly in the background. Some of the media interpretation and discourse about the WCRF study and its press releases can be found in Riesch and Spiegelhalter (2011).

## Carbon Capture and Storage

Carbon Capture and Storage (also referred to as Carbon Capture and Sequestration), or CCS, is a technology designed to reduce carbon emissions from fossil fuel burning power plants by capturing the CO<sub>2</sub> through various processes and storing it underground in depleted natural gas reservoirs or other suitable storage sites. The technology is seen by its proponents as an important and technically feasible way to lower carbon emissions because it relies on already fairly well-known mechanisms. While it is admittedly only planned as a relatively short-term solution to be deployed while renewable energy sources are being developed further, it solves some of the problems of “technology lock-in” that could happen if we concentrated only on a few favored energy sources such as solar- and wind-power which we have no guarantee yet that they will be deployable at a large enough scale to reduce carbon emissions in time to avert catastrophic climate change. Therefore, it is seen as part of a necessary portfolio of energy technologies that needs to be included if we want to avoid putting all our eggs in one basket (as argued for example by the influential Stern report, Stern 2007). Further benefits of the technology include more security in the energy supply because it would make burning coal an environmentally sound energy option again and, therefore, reduce the dependency some countries with large coal resources like the UK have on foreign gas imports. A more environmentally appealing further benefit of CCS is that when the technology is developed far enough, it can be used in conjunction with biomass burning power plants and therefore represents one of the few currently technically feasible ways of removing carbon from the atmosphere.

Despite these advantages, CCS has many opponents principally among the environmental community, who argue that it merely propagates our dependency on fossil fuels and drives funds away from developing more promising energy technologies which need to be developed anyway because even proponents of CCS see it only as a short term solution (Greenpeace 2008). Finally, one objection to CCS which threatens to be a show-stopper is the safety risks to local people and the local environment that are posed by possible CO<sub>2</sub> leakage from the storage reservoirs and the pipelines that transport the CO<sub>2</sub> from the power plants to these

sites. It is these safety risks of CCS that I will concentrate on here; however, the other arguments for and against CCS are relevant here because it is our background worldviews, knowledge, and assumptions which color the way we perceive specific risks. One immediately obvious example of how our background knowledge may color our perception of the risks of the technology concerns our knowledge (and uncertainties within that knowledge) of the toxicity of CO<sub>2</sub>. While CO<sub>2</sub> is not, in fact, neither toxic nor flammable (it does, however, act as an asphyxiant, and therefore still represents a potential though somewhat lessened danger to people living near leakage sites), public opinion surveys on perceptions of CCS have shown that worries over CO<sub>2</sub> are very much in the forefront of public safety concerns (Itaoka et al. 2004; Mander et al. 2010).

Level 1 uncertainty in this case is the final number of the risk assessment, which is usually the basis on which politicians or energy companies would claim that experts find the technology to be very low risk. These numbers are arrived at through models which make of course several assumptions. A general model of how carbon storage works depends very much on local conditions if we want to arrive at numbers for any particular reservoir and the surrounding area. The local conditions vary in great detail, and therefore experts who perform risk assessments of prospective sites need to investigate them very closely—what are the exact geological formations the CO<sub>2</sub> would be stored in, what are the properties of the cap-rock formations that are needed to keep the CO<sub>2</sub> from traveling up, are there seismic fault lines and if so how would they affect the storage, how many man-made injection wells are there, and how exactly are they going to be sealed once the CO<sub>2</sub> is injected, what is the general three-dimensional shape of the landscape above the possible leakage sites (since CO<sub>2</sub> is heavier than air, there is a chance that it might stay if it leaks into a valley and thus cause greater potential health risks). All these are in a sense parameters that need to be put into the general models if we want to arrive at a final number of expected deaths per year. All of these are subject to their own uncertainties, either because of potential measurement error, or even a more general lack of understanding of the local conditions which in practical terms can only be estimated.

At level 3, there is the choice of general model. In the case of CCS, there is still some argument over whether models developed by the gas and oil industries are really applicable to the storage of CO<sub>2</sub> (Raza 2009). There are also potential debates to be had over precisely what statistical methods should be used and their applicability. Writing in a Dutch popular science magazine article about the proposed (now canceled) CCS storage site under the town of Barendrecht near Rotterdam, Arnoud Jaspers felt that there is some additional uncertainty over model choice when he interviewed modeling experts, for example, some of the models simply did not take into account the three-dimensional structure of Barendrecht and therefore arrived at unreliable scenarios of what would happen should CO<sub>2</sub> leak (Jaspers 2009, 2010). As this shows, not every relevant bit of information makes it into all the models, and there is therefore some uncertainty about which model would be the best to use.

Furthermore, there are other things that are not considered in any model because we simply do not know enough about them or their relevance to be able to model them, these then are the acknowledged inadequacies we term level 4 uncertainties. A recent draft guidance document by the European Union on the implementation of Directive 2009/31/EC concerning geological storage of CO<sub>2</sub> (EU 2010), for example, divides the types of risk expected from reservoir leakage to be “Geological leakage pathways,” “leakage pathways associated with man-made systems and features (i.e., wells and mining activities),” and “other risks such as the mobilization of other gases and fluids by CO<sub>2</sub>)” (p. 31). While the first two are routinely part of the models, the “other” category provides more of a problem because, other than listing some possible scenarios that can only to be considered on a “case by case” basis, there is not that much additional analysis that can be introduced. Therefore (as a brief glance through the technical papers on CO<sub>2</sub> storage at the GHGT10 conference has shown—discussed in more detail in Riesch and Reiner 2010), most risk models of CO<sub>2</sub> storage consider leakage pathways along geological fractures or man-made boreholes but do not as such feature other possible leakage pathways either because they are judged to be not very important or, more worryingly, not enough is known about them to include them in the models.

Lastly, there are level 5 unknown unknowns which are not part of the modeling process, not because we do not know enough about them, but because we do not know about them at all. Giving a concrete example is, again, impossible since simply by thinking about them they become known unknowns. However, there are scenarios that can be imagined by the public that are never even considered in the expert literature. For example, the cover illustration to Jaspers’ (2010) popular science article on CCS features a huge “blow-out” scenario with a vast amount of CO<sub>2</sub> escaping explosively and destroying large parts of Barendrecht’s neighborhood with rescue helicopters hovering around the scene like tiny flies in relation to the explosion; painting the picture of carbon storage as a huge shaken soda bottle bubbling menacingly underneath the town which will explode spectacularly as soon as there is any kind of leak. This scenario was emphatically not considered in any risk assessment partly because such an explosion would be contrary to anything we know about the behavior of stored CO<sub>2</sub>, and it is probably fair to say that therefore it is not included in risk assessments nor even considered as a known inadequacy of the models. Nevertheless, this scenario of something going horribly wrong somehow is a valid concern especially for those who do not possess the expert background knowledge to adequately judge it as so unlikely as to not even be worth including in the model.

This rough overview on the risk debates on CO<sub>2</sub> storage is meant to show that there are very much different perspectives we can take on the risk of CCS, and which ones we put most stock in depend on our background knowledge and ideologies. As in the red meat and cancer example above, different actors in the debate have emphasized different types of uncertainties in this case. Energy companies like Shell or BP who are developing CCS projects, as well as those politicians who are keen on promoting it take comfort from the fact that final risk

assessments put the safety risks of the technology as very low, and in much of the industry communication literature on CCS, it is these figures that get mentioned rather than any more technical discussion surrounding how they were acquired. Literature from environmental groups on the other hand see the uncertainties in a different context by highlighting the potential of measurement errors in local evaluations, casting doubt on the modeling processes involved (for example, Greenpeace 2008).

## Climate Change

This leads us to a final very brief example because the one thing that complicates debates about CCS is its relation to the mitigation of climate change. Man-made climate change presents a particular problem, not because there are by now any doubts left that it is happening, but because the forecasting of how bad it will be under various scenarios is a very imprecise business. Scientists who try to predict possible climate futures almost always start by admitting that they are working from one particular model, and that there are several that we know we could use instead that give different estimates, and we rarely have any good handle of working out how likely each model is to reflect reality. In fact, it is often acknowledged that we know so little and climate and weather patterns are so complex that there are always possible factors that we do not even know that we do not know about which can take us by complete surprise.

Uncertainties in climate change modeling are thus dominated by levels 4 and even 5: We have several models to choose from, but not much knowledge on how well they are doing their jobs, and even then we are well aware that our forecasting is hostage to completely unforeseen things as well. Social science research on climate modelers themselves has shown that there is a wide range of expert opinions on the best modeling process and that, moreover, experts themselves will have an unreliable estimation as to the possible shortcomings of their own models (Lahsen 2005). Even then, if they have an adequate estimation of the reliability of their models, experts find it hard to communicate them, especially when these estimations cannot be quantified (Hillerbrand 2009). Therefore, yet again, for the nonexpert observer of these debates, the most important uncertainty is not over which model is best, but over which expert to trust most, and because there is so much disagreement and unreliability of the experts' own assessments of their models, this uncertainty weighs more for the nonexpert than for the experts.

## Further Research

By conceptualizing uncertainty along the various dimensions introduced above we can gain an appreciation of what the different disciplines find interesting about risk and, hopefully, find how they interconnect. The risk society literature, for example,



as I outlined above is interested in different aspects of risk than the risk management literature. Considering the particular dimension that I think is most important in my map, the objects of uncertainty, gives us an idea of how and why different people estimate uncertainties differently even when presented with the same information and furthermore shows how different disciplines can themselves study risk from different perspectives.

This is therefore more useful for sociological research than the otherwise admirable combined system of van Asselt and Rotmans (2002) and Walker et al. (2003), whose scheme was designed primarily for use by experts in integrated assessments. Unlike Wynne's and Stirling's systems, which were designed for sociologists to understand and analyze different reactions to uncertainty, my map and classification tries to be more inclusive and meticulous in teasing apart different aspects of uncertainty, while hopefully still being simple enough to be useful for gaining an immediate and intuitive understanding of why and how opinions on risk so often differ. This can then become a useful tool when social scientists communicate the problems of the social contexts of risk to technical experts—this is after all a role performed very often by social scientists who have been funded by scientific research institutions and funding agencies “to look at the social side of things,” but who often struggle to make the social insights relevant and intuitive to the technical experts they work with. This is a main reason why we (Riesch and Reiner 2010; Upham et al. 2011) use this framework in the study of risk opinions on energy infrastructure (CCS and biofuels, respectively), to so far very positive reactions from the technical communities. This approach therefore tries to marry the sociological usefulness of Wynne with the technical relevance of Walker and van Asselt and their colleagues.

Though the scheme presented here is meant to be illustrative rather than prescriptive, it shows some clear lessons for risk communication strategies. Since, as I have argued here, different people are worried about different aspects of the risks and, in particular, attach the uncertainty to different objects, risk communication strategies often fail to convey the information that people actually find important. While there is no silver bullet with which to persuade people who simply do not trust the experts, or who's understanding of technological risks gives a higher importance to unforeseen events, taking these different perspectives into account will ensure that the conversation at least does not disintegrate into different actors failing to understand each other. In designing a communication tool about the risks of CCS, for example, we may want to pay particular attention not just to the risks as calculated by the risk assessment, but also how it was arrived at, what the uncertainties with the parameters are, what was the choice of models available, and why was this particular one chosen, what possible inadequacies were not modeled and finally what are the plans for action should unforeseen consequences occur.

Future research will hopefully develop some of the other dimensions along the more detailed level as I have tried with the objects of uncertainty dimension (and as van Asselt and Rotmans have already done with the sources of uncertainty). This would then allow us to construct a more detailed table through which we can

map, at a glance, where the different academic literatures on risk lie and intersect and which may help researchers in finding connections and future ideas for more integrated interdisciplinary research on risk.

Work is also underway to develop case studies which apply the objects classification to different risk scenarios. I have summarized here the application to CCS (Riesch and Reiner 2010); furthermore, we are applying it to the problem of indirect land use change for the biomass energy industry (Upham et al. 2011). Furthermore, detailed studies on more diverse risk situations will hopefully be able to tease out more of the potential but also limitations of our scheme.

## References

- Aven T, Renn O (2009) On risk defined as an event where the outcome is uncertain. *J Risk Res* 12(11):1–11
- Bailer-Jones DM (2003) Scientists' thoughts on scientific models. *Perspect Sci* 10:275–301
- Beck U (1992) *Risk society: towards a new modernity*. Sage, London
- Campbell S, Currie G (2006) Against beck: in defence of risk analysis. *Philos Soc Sci* 36(2): 149–172
- Carlyle T (1871) *Oliver cromwell's letters and speeches: with elucidations*. Scribner, Welford, New York. [http://www.gasl.org/refbib/Carlyle\\_\\_Cromwell.pdf](http://www.gasl.org/refbib/Carlyle__Cromwell.pdf). Accessed 1 Sep 2010
- Douglas M (1992) *Risk and blame*. Routledge, London
- Edwards A, Elwy G, Covey J, Matthews E, Pill R (2001) Presenting risk information—a review of the effects of “framing” and other manipulations on patient outcomes. *J Health Commun* 6:61–82
- EuropeanUnion (2010) Implementation of directive 2009/31/EC on the geological storage of carbon dioxide. <http://ec.europa.eu/clima/policies/lowcarbon/docs/GD1-CO2%20storage%20life%20cycle%20risk%20management-consultation.pdf>. Accessed 31Dec 2010
- Funtowicz SO, Ravetz JR (1990) *Uncertainty and quality in science for public policy*. Kluwer, Dordrecht
- Funtowicz SO, Ravetz JR (1993) Science for the postnormal age. *Futures* 25:739–755
- Giddens A (1999) Risk and responsibility. *Mod Law Rev* 62:1–10
- Gigerenzer G (2002) *Reckoning with risk*. Penguin, London
- Gillies D (2000) *Philosophical theories of probability*. Routledge, London
- Greenpeace (2008) *False hope: why carbon capture and storage won't save the climate*. Greenpeace International, Amsterdam
- Hacking I (1975) *The emergence of probability*. Cambridge University Press, Cambridge
- Hillerbrand R (2009) Epistemic uncertainties in climate predictions. A challenge for practical decision making. *Intergenerational Just Rev* 9(3):94–99
- Hoeting J, Madigan D, Raftery A, Volinsky CT (1999) Bayesian model averaging: a tutorial. *Stat Sci* 14:382–417
- Hogg MA (2007) Social identity and the group context of trust: managing risk and building trust through belonging. In: Gutscher H, Siegrist M, Earle TC (eds) *Trust in cooperative risk management*. Earthscan, London, pp 51–72
- Itaoka K, Saito A, Akai M (2004) Public acceptance of CO2 capture and storage technology : a survey of public opinion to explore influential factors. In: Rubin ES, Keith DW, Gilboy CF (eds) *Proceedings of 7th international conference on greenhouse gas control technologies*, vol 31, 1: Peer-reviewed papers and plenary presentations, IEA Greenhouse Gas Program, Cheltenham

- Jasanoff S (2003) Technologies of humility: citizen participation in governing science. *Minerva* 41:223–244
- Jaspers A (2009) Slapen met de ramen dicht. *Natuur-Wetenschap & Techniek (NWT)* 77(4): 24–33
- Jaspers A (2010) The view from technological journalism. FENCO workshop CCS and public engagement, Amsterdam
- Joffe H (1999) Risk and the other. Cambridge University Press, Cambridge
- Keynes JM (1937) The general theory. *Q J Econ* 51:209–233
- Knight FH (1971) Risk, uncertainty and profit (reprint of the 1921 edn). University of Chicago Press, Chicago
- Lahsen M (2005) Seductive simulations? Uncertainty distribution around climate models. *Soc Stud Sci* 35(6):895–922
- Mander S, Polson D, Roberts T, Curtis A (2010) Risk from CO<sub>2</sub> storage in saline aquifers: a comparison of lay and expert perceptions of risk. GHGT10, Amsterdam
- Moscovici S (2000) Social representations. Polity Press, Cambridge
- Norton J, Brown J, Mysiak J (2006) To what extent, and how, might uncertainty be defined? Comments engendered by Defining uncertainty: a conceptual basis for uncertainty management in model-based decision support: Walker et al., *Integrated Assessment* 4: 1, 2003. *Integrated Ass J* 6(1):83–88
- Ravetz JR (2006) Post-normal science and the complexity of transitions towards sustainability. *Ecol Complex* 3:275–284
- Raza Y (2009) Uncertainty analysis of capacity estimates and leakage potential for geologic storage of carbon dioxide in saline aquifers. Masters Thesis, MIT, Cambridge
- Renn O, Klinke A (2004) Systemic risks: a new challenge for risk management. *EMBO Rep* 5(S1):S41–S46
- Riesch H, Spiegelhalter DJ (2011) Careless pork costs lives: risk stories from science to press release to media. *Health Risk Soc* 13(1):47–64
- Riesch H, Reiner D (2010) Different levels of uncertainty in carbon capture and storage (submitted)
- Roeser S (2009) The relation between cognition and affect in moral judgments about risk. In: Asveld L, Roeser S (eds) *The ethics of technological risk*. Earthscan, London, pp 182–201
- Roeser S (2010) Intuitions, emotions and gut reactions in decisions about risks: towards a different interpretation of neuro ethics. *J Risk Res* 13(2):175–190
- Rumsfeld D (2002) Defense.gov news transcript. <http://www.defense.gov/transcripts/transcript.aspx?transcriptid=2636>. Accessed 3 Dec 2010
- Schneider SH, Turner BL, Garriga HM (1998) Imaginable surprise in global change science. *J Risk Res* 1(2):165–185
- Shackley S, Wynne B (1996) Representing uncertainty in global climate change science and policy: boundary ordering devices and authority. *Sci Technol HumVal* 21(3):275–302
- Slovic P (2000) The perception of risk. Earthscan, London
- Slovic P, Finucane M, Peters E, MacGregor D (2004) Risk as analysis and risk as feelings: some thoughts about affect, reason, risk and rationality. *Risk Anal* 24(2):311–322
- Sorensen R (2009) Epistemic paradoxes. In: Zalta E (ed) *The Stanford encyclopedia of philosophy*. <http://plato.stanford.edu/archives/spr2009/entries/epistemicparadoxes/>. Accessed 31 Dec 2010
- Spiegelhalter DJ (2010) Quantifying uncertainty. In: Paper presented at handling uncertainty in science, Royal Society, London, 22–23 Mar 2010
- Spiegelhalter DJ, Pearson M (2008) 2845 ways to spin the risk. <http://understandinguncertainty.org/node/233>. Accessed 31 Dec 2010
- Spiegelhalter DJ, Riesch H (2011) Don't know, can't know: embracing deeper uncertainties when analysing risks. *Philos T Roy Soc A* 396(1956):4730–4750. doi:10.1098/rsta.2011.0163
- Stern N (2007) *The economics of climate change*. Cambridge University Press, Cambridge
- Stirling A (2007) Risk, precaution and science: towards a more constructive policy debate. *EMBO Rep* 8(4):309–315

- Sunstein CR (2005) *Laws of fear*. Cambridge University Press, Cambridge
- Tajfel H (1981) *Human groups and social categories: studies in social psychology*. Cambridge University Press, Cambridge
- Taleb N (2007) *The black swan: the impact of the highly improbable*. Penguin, London
- Upham P, Riesch H, Tomei J, Thornley P (2011) The sustainability of woody biomass supply for UK bioenergy: a post-normal approach to environmental risk and uncertainty. *Environ Sci Policy* 14(5):510–518
- van Asselt MBA, Rotmans J (2002) Uncertainty in integrated assessment modelling: from positivism to pluralism. *Clim Change* 54:75–105
- Walker WE, Harremoes P, Rotmans J, van der Sluijs JP, van Asselt MBA, Janssen P, Kraye von Krauss MP (2003) Defining uncertainty: a conceptual basis for uncertainty management in model-based decision support. *Integrat Ass* 4(1):5–17
- Washer P (2004) Representations of SARS in the British newspapers. *Soc Sci Med* 59:2561–2571
- WCRF (2007a) Food, nutrition, physical activity, and the prevention of cancer: a global perspective. WCRF, Washington
- WCRF (2007b) Landmark report: excess body fat causes cancer. [http://www.wcrf-uk.org/press\\_media/releases/31102007.lasso](http://www.wcrf-uk.org/press_media/releases/31102007.lasso). Accessed 13 Aug 2008
- WCRF (2007c) Media quotes. [http://www.wcrf-uk.org/press\\_media/quotes.lasso](http://www.wcrf-uk.org/press_media/quotes.lasso). Accessed 13 Aug 2008
- Wynne B (1992) Uncertainty and environmental learning: reconceiving science and policy in the preventive paradigm. *Global Environ Chang* 2(2): 111–127