

## Taming chance: Risk and the quantification of uncertainty

JOHN KADVANY

*Applied Decision Analysis, Inc., 2710 Sand Hill Road, Menlo Park, CA 94025, U.S.A.*

**Abstract.** The quantification of uncertainty through probability is central to definitions of risk used in environmental policy analysis. This essay explores the translation of unquantified uncertainty into probability and the expression of allied philosophical problems in the practice of environmental risk and decision analysis. First we look at means used in science for handling uncertainty associated with some major risks and which are not well represented through probability. 'Saving the science' without quantified probability is addressed through the role of probabilistic events in risk analysis, suggesting the need to expand the scope of risk analysis to include some types of unquantified reasoning about adverse events. Next we look at uses of subjective probability and decision analysis to overcome problems of unquantified uncertainty in science, where we argue that a constructive conception of probability judgments, based in the foundations of decision analysis, provides the most useful approach for such methods. A theme throughout is the role of intellectual control implicit in our efforts 'to tame change' through the representation of uncertainty through probability.

### I. Defining risk

Though alternative definitions of risk are adopted by both individual analysts and institutions, a single definition sufficient for most applications is that of risk as *adverse consequences under uncertainty*. This definition, which I will call *relational risk*, succinctly identifies the dependent categories of *value*, *events* (including *actions*), and *knowledge* which jointly characterize particular instances or classes of risk as a complex relation between humans and nature. Depending on preferred methods for dealing with uncertainty, theories of human action, ideologies of value, and accounts of the social or natural worlds, specialized definitions can be used in place of relational risk. The best known is a special case of relational risk, namely *quantified risk*: adverse consequences *plus* a probability assignment over events.<sup>1</sup>

This is essentially the definition adopted by the U.S. Environmental Protection Agency (EPA), the National Academy of Sciences (NAS), and many others. Other definitions may also be found, such as the Council on Environmental Quality's (CEQ) '*possibility* of suffering harm from a hazard,' (Cohrsen and Covello, 1989: p. 1) which attempts not to enforce the mathematical representation of uncertainty from the start, though an interpretation of '*possibility*' significantly different from probability is not provided by them. For some, there are no alternatives to probabilistic uncertainty. One risk analyst asserts that to be '[mathematically] coherent ... is synonymous with being objective,' and that 'probability theory is still the only rational way that is available to us for handling uncertainty' (Apostolakis, 1990: pp. 1359,

1364).<sup>2</sup> Taken literally, this comment implies that there was no rationality, nor rational means of dealing with uncertainty, before, say, the seventeenth century, when probability was invented: Aristarchus, Archimedes, Copernicus, Galileo, Tycho Brahe were 'irrational' insofar as their scientific research contained significant uncertainties. This example is extreme, but points up the curious consequence that, taken seriously, the EPA/NAS definition implies that risk, or dangers under uncertainty, either did not exist, or could not be rigorously thought about, for most of human history.

The examples of risk definitions illustrate two issues to be addressed by this essay, and two ways in which uncertainty plays a fundamental role in the quantification of dangers or pollution. The CEQ definition shows that it's not obvious what is meant by quantified versus unquantified risk, and what relationships may exist between the two. Second, as indicated by the accusation of 'irrationality,' the value of risk analysis for some depends on whether, or to what extent, you can represent uncertainty through probability. The relationship between quantified and unquantified uncertainty, the 'taming of chance,'<sup>3</sup> so to speak, apparently concerns the legitimacy of risk analysis as a science, and the role of science in environmental policy analysis.

But what does the requirement of quantified uncertainty amount to? Does it somehow 'follow from' the idea of risk analysis as science? Or is it the case that the quantification of uncertainty places certain *demands* on science? If that is the case, can the demands always be met?

To answer these questions, I'll argue in Section 2 that *extensionality*, meaning the provisional characterization of dangerous or polluting *events* in what are intended to be largely description-independent and intentionless terms, is more fundamental to risk analysis than the probabilistic quantification of uncertainty over such events. That is, probability models of any kind assume a 'universe' (sometimes called a 'sample space') of mathematical events to which probabilities are assigned, and no events means no probability, though there still may be uncertainty. Also, *provisional* means that such characterizations are fallible and theory-laden through the theories defining them; there are no privileged 'Events,' only 'events.' Section 2 summarizes why the concept of risk provides no substitutes for this general methodological requirement. Section 3 then illustrates how extensional events are defined, discovered, or otherwise constituted through scientific research, while quantification of uncertainty associated with such events may lag considerably or not be forthcoming from science at all. The availability of such events is the 'demand' referred to above. The examples in Section 3 summarize aspects of this process of event-creation, so to speak, for three important environmental risks (nuclear winter, global warming, electromagnetic field exposure) and shows that science is not always completely successful in saying just *what risk is*. The consequence is that limiting risk analysis for policy purposes to quantified risk, therefore, at the least confuses the role of science in policy, and at worst effectively eliminates valuable scientific research from decision-making by agencies or the public. Not that probabilistic quantification of uncertainty

should be 'given up,' whatever that might mean. We all want estimates of HIV diffusion rates, the incidence of heart disease among various populations, and the frequency of automobile accidents as a function of time of day and blood-alcohol levels. The point is that the role of uncertainty in science and our understanding of dangers can go significantly beyond the standard methodological assumptions of risk analysis.

Section 4 then takes up subjective (or 'Bayesian') probability theory and decision-analytic approaches to chance as a surrogate for unquantified scientific uncertainty in environmental policy. Such approaches to the quantification of uncertainty are widely used in environmental policy as an alternative to probabilistic models created 'naturally' through scientific research. I argue that applications of subjective probability *may* mistakenly attempt to represent a social-psychological process, through which quantification is dynamically constructed and not neutrally 'assessed,' with a kind of specious objectivity. As summarized in Section 5, taken together, these two explorations into the contemporary quantification of uncertainty demonstrate two means of control, or the 'taming of chance,' in environmental policy and risk analysis: either the taming of real unquantified scientific uncertainty, or the taming of psychological, social, and political processes needed to make policy decisions in spite of uncertain knowledge. The terms of my discussion include the logic of risk and probability and methodological criteria in science, and are intended to place this argument at the foundations of risk analysis as opposed to some vaguely removed 'analysis of values,' or the sociology of scientific or policy 'practice.' We can settle our value and institutional differences and there still will remain fundamental problems in quantifying uncertainty.

## II. Extensionality in risk analysis

How might one attempt to *justify* the role of probability in risk analysis, as opposed to simply *assuming* its central role, as is often done?

A basic goal of risk analysis is to provide means for some type of canonical representation of risk, or 'risk-in-itself,' whether the risk be attributable to lung cancer, core melts, atmospheric CO<sub>2</sub>, automobile accidents, or AIDS. *We need something to count.* Even if, as analysts, we happen to lack confidence in a risk estimate, we think that 'a risk still exists, however uncertain its calculation,' (Wilson, 1987: p. 45) and a central purpose of risk analysis is to characterize, causally and quantitatively, what that risk *is*. But while this broad-brush characterization of risk analysis may seem too uncontroversial or general to be of importance, it also does not imply that the 'natural' or 'obvious' treatment of uncertainty is through probability. Nor are there alternative 'appeals' to the role of science in risk analysis which might justify that 'derivation.'

For example, risk is not amenable to the kind of axiomatic analysis provided for concepts like mass or force, in the sense that risk is not reducible to a calculus of physical abstractions. Neither is risk a 'natural kind,' like water,

fungi, neurons, or stars. That is, 'risk' does not denote a class of objects, substances, events, or processes falling under a readily identified observational or functional group which then can be analyzed through observed or theorized commonalities. Indeed, risk analyses utilize a range of principles associated with different paradigmatic contexts, such as ecological versus human health risks, and depend very much on the type of risk at issue. Neither are there genuine conjectures and refutations, nor crucial experiments, nor scientific research programs about risk *per se*. There is not, for example, nor will there ever be, a set of 'risk constants' ( $10^{-6}$  notwithstanding)<sup>4</sup> nor 'conservation laws' as part of the analytical theory of risk. Nor is there a characteristic explanatory style, such as mechanistic explanations in physics, or functional explanations in biology, or holistic-mathematical explanations in theoretical ecology. Rather, the goal of defining 'risk as it is' is achieved through an array of methodological techniques used to create risk analyses using many kinds of theories and styles of scientific reasoning. Often, risk analysis adds no fundamentally new consequences to the theories it uses; as stated by the NAS, 'In principle, risk analysis does no more than organize information from substantive disciplines in a way that allows overall estimates of risk to be computed' (National Research Council, 1989; p. 223). There is also no special category of *causality* addressed by risk analysis. Causality, for a discipline as broad as risk analysis, is largely an extra-scientific concept, added almost as a sociological honorific to a compelling body of research, say that associated with the conjecture that 'smoking causes cancer.' The risk analyst Lester Lave, for example, says that causality is 'a theoretical construct; it cannot be proved empirically. When scientists speak of proving causality, they are referring to a consensus among scientists.... Neither laboratory experiments nor epidemiology can "prove" causality' (Lave, 1982: p. 36).<sup>5</sup> *Therefore: the use of probability follows neither from a risk-specific methodology nor a risk-specific conception of causality.*

The definition of risk as a *relation*, quantified or not, also has similar consequences for the scope of risk analysis and the role of uncertainty. Environmental risk, for example, is a relation among a hazard source, exposure pathways, exposed populations or environments, and health or environmental consequences (Merkhofer and Covello: 1994, p. 2).<sup>6</sup> As may be said to make the point, a plutonium baseball in outer space poses less risk than does a can of cleanser in a child's hand. This variable relational structure results in an almost unlimited variety of causal models, ranging from fault-trees for nuclear reactors, to pharmacokinetic conveyor mechanisms in the body, to theories of chemical fate and transport in soils and groundwater. While such model structures may be organized in taxonomies, the taxonomy is purely a (valuable) organizational schema which does not provide a further theoretical basis for risk. That is, there is no essential 'structure' to risk relations, even *limiting* ourselves to quantified risk. In any instance, the 'right' model of what a risk is depends on the types of sciences involved, whether the consequences affect health or the environment, the likely principal causes of the risk, the

choice of risk metric (e.g., individual versus population risk), the current level of scientific knowledge, and heuristics used in various biological, medical, or physical sciences. While applications of, say, mass, motion, or momentum have, at least in *relative terms*, an overarching and somewhat ‘essentializing’ set of fundamental relations connecting them together, the relational concept of risk has no such framework, *in principle*. Where science lags, fails, or falls short for policy analysis, value judgments on risky outcomes, institutional control, and politics often fill the void, for better or worse. But the void is there because relational risk provides few critical methodological, explanatory, or causal constraints.

To repeat: there are no fundamental constraints on risk analysis which ‘lead to’ the quantification of uncertainty via probability, or no more so than science in general; the quantification of uncertainty through probability is a choice. What conceptual feature, then, *should* risk analysis appeal to in order to justify probabilistic analysis, and to distinguish itself from, say, interpretive ethnographies of risk, or the informal judgments sometimes called risk perceptions?

The only real conceptual requirement on quantified risk analysis, or more broadly, the non-interpretative analysis of pollutions and dangers, that is consistent with the varieties of risks and their causes is that of *extensionality*. Extensionality means that a theory takes *no* account of how objects or processes are described or represented. Human deaths ‘are’ no more than counts of individual bodies falling under mortality criteria, and may be equally-well specified as ‘lives lost’ or ‘lives saved.’ Alternative characterizations of human life should not make a difference to the outcomes studied in a quantified risk analysis. Properties that *cannot* be described in extensional terms include voluntariness, controllability, equitability, or dread, i.e., all and more of the standard ‘qualitative’ risk factors prominent in risk perception research (e.g., Slovic et al., 1990). Whether a risk is voluntary, for example, is a social fact that cannot be explicated only in extensional terms because the relevant criteria of truth may depend upon the intentional understanding of human agency or causation in social acts. The non-interpretative (i.e., not involving issues of meaning and intention) analytical task is to describe the risk to living organisms but not human beings, or to the environment but not nature. Such descriptions are typically extensional, while those theories describing *intentional* human relations are typically *intensional*, meaning that the theories *do* depend on choices for how intentional relations are described. The role of qualitative risk factors such as voluntariness or equity, for example, may be created through the language of ‘rights,’ ‘obligations,’ and ‘justice,’ all of which are manifested through the ways in which a society and its members choose to describe themselves, their rule-governed acts, and teleological behavior in general; the categories are not reducible to extensional terms *only*.<sup>7</sup>

A risk analyst would likely be unsurprised to hear that she’s been speaking extensionally all her life. Indeed, much of the best-known work on cognitive judgments under uncertainty may be summarized as saying exactly that common-sense expressions of quantified uncertainty, at least in ‘experimental’ set-

tings, reflect non-extensional reasoning.<sup>8</sup> *Most importantly, an extensional description of risk outcomes and their causes is the fundamental prerequisite to the use of probability in risk analysis; that's just how probability is defined, over mathematical 'events.'* This requirement is also *all* that distinguishes risk analysis, in general, from other risk representations or languages of risk. The reliance on probability as a measure of uncertainty will be useful only so far as there are the types of (theory-laden and revisable) events needed for probabilistic reasoning. So to the extent that scientific uncertainty is *exactly about* the definition or characterization of dangerous events, we may know just enough to be concerned about a danger, yet not enough to quantify that uncertainty via probability. That should not mean we ignore or even distort the science. Rather, we need to extend risk analysis to accommodate such unquantified scientific uncertainty, which may even be the rule rather than the exception with regard to many contemporary hazards.<sup>9</sup>

The next section supports this claim using three examples of unquantified scientific uncertainty which today are not well-represented through any useful probability distribution, and which illustrate the role of science in 'providing' or 'not providing' events for risk analysis. The point of these examples will be that the uncertainty associated with many scientific (e.g., physical, biomedical, geophysical, etc.) analyses of risk is not usefully cast in probabilistic terms because of the provisional success science has in defining events suitable for probabilistic models. But risk analysis cannot avoid extensionality insofar as probabilistic or other quantification is a goal; and as there is no special foundation otherwise in risk analysis for characterizing 'risk as it is,' the proper scope of risk analysis, especially for policy purposes, *should* be relational risk in general and include the unquantified expression of scientific uncertainty.

### III. Unquantified scientific uncertainty

Quantified risk represents a considerable narrowing of the more general definition of relational risk. While relational risk has implicitly built into it any theory of the natural world or human action, and the uncertain knowledge of both, the restriction in quantified risk to events amenable to probabilistic modeling, and uncertainty restricted to uncertainty as represented by probability theory, means that the transition from unquantified risk to quantified risk may be problematic when the special conditions of mathematical probability are not met. While quantified risk is the most natural definition for many risks, and the one almost always adopted by analysts (at least as a goal), it fails to represent all types of uncertainty found in the scientific analysis of danger. The three examples following illustrate three 'grades' of scientific uncertainty and their qualitative features not attributable to value choices, subjective bias, or 'technical versus cultural' rationality, all of which are typical 'external' explanatory notions applied to communication or policy problems associated with risk analysis.<sup>10</sup> Rather, this qualitative or structural uncertain-

ty is largely 'internal' to the science itself and is intrinsic to some of the best science available to the policy analyst today.

*Example 1.* Climate forecasting for nuclear winter, or the global weather and climate changes expected to occur following a large-scale nuclear war, is a relatively simple example of unquantified risk not readily expressed as a whole through quantified risk (Turco et al., 1990). While there may be a consensus that 'the basic physics of nuclear winter appear to be firmly established,' and specific uncertainties can be calculated for smoke emission, soot dispersion, fire ignition, and other consequences of a large-scale nuclear exchange, there is, 'as in any complex geophysical problem ... no reliable quantitative methodology for determining the overall uncertainty in existing climate forecasts for nuclear winter.' Nonetheless, 'existing data indicate that, should substantial urban areas or fuel stocks be exposed to nuclear ignition, severe environmental anomalies – possibly leading to more human casualties than the direct effects of nuclear war – would be not just a remote possibility, but a likely outcome.' That is, there is a strong likelihood of climatological perturbations accompanied by devastating consequences 'unprecedented in human history,' but of such a redoubtable complexity and uniqueness that the uncertainty associated with the hazard event – at least for one research group – resists meaningful probabilistic description: there's no probability distribution integrating component uncertainties in a non-trivial or useful way. According to one set of researchers, for theoretical (though possibly not pragmatic) purposes, the *theories* of nuclear winter may not be completely represented in a quantified risk calculus. This risk analysis contains several probabilistically quantifiable components, but no quantified global assessment, and nonetheless, a clear pragmatic recommendation on the level of risk by combining quantified and unquantified risk. A desired integration of the quantified climate indicators in an overall distribution, it is believed, has no credible basis. For these researchers, there is no useful quantitative bridge between their holistic judgment of danger and several interdependent probabilistic factors: this is a point at which the quantification of uncertainty stops. For those agencies or individuals strictly defining risk as quantified risk, strictly speaking, no total risk has been identified. At the least, in this case the possibility of communicating scientific uncertainty is hampered by identifying risk *only* with its quantified representatives.<sup>11</sup>

*Example 2.* Global warming of the earth's atmosphere and the attendant potential for climate and environmental changes on a mass scale is another risk partaking of relational risk in essential ways, notwithstanding James Hansen's polemical and widely reported expression of quantified uncertainty before Congress that he was '99%' sure that global warming had occurred (Schneider, 1989: Chapter 6). Similar to nuclear winter, there is today no accepted methodology which integrates results depending on pure numerical simulation, Bayesian models of historical temperature trends, physical-chemical atmospheric models, and theories of ecological-economic transformation, to mention just a few modeling techniques found in areas of active research.

More cautious advocates of the global warming hypothesis are less sanguine than Hansen about summarizing an entire research program in global climatology by a single number or probability distribution (Schneider, 1989: pp. 196–197), and offer a number of rationales.<sup>12</sup> For example, we may identify specific technical problems such as unrealistic models of cloud dynamics, or inadequate temperature records, as the root of skepticism in climate models; or, one may question whether asking if the climate year 1987 was, statistically speaking, ‘1% hot,’ is a valid test of the warming conjecture; or one may question specific assumptions implicit in Hansen’s claim, such as the statistical independence of climate years, or suggest correcting for urban bias induced by warm buildings; or one may propose a competing non-warming research program based on the assumption of negative feedback loops to the atmosphere through ocean cooling.

In the early 1990s this list could be continued at length, and distinguishes the risk communication on global warming from that of nuclear winter and many other risks. For global warming, the proliferation of uncertainties reflects in part the emergence of numerous research programs crossing theoretical and technological boundaries, rather than (or at least, in addition to) the intrinsic complexity of an otherwise well-understood hazard. At the same time, basic research questions are not reducible to a ‘naive falsificationist’ view of scientific progress, in the sense that the global warming conjecture can be confirmed or refuted through a crucial experiment or test criterion, specifiable in advance of carrying out the test.<sup>13</sup> With such a test defined, one might justify Hansen’s ‘99% sure’ with some rigor, but without an implied test *event*, there is no probability. More specifically, a key problem in evaluating the risk of global warming is whether, e.g., warming (if it is occurring) is due significantly to carbon emissions from fossil-fuel electric power plants, not whether warming is occurring per se. But defining a simple ‘test’ for that conjecture is not what Hansen’s ‘99% sure’ could ever mean. This requirement for a critical event is called the ‘clairvoyance test’<sup>14</sup> in decision analysis, and is precisely what a policy analyst should look for in probabilistic assertions of uncertainty. That is, one should be able to ‘look into the future’ like a clairvoyant and be able to specify the conditions which, in principle, show whether the probabilistic event has occurred or not. But the clairvoyance test often fails when applied to scientific research. Rainfall tomorrow in California is OK, temperature change this year in a well-defined locale is also fine, but *global warming caused by excessive carbon emissions*, and many risk hypotheses of the form ‘physical process *X* causes risk *Y*’ cannot be so construed, as an unknown causal structure is just what is needed to define relevant *events*.

Indeed, one function of global warming research programs is to *define* just the tests which will serve as corroborations or refutations of the existence, mechanisms, and causes of global warming. In other words, Hansen’s famous risk message implicitly implies the maturing, rather than the beginning, of research programs whose goal is to *find* the events making a claim like his possible, regardless of its numerical meaning. Unlike nuclear winter, absence



of probabilistic quantification for global warming reflects perhaps less the inherent complexity of greenhouse phenomena than the level of scientific progress. There are here, then, two distinct types of qualitative scientific uncertainty.

Rather than there being a single testable conjecture for global warming and greenhouse effects, there is a group of developing *research programs*, whose observable or verifiable consequences are steps removed from the core of competing assumptions driving the programs forward.<sup>15</sup> These research programs and the competition between them consist of unfinished *series* of theories developing over time, connected through a network of empirical and theoretical conjectures, corroborations, anomalies, technical puzzles, outright inconsistencies, shared and competing notions of physical and quantitative method, and varying heuristics for moving the programs forward in time and defending them from attack; all these categories of scientific research programs are amply illustrated by popular and technical accounts of greenhouse theory and debates, and serve to organize much of greenhouse risk communication. The methodological categories of greenhouse research programs are also used to justify much of current scientific analysis of global warming, but these substantive aspects of unquantified risk are not expressed well, if at all, through probabilistic quantification, and their roles are not necessarily reducible to the subjectivity or politics of scientific judgment, as often suggested: the categories listed are categories of theories and representations *used* in research programs. These are the ‘qualitative dimensions’ of greenhouse risk analysis, and concern resolving technical puzzles and theoretical anomalies, creating theoretical and empirical corroborations, elaborating core assumptions into working models, and so on. Such qualitative dimensions exist throughout modern science, and in the case of global warming, happen to drive much risk communication.

*Example 3.* The final example of a risk whose scientific uncertainty is poorly represented in probabilistic terms is the potential cancer risk due to 60 hertz, alternating current electromagnetic fields (EMF). These fields are created by household electrical appliances, by industrial machinery, and by transmission and distribution lines; EMFs are ubiquitous in electrified countries world-wide (Nair et al., 1989). Two kinds of methodological issues are involved in quantifying the health risk, if any, due to EMF exposures. The first set of issues has to do with the anomalous character of EMF risk compared to most chemical or environmental risks. For example, biological effects of EMF on the ability of cells to communicate with one another may occur in highly non-linear patterns, such as ‘windows’ of field frequency and intensity, between which no biological effects do occur. Furthermore, even for windows in which effects occur, increasing frequency or intensity does not necessarily lead to stronger effects; at least at the cellular level, it is not apparent that ‘more is worse,’ as is often assumed for better-understood and regulated risks. It is also not apparent what measure of dose is relevant for epidemiological studies, among time-averaged, threshold, cumulative, or other EMF exposure

metrics. For these reasons, the problem of simply defining, much less assessing, EMF health risks in terms of sources, exposure, and consequences poses technical and theoretical difficulties of the first order, while sufficient concern has been raised by biological and epidemiological studies to rule out the categorical denial that EMF could possibly pose any health risk, as was asserted openly by many until only recently.<sup>16</sup> As with global warming, an important focus of non-quantitative, scientific risk discourse on EMF is the structure and progress in EMF research; as for nuclear winter, there are fundamental theoretical complexities resisting solution as well.

The second set of issues involved in quantifying EMF health risk consists in how causal theories may direct the transition from unquantified to quantified risk. No biophysical mechanism is now known which would account for EMF health effects, and it is possible that the discovery of an explanatory mechanism will generate a significant reorganization<sup>17</sup> in cellular biophysics, meaning a major change in the fundamental assumptions on which cellular research programs are based, with a scope going well beyond the puzzles presented by exposures to appliances or power-lines. This situation differs considerably from that of global warming, where there is not an expectation that, say, some physical conservation law will require ‘exceptions’ to explain greenhouse phenomena. For EMF, the shift would consist in the rejection of the core assumption, widely held until only the last decade, that EMF, as a form of nonionizing radiation, could have no significant biophysical (*a fortiori*, health) effects because it does not carry sufficient energy either to heat tissue or break chemical bonds.<sup>18</sup> Unlike global warming or nuclear winter, the quantitative assessment of EMF as a public health risk may depend on a radical and epoch-making reorganization of some scientific and medical knowledge. Hence we have a third level of qualitative uncertainty, beyond analytical complexity and progress in evolving research programs, corresponding to epistemic change on the order of replacing or significantly revising core assumptions of multiple research programs extending beyond EMF research on its own. This unquantified uncertainty also plays a key role in EMF risk communication and public policy. For example, consider the statement that ‘researchers ... put the chances at between 10% and 60% that EMFs will turn out to have some health effects.’<sup>19</sup> Either there is an *event* to which ‘10%’ and ‘60%’ refer – in which case somebody knows what ‘EMF turns out to have health effects’ should be taken to *mean* – or this is just another way of expressing a sociological insight of ‘which way this thing is going to go,’ which is not the same at all as ‘EMF has health effects’; the latter is a natural, not a social, event.

The moral of these three illustrations is that important practical cases of unquantified risk may be understood in terms of the qualitative uncertainty used to represent scientific progress and the passage from unquantified to quantified risk. As mentioned, this scientific uncertainty is not explicable in terms of value-laden choices, qualitative-perceived versus quantitative risk, nor subjective bias. The uncertainty summarized here *just is the process of*

*science itself*, and is not especially well-suited to representation through probability. The thinking is extensional, and that is what a risk analyst can draw upon; but there is no consistently meaningful quantification of key aspects of the scientific uncertainty, and the policy analyst waiting for quantification ignores useful science at her own risk.

In the next section we look at an approach with deep philosophical roots providing a different approach for the policy-maker to negotiate with unquantified uncertainty.

#### IV. Ramsey's choice

The examples of the previous section tacitly assume that scientific credibility *per se* is a critical measure of value for quantified uncertainty. Alternatively, policy- and decision-makers faced with environmental and technological problems have to make decisions, and not 'contribute' to science. The role of science for policy is to help make good decisions, not to create knowledge for its own sake. The competition of scientific research programs or the impetus to reformulate fundamental principles may be driven by imagination tangling with nature, rather than the need to allocate scarce resources.

The policy-maker then has two options available. One option used often is to relax research standards so that answers are forthcoming for policy purposes but are not defensible as science *per se*. For example, the work might be conducted by consultants and published as an agency technical report, but not be submitted to a process of peer review (Jasanoff, 1990). The characteristic feature of this approach is to make enough assumptions – whether of exposure to pesticides, groundwater transport, accidental human intervention, whatever – and assign, by estimation or fiat, numbers such that a 'reasonable' answer is forthcoming for the problem posed. The need for a *type* of answer is used to back-infer assumptions needed to create it.

U.S. EPA has used such models by adopting 'conservative' parameter estimates when needed, that is, estimates which will almost surely *overestimate* the ultimate risk, such as estimates for pesticide toxicity, exposure, or species-sensitivity. Then, if the final calculation of risk is 'low,' there will be informal confidence that the expected use is 'safe,' though a true risk estimate has typically been lost through the uncoordinated use of 'safety factors' in model development. Among the difficulties encountered by such an approach is what happens if the policy estimate is at the borderline of what is deemed 'safe' versus 'unsafe,' say a calculation of (to within an order of magnitude) a  $10^{-6}$  lifetime cancer risk, often used to define a *de minimus* level of regulatory concern from pesticide exposure. Then, obviously, a non-organic grower or pesticide manufacturer may argue that the estimate is 'really' conservative, and that the pesticide is 'really' safe, while environmentalists may be outraged when regulators broker an agreement to minimize economic impacts by slowly phasing out a chemical instead of banning it outright. This scenario is

essentially what occurred in the late 1980s for the growth-enhancer sold commercially to apple growers as Alar. But when models are allowed to introduce safety factors for conservatism into parameter estimation, it's very difficult to untangle legitimate assumptions from those invented for policy 'science.'

As another example, which will be continued below, the Department of Energy (DOE) has wanted to know the likely patterns of seismic activity, among many other physical conditions, to occur in the vicinity of the proposed high-level nuclear waste repository at Yucca Mountain, Nevada *for the next 10,000 years*. The repository is intended as permanent storage for the nuclear waste which has been accumulating for decades at the nation's nuclear power plants. Through the Nuclear Waste Policy Act and its amendments, the federal government has guaranteed that it would develop a repository for nuclear waste generated by the nation's utilities using a tax levied on electric energy production. The 10,000 year value is not arbitrary, but is defined by environmental regulations administered by the EPA. Part of the solution for the DOE, in brief, was to extrapolate known geological trends into the future, which certainly provides *an* answer, but one which has been severely criticized as geological science: 10,000 year site-specific forecasts in geology simply lack general scientific credibility (Shrader-Frechette, 1993). While for DOE the computation apparently had temporary pragmatic value, the approach leaves much to be desired in intellectual terms. Nonetheless, along with the Alar example it is illustrative and typical of one means used by policy-makers to tame scientific uncertainty for public policy. In some cases, if scientists won't produce estimates on their own, policy-makers or their proxies will have it done for them.

A second option for the policy-maker is to eschew science as the correct *ultimate* context for environmental policy. Not that science is to be discarded; far from it. But the results of science are to be seen specifically as a hand-maiden to decision-making: scientific facts, theories, research programs, and conclusions are of interest only insofar as they contribute to, say, selecting a site for the high-level waste repository; or deciding whether to retrofit miles of electrical transmission or distribution lines to mitigate potential electromagnetic field risk; or allocating more resources for ecological habitat restoration. *Selecting, retrofitting, and restoring* are all *actions* to be *decided* upon, and the knowledge one has today, scientific or otherwise, is what one has for decisions or choices that must be made *now*; the excuse that 'more research is needed' is off the table.<sup>20</sup> Contextualizing knowledge with respect to decisions, outcomes, and preferences is the perspective of *decision analysis*, and our philosophical interest is how the transition occurs between unquantified uncertainty and the probabilistic judgments required to make a decision analysis complete.

In this century, it was the logician Frank Ramsey who thought first and deepest of uncertainty in the context of action or decision. 'Let us,' wrote Ramsey in 1931, 'try to find a method of measuring beliefs as *bases of possible*

*actions ... our judgment about the strength of our belief is really about how we should act in hypothetical circumstances...*, and not by the ‘introspection’ of ‘belief-feelings’ and the like (Ramsey, 1988: p. 29, emphasis added). Thus began one major branch of modern decision theory and the formalization of subjective or ‘Bayesian’ probability theory. Ramsey provided the basic arguments and a set of axioms involving preferences for uncertain choices needed to show that mathematical probability could be interpreted as representing ‘degrees of belief.’ Regardless of the quality of our science or other knowledge base, environmental risk experts, environmentalists, policy-makers, whoever – could, in principle – quantify their beliefs as Ramsey’s probabilities as long as these numbers were derived from possible choices among variously preferred outcomes. Since ‘choosing’ a ‘preferred’ scientific hypothesis makes little sense, Ramsey’s paradigm is an alternative to scientific practice in the context of policy and decision-making.<sup>21</sup> Ramsey and others of the early twentieth century assumed that degrees of belief took as their ‘objects,’ or what degrees of belief were degrees of, *propositions*, meaning not intuitions or feelings, but idealized descriptions of outcomes or possible states of the world which occur under uncertainty. Ramsey not only showed how to interpret probability through preference and decision, he saw that the means for associating such preferences with individuals – hence ‘subjective’ probability – was the language representing such preferences.

A marvelous consequence of Ramsey’s approach, later finessed by von Neumann and Morgenstern, de Finetti, Savage, and others, is that the rule for selecting the optimal choice among competing alternatives is simple. One chooses on the basis of highest mathematical ‘expected value,’ basically, the sum of numerical values assigned to outcomes weighted by their probabilities.<sup>22</sup> The simplicity and elegance of Ramsey’s approach is unsurpassed. His theory’s computational transparency and constructive content make it entirely suitable for practical problems. For example, Ramsey’s method for assigning probabilities to propositions, or statements, is to ask for the decision-maker’s (expert’s, environmentalist’s, etc.) preferences between outcomes: ‘Are you indifferent between (A) a sure loss of \$100,00 *versus* (B) a 1/10 chance of losing 10 acres of wetlands, *or* do you prefer one prospect over the other?’<sup>23</sup> An adequate set of propositions or statements expressing, e.g., ‘option B is preferred to option A,’ can be used to derive more complex preferences among alternatives associated with uncertain outcomes, say, the prioritization of wetland restoration sites given a limited budget. Such is the basis in language for an empirical process of subjective probability assessment.<sup>24</sup>

To give a sense of the type of problem that can be addressed through Ramsey’s paradigm, and subjective probability judgments, consider the choice of which of five sites to select for the high-level nuclear waste repository mentioned above, from among Yucca Mountain (a site in volcanic tuff, and the site tentatively selected), Davis Canyon (a bedded salt site in Utah), Deaf Smith (a bedded salt site in Texas), Hanford (a basalt site in Washington), and Richton Dome (a salt dome in Mississippi) (Merkhofer and Keeney, 1987).

The possible siting 'outcomes' are complex, in fact multiple-valued, involving the aggregation of potential health impacts, environmental impacts, socio-economic impacts, and other categories.<sup>25</sup> But the basic decision-analytic approach of valuing outcomes and comparing outcomes weighted by their probabilities is the same. The decision alternatives have been defined in advance and the problem is to compare the *relative* benefits and *relative* probabilities of choices, which is all that Ramsey and his successors can provide. That is, there's nothing in decision analysis axioms telling you in some *absolute* sense whether you've assigned the right probabilities, and obviously the outcome values or utilities are wholly value-laden, which is admirably explicit in the decision-analytic approach. All you are told is how to construct *a*, not *the*, *consistent* set of probabilities, and how to calculate optimal choices from them, assuming Ramsey's axioms. In this case, it turns out that of the five options, Yucca Mountain is optimal. Hanford was selected first for political reasons, but the choice was changed later to the Nevada site; a clear improvement according to the analysis, showing its worth, but the site selection process was still overwhelmingly politicized, and the utilization of Yucca Mountain is not proceeding as scheduled today (Jacob, 1990).

As mentioned, all that decision analysis can tell you – which is a lot – is how to make complex judgments under uncertainty *consistent*. For concreteness, Table 1 provides illustrative likelihoods from the repository siting analysis, assessed from a panel of geologists:<sup>26</sup>

Table 1. Likelihoods of outcomes at various sites.

Uncertain outcome associated with site	Davis Canyon	Deaf Smith	Richton Dome	Hanford	Yucca Mtn.
Large-scale exploratory drilling occurs	$2 \times 10^{-3}$	$2 \times 10^{-3}$	$2 \times 10^{-3}$	0	0
Repository-induced dissolution of the host rock	0	0	0	0	0
Movement on a large fault inside the controlled area but outside the repository	NA	NA	0	$3.2 \times 10^{-3}$	NA
Movement on a large fault outside the controlled area	0	0	0	0	0
Extrusive magmatic event that occurs 500 to 10,000 years after repository closure	0	0	0	0	$1 \times 10^{-6}$
Incomplete sealing of the shafts and the repository	$1 \times 10^{-4}$	$2 \times 10^{-4}$	$5 \times 10^{-4}$	$1 \times 10^{-2}$	0
Unexpected features	$1.4 \times 10^{-2}$	$1.6 \times 10^{-2}$	$1.3 \times 10^{-2}$	$2.4 \times 10^{-2}$	$1.0 \times 10^{-2}$
Movement of a large fault inside the repository	0	0	0	$3.2 \times 10^{-4}$	NA
Intrusive magmatic event	0	0	0	0	0
Expected conditions, based on available information	$8 \times 10^{-2}$	$8 \times 10^{-2}$	$8 \times 10^{-2}$	$6 \times 10^{-2}$	$8 \times 10^{-2}$

What can be said about these numbers as relative or absolute likelihoods is instructive in terms of the transition from unquantified to quantified risk. The strength of the decision analysis depends, in part, on the strength of these numbers. If the numbers are credible in some absolute sense, then that strength increases; if the numbers are only relatively consistent, the strength diminishes, but not to zero by any means. Relatively consistency, however, is often enough for decision-making. For example, Table 1 shows that the ‘incomplete sealing of the shafts and the repository’ is twice as likely to occur at Deaf Smith than at Davis Canyon; and is about *two orders of magnitude* more likely to occur at Hanford than each of these three sites. More than the point estimates of likelihoods, these relative comparisons are just what Ramsey was looking for, and are fully in line with the mathematics of subjective probability. Indeed, when Hanford was selected by Congress as the initial site, this analysis was used to discredit that alternative,<sup>27</sup> and it is hard to fault this use of the study. From such relative judgments alone, one can infer comparative benefits, given a fixed set of decision alternatives.

Mathematically and axiomatically, nothing stronger than relative consistency of judgments and the comparative value of choices can be asserted. But, as occurs for many descendants of Ramsey, such a soft transition from unquantified to quantified risk is not enough, as great an improvement as it is, say, over EPA’s ‘safety factors.’ A claim of the absolute accuracy of such subjective judgments is desired, in the sense that the experts’ judgments are hoped to be *calibrated* to the *true values*. That is, in contrast to Bruno de Finetti’s polemical assertion that ‘probability does not exist’ – i.e., that subjective probability is completely a matter of self-consistent, personal judgments – in many applications of subjective probability a claim for such internal consistency is implicitly considered insufficient. Instead, some intimation is made for the absolute correctness of the judgments involved, thus creating a hybrid probability theory which is half frequentist-objectivist, and half-personalist-subjective, and, one supposes, is intended to regain some sliver of quasi-scientific objectivity and thereby to tame uncertainty.

In the repository siting case, for example, it is mentioned that psychological studies ‘suggest that professionals with training in assessing probabilities can conduct this task in a reliable manner,’ referring to the process of improving weathermen’s personal probabilistic forecasts (e.g., the likelihood of rain on a given day in a town) which, unlike the high-level repository, can be improved through empirical observation and feedback to forecasters of weather outcomes: rain, hail, sun, clouds, etc. With a 10,000 year horizon for the events at of concern, such feedback is not a serious option, and the probabilities may not be ‘updated’ from prior distributions to posterior distributions; to finesse the point, the likelihoods are ‘subjective’ but not ‘Bayesian.’ Whatever calibration is possible must be created by ensuring consistency on prior probabilities. But consistency does not imply uniqueness, and so calibration is largely a hope, though one which may be guided by knowledge of typical heuristics for judgment under uncertainty (Kahneman et al., 1982). While that limitation

has no implications for the relative ranking of sites for a high-level waste repository, there is a significant limitation on the political force that can be claimed for the analysis, and the type of political control that might be expected to ensue through this interpretation of nuclear uncertainty. As far as the axiomatic theory goes, there is no fundamental basis for calibration, no ‘third way’ between the limited, self-consistent judgments of Ramsey, de Finetti, or Savage, and the frequentist view of traditional statistics.<sup>28</sup>

Yet, the Janus-faced figure of calibrated objective-subjective probability appears throughout the probability assessment literature. For example, a distinguished review article opens with the assertion that ‘probability *encoding*, the process of *extracting* and quantifying individual judgment about uncertain quantities, plays an important role in the application of decision analysis’<sup>29</sup>; yet ‘encoding’ has no mathematical basis apart from the classical sources, nor does the ‘extraction’ metaphor, which, one presumes, is supposed to make the assessment analogous to the determination of something like a psychophysical magnitude, like perceived hue or pitch. The metaphors of extraction and encoding are subtle means of enhancing objectivity and taming uncertainty in ways not guaranteed by mathematical theory alone. There is also a large literature on ‘bias’ and ‘debiasing’ (Fischhoff, 1982) which can be taken in a relative or absolute sense. That is, there are means for eliminating many *relative* biases which tend to create inconsistencies, or, taken more strongly, there are means for eliminating biases tending to create deviations from a ‘true value.’ The former is what should be meant while the latter may often be *allowed* to be meant. In passing from unquantified to quantified risk it is more desirable, for some, to provide control and tame chance through a synthetic objectivity, rather than being satisfied with the relative improvement provided to the decision-making process by being explicit, by articulating consequences, and by identifying ranges of outcomes and the tradeoffs associated with them.

Let’s see how deeply the metaphors persist. The guidance given to encode or extract is not always that of enlightened dialogue. One source provides instruction on how to proceed, somewhat reminiscent of guidance once provided for eliciting witchcraft accusations, in the case of ‘experts who “can’t or won’t play the elicitation game”’: that is, scientists who are ‘culturally or esthetically unable to put probabilities on things [sic].’<sup>30</sup> These scientists need to be tamed:

Depending on the circumstance of the elicitation there are a number of ways in which an analyst can ‘turn up the heat’ on an expert. You can go through a series of arguments to show that answering the questions will not violate the expert’s scientific integrity. You can point out that some decisions have to be made before all the scientific facts are in hand, and ask if the expert would really feel better if people with less scientific understanding of the problems than they provided the scientific judgments on which society bases its decisions. You can try to make it clear that in this frame-



work, probability is not a statement about the real world but is, rather, a statement about the expert's belief.... *And you can simply persist until it becomes easier for the expert to cooperate than continue to resist.*

For brevity, let's call this assessment step, conducted with reluctant or intransigent assesseees, *inquisition*.

To sort out the difficulties in what is occurring in this somewhat outrageous passage, it's necessary to get back to what a subjective probability and a probability assessment are supposed to be, which will return us to the philosophical focus of our exploration. Frank Ramsey warned that these are not introspective extractions of an internal state, but objective preferences, or desired choices between alternatives, and that such choices are expressed in terms of *propositions*. A proposition, unfortunately, is a relatively mystical entity. Propositions are an abstraction of sentences and have long been eschewed by logicians, if not by as many analytic philosophers.<sup>31</sup> More straightforwardly, preferences are expressed through sentences, being simply statements of preferences, or other judgments involving uncertainty and outcomes. A probability assessment may be defined then as a consistent set of sentences, being a reconstructed fragment of discourse, involving relevant decision alternatives, outcomes, and uncertainties, and expressing the decision-maker's preferences. Interestingly, there's nothing in the notion of subjective probability and utility that is inherently cognitive; perhaps a better term of art would be *intersubjective* probability. But over the years that decision analysis has developed, vague phrases like 'measuring degrees of belief,' which Ramsey used informally, other introspective-sounding locutions, or 'extraction,' have unfortunately obfuscated the logical basis underlying assessment. An assessment is, first, *a set of sentences*, and theories of subjective probability are therefore metatheories about the sentences produced by people who make assertions involving probability. These metatheories may be purely mathematical, cognitive, psychoanalytic, social-psychological, psychophysical, mystical – take your pick; but none of that is implied in assessment per se.<sup>32</sup> For example, assertions included in an assessment involve probabilities about events, but typically, these events never directly become part of our experience, and hence there is not necessarily a psychophysical basis, as often suggested, for a kind of introspected 'hedonic' value<sup>33</sup> – yet another attempt to tame uncertainty by creating for it a quasi-physical status. Rather, what we experience primarily is *information* about events, meaning the uncertain status of events and outcomes. If my decision or preferences depends on some conjecture or hypothesis *H*, whether on the state of ecosystems, high-level waste repositories, or cancer cases, I don't act directly on whether *H* is true or false; *H* must be interpreted to be true or false through its expression in language. A decision-maker can consider preferences among only those possibilities expressed in sentences she understands. An assessment tests candidate sets of sentences for consistency under Ramsey's (or Savage's, or whoever's) rules. If inconsistent, they may be modified to make them con-

sistent; that's part of the normative content of decision analysis. Hence typically, an assessment will take the form of a dialogue which eventuates in a final set of judgments expressed as sentences involving preferences, uncertainties, and actions. This 'convergent' set, when it exists, may conventionally be called 'the assessment'; it is neither an 'extraction,' nor an 'encoding,' and its quality is not improved by inquisition.

Once assessment is framed logically, in terms of sentences and sets of (consistent) sentences, or what logicians call a *theory*, one question is immediate: Is the theory *complete*? That is, is a positive preference between alternatives, or associated probability judgments, *always* implied by the theory 'held by' the assessee? Can I always provide a judgment of preference *one way or another*? Suppose it is said that, 'I can't choose between the \$100,000 and the uncertain prospect on the ecosystem, but I'm not indifferent between them, I have no preference.' In terms of a theory *T* expressing these preferences, this is to say that the theory is incomplete, which is a legitimate fact for a person to have about their beliefs: for a sentence *s* couched in the language of the theory, neither *s* or *not-s* is contained in *T* or derivable from *T*. Our inquisitor's plight above, to wrest a judgment from the recalcitrant assessee, then has no *logical* foundation, though it may have its pragmatic motivations. It is an empirical question whether, at the start of a probability assessment, individuals 'have' a theory to be 'extracted.' Indecision does not imply indifference.<sup>34</sup> Stated positively, this means that subjective probabilities have to be constructed, they don't spring fully-blown, extracted, from anyone's head, and the preferences we construct may depend on which questions we ask ourselves. As stated by Glen Shafer, the *false* idea is that

a person's true preferences are well-defined before he deliberates and that he just needs to ask himself the right question in order to find out this true preference.... From our constructive viewpoint, we see quite a different picture. The man [in Shafer's example] does not really have a true preference, and he is looking to various arguments ... in an effort to construct one.... For the constructive view, indecision is the starting point. Before we start to work constructing preferences, we may be undecided between all pairs of acts. We may not even have thought of all the possible acts. But this does not mean we are indifferent. As we construct preferences, we eliminate some indecision. In the end we may eliminate all indecision; we may, that is to say, rank all acts in a strict order of preference.... But Savage has given us no reason why we should feel compelled to carry our elimination of indecision so far.... Why is it normative to ... rank all acts? (Shafer, 1988: pp. 197, 204).

The *elimination* of indecision, or its transformation into indifference or strict preference, is empirical and pragmatic, not logical or mathematical, while the constructive process of eliminating indecision, through an assessment, may still be cast logical terms: a 'dialogue' is simply a shared set of

sentences created by the assessee and assessor, and through which a theory is constructed. *A constructive account of preference and uncertainty therefore is built into the foundation*, where 'constructive' refers to the process of enumerating, revising, and potentially converging on an assessment. Constructive accounts of preference are now finding favor among cognitive researchers (Gregory et al., 1993) who criticize the essentialist, 'extraction' perspectives of some economists who may be asked by the government to 'estimate' the environmental value to the public, for example, of damage associated with disasters such as the Exxon Valdez Alaska oil spill. Our account shows that the constructive approach is built into the mathematical foundations. The construction may involve a single decision-maker or a politically active group. In the former case 'cognitive' issues may be important, but there is no reason that political or social-psychological factors should not play roles as well, and in the latter case they likely will.

There is then, no 'extraction' or 'encoding' as is often expressed; these are misleading mechanistic, information-theoretic metaphors, and are implicit attempts to control a process that is open-ended and collaborative. This philosophical perspective is consistent with what is being discovered by psychologists as requirements for meaningful assessments. There is also no reason to side with the inquisitor who believes that 'experts' should play the elicitation game as presented. Structurally, the interview process used to elicit expert judgments is better compared to psychoanalysis, with an attendant interest in the assessment process as an evolving dialogue (Keeney, 1977, 1982). From a communicative perspective, there are fundamental differences between an interview conducted as psychoanalytic dialogue, in which truth-claims are as much created as found, and one aimed at objectively 'extracting' pre-existing values otherwise distorted through eliminable biases, analogous to 'changing the perspective' on a visual scene. It is a fundamental and perhaps moral choice whether communication is construed as a meaning-laden conversation, or the adjustment of a mechanism. It is also the former, not the latter, which has logical priority.

Let's finally return then to Frank Ramsey and the origins of decision theory, for the wrong choice occurring there, and leading, I believe, to wayward opportunities for subjective-objective probability. Ramsey's choice for his time and place was unavoidable. His theory of language was to 'assume here Wittgenstein's theory of propositions,'<sup>35</sup> meaning, roughly, that states of the world, and preferences among them, may be thought of in terms of idealized descriptions contained in some unspecified logical calculus. This calculus is thought to be a fairly pure 'picture' of the world, or a kind of logical isomorphism, while its adequacy as an empirical representation of preferences, and even some of its logical shortcomings, such as indifference versus indecision, is largely left unquestioned. Ramsey, for all his genius, was perhaps too early in the evolution of mathematical logic to have a less speculative understanding of his own theory. The difference between what Ramsey assumed, and what later analysts and cognitive researchers have shown is needed for deci-

sion analysis practice, is the difference between a static, abstracted view of language, and a dynamic, constructive view of language, which recognizes that the objectivity of expression evolves through a temporal dialogue. It is the difference between logical positivism and a kind of austere ethnography, or the difference between the early Wittgenstein of the *Tractatus*, which Ramsey knew well, and the late Wittgenstein of the *Philosophical Investigations*, which Ramsey would never know having died so young. The contrast may also be conceived in terms of a view of language as almost a dispensable medium of expression for mathematized propositions, versus a process requiring the interaction of individuals developing rules for understanding the world as they make their way in it. The former philosophy at times also becomes, through interpretations of 'assessment,' a means of taming the latter. It is remarkable then that from one of the most formalist branches of modern philosophy, a radically anti-formalist conclusion should ensue: that making a formalism itself work requires a constructive, temporal, and social process. It is not recognizing *that* which leads to mistaken opportunities of control, extraction, and inquisition.

## V. Taming chance

By the taming of chance<sup>36</sup> I mean the use of probabilistic or statistical methods to represent uncertainty when occurring largely to the exclusion of unquantified, less structured, and less codified methods which may only be implicit in human behavior or scientific practice; typically within some extensive institutional context, such as that of a federal regulatory agency; and involving detailed procedures for defining and elaborating facts or categories associated with an uncertain science of pollution or dangers of concern. As such, the risk representations involving the quantification of uncertainty provide important means, for better or worse, of environmental control. The taming of chance is one among many intellectual technologies<sup>37</sup> guiding the control of society and nature through the processing of uncertainty and interpretations of danger.

Many risk debates are primarily value-driven, in the sense that what we choose to do, for example concerning the restoration of contaminated hazardous waste sites, is going to depend primarily on our willingness to pay for certain levels of environmental restoration: the costs are high, and are they worth it? Uncertainty, along with value, is also at the conceptual center of risk, and the taming of chance is near the center of risk debates which focus on uncertainty. When the focus of debate moves to uncertainty, then, the topic has become that of *control*. From an engineering perspective this should be clear, as uncertainty and control essentially complement one another. High uncertainty results from low information content in a message, and both imply high entropy and lack of control; conversely, when control is possible, uncertainty, of certain kinds, is reduced. The information-theoretic analogy is

only metaphorical for most environmental risks, but the relationship between control and chance is largely the same: we tame chance in order to gain control.

In this essay, we have explored some aspects of the translation process between unquantified risk and its quantified representatives. The examples of unquantified scientific uncertainty in Section 3 demonstrate that problems in this translation process lead to the conclusion that either: (i) risk analysis has no use for much of what is known as ‘science’ because the ways in which uncertainty is interpreted resist significant quantification; or (ii) risk analysis has to work within more flexible and less controlled boundaries for science in providing advice to policy-makers. Option (i) is clearly a mistake, and is not what often occurs in practice anyway. But a belief against (ii) is canonized in the narrow definition of quantified risk. I suggest this belief results from a mistaken scientific application of probabilistic methods to science itself, and is a way of creating more control than may be delivered by scientific knowledge alone.<sup>38</sup> A fear, I suspect, lurking behind acceptance of option (ii) is the bogey-man of ‘constructive’ science. The idea that scientific methodology itself is in flux, that rules are often made up as we go along, is anathema to late children of logical positivism. It is somehow believed that the historical lessons of contemporary philosophy of science – respected as it often is by practicing mathematicians and scientists – will lead to some kind of degrading anarchy. But when you look closely enough at any human activity, intellectual or otherwise, there’s a good deal of anarchy *in the sense that* rules for the activity get made up, normalized, criticized, and changed all the time. That is, it’s a constructive process, with no implications one way or the other regarding solipsism, anti-materialism, or vulgar relativism.

Which brings us to our ending point in Section 4 on intersubjective probability and its constructive assessment. Here too, there is a dependency between the representation of uncertainty and control: that is, the process must appear to be non-constructive in order to maintain or leverage certain types of institutional control. Consider, for comparison, the idea that natural resource values may, with sufficient methodological skill applied to surveys and data analysis, be more or less neutrally assessed or ‘extracted’; this being the presumption behind most economic valuations of environmental damage, including those used in courts. While it is often acknowledged that this valuation is problematic, its behavioral rationale is rarely raised. But there has been a debate, now proceeding for some years, between behavioral researchers and economists, the former arguing, largely without rebuttal, that judgments of environmental value (e.g., of the damage incurred by the Exxon Valdez) are driven by constructive processes (Fischhoff and Furby, 1988; Fischhoff, 1991).<sup>39</sup> The implicit threat here to control is that the courts and economists are not in a position to assess the true value of damaged resources. Constructive processes in probability assessments, as occur in expert judgment of uncertainty, similarly makes apparent the taming of chance and limits certain types of control over environmental risks.

What all this expresses, I believe, are some of the ways in which we *cannot* indeed tame chance or control uncertainty to the extent we, as a society, would like. In a sense, that's why the big risks we face, such as global warming or species extinctions, *are* risks: they are really not yet under control. The examples of this essay show that, in spite of the undoubted improvements brought to policy through the risk concept, we have also reached some limits of usefulness for quantifying uncertainty, and these limits are serving to buttress outmoded ideologies of knowledge, representation, quantification, and judgment. Overcoming these limits means making policy more 'scientific' by directly participating in scientific debates *without* inappropriate proxies for uncertainty. It also means engaging in policy making by fully accepting the constructive, participatory, ultimately open-ended and untamed nature of judgments under uncertainty.

## Notes

1. For example, 'By risk is meant the likelihood, or probability, that the toxic properties of a chemical will be produced in populations of individuals under their actual conditions of exposure' (Rodricks, 1992: p. 48).
2. This article also contains some odd mathematics, such as a functional definition containing a 'variable' for 'the entire body of knowledge and beliefs of the modeler' (p. 1360).
3. The phrase is due to (Hacking, 1990), being a historical account of the introduction of statistical methods and concomitant technologies of social control during the nineteenth century. Hacking's title borrows from the French poet Stéphane Mallarmé's lines that 'dice thrown never will annul chance.'
4.  $10^{-6}$  is often used as a 'de minimus' level for acceptable individual risk, say for the lifetime cancer risk posed by pesticide applications to agricultural produce.  $10^{-6}$  is to risk management what a 95% confidence interval is to applied statistics. For approaches and criticisms, see (Whipple, 1987).
5. Unfortunately, exactly the opposite of this progressive view is promulgated through the mass media; for an example of a journalistic 'proof' of cause and effect see *New York Times* (1994).
6. 'A characteristic of a situation or action wherein two or more outcomes are possible, the particular outcome that will occur is unknown, and at least one of the possibilities is undesired' (Merkhofer and Covello, 1994: p. 2). This definition neatly subsumes the probabilistic EPA/NAS definitions as a special case, and does not assume quantification of uncertainty. The relational aspect of risk is somewhat obscured by the unneeded notion of 'characteristic.' In more general terms pollution may be defined with the phrase 'dirt is matter out of place' (Douglas, 1966: p. 48), due originally to William James' *Varieties of Religious Experience*.
7. Because risk outside of risk analysis is immediately implicated in many kinds of social, psychological, and historical settings that *cannot* be explicated in extensional terms, extensionality becomes a fundamental feature differentiating the language of risk analysis from many other ways of representing dangers and pollutions, and clashes between 'lay' and 'expert' definitions of risk may sometimes be explicated as clashes of discursive styles, but *not* necessarily because of quantification. One cannot then simply be 'more sensitive' as a policy-maker to, for example, a lack of public 'numeracy.' As far as risk debates, at least in the 1990s, involve subtle interpretations and perceptions of risk in people's lives, this implies that the world is publicly treated simultaneously as a mechanical-causal and an intentional system, or simultaneously in extensional *and* meaning-laden terms. The task

entails building bridges between very different languages of risk, and can easily not work at all. That such an integration may be *routinely* required in risk communication reflects a root issue and not a superficial aspect of, e.g., the public's scientific literacy, or numeracy, or communicators' receptivity to 'feelings and fears.'

8. See Arrow (1988: p. 505): '...an element of rationality, so obvious to the analyst as to pass almost unnoticed, is its extensionality, to use the language of logic. That is, if a choice is to be made from a set of alternatives, the choice should depend only on the membership in the set and not on how the set is described ... yet the lesson of the framing experiments is precisely that [extensional] invariances do not hold. How the choice question is framed affects the choice made.' Much of the research Arrow refers to is included in Kahneman et al. (1982).
9. See Funtowicz and Ravetz (1990) for arguments for a 'post-normal' science involving both high stakes and significant unquantified uncertainty (e.g., as in global warming or EMF). My account differs in focusing solely on uncertainty and on tracking problems back to some fundamental methodological issues.
10. See National Research Council (1989: pp. 254ff.) for an exposition of the role of value judgments in risk analysis. My goal is not to disagree with such an analysis, but to help sort out difficulties associated with uncertainty only, not uncertainty in the context of values. For other accounts on value judgments and subjectivity in risk assessment see Jasanoff (1986), and O'Brien (1987); on contrasting 'modes of rationality' in risk communication see Krimsky and Plough (1988).
11. There's always the possibility of 'trivial' quantification, such as representing the likelihood of various global catastrophes with a uniform distribution between zero and one, but such a translation has barely any relation to the science: its models, types of evidence, background theories, etc.
12. On the various topics mentioned in the text below see *Science Research Reports* (1989).
13. Arguments against 'naive falsification' and, more strongly, arguments that aspects of scientific progress are only rationalized 'in hindsight,' or after the fact, are developed by Lakatos (1978: pp. 68ff.). The relevance of such *post festum* rationality for risk analysis has been noted in National Research Council (1989: p. 246).
14. The role of the clairvoyance test is clearly expressed in foundational accounts of subjective probability, e.g., DeFinetti (1977). See Koopman (1977: p. 122) for an explicit rejection of the possibility of attaching probability values to scientific conjectures violating the clairvoyance test.
15. The theory of scientific research programs developing the categories enumerated here is contained in Lakatos (1978: ch. 2), being largely a critique of Karl Popper's falsificationist account of scientific method, but also, like Popper's, a theory of knowledge without foundations, of *science under ubiquitous uncertainty in observation, modeling, and theory*. Risk analysis needs to catch up with the empirical insights of such contemporary philosophy of science.
16. See Liboff (1987: pp. 11–12) on the 'continuing tension between the "thermal" and "non-thermal" bioelectromagnetic camps. Because the engineering community has a firm grasp of what joule heating is about it has consistently attempted to explain more biological facts than it should using this parameter; when scientists have reported effects at intensities lower than can be explained thermally, these reports have either been discounted as somehow wrong or they have simply been ignored ... the author steers away from scientific discussion ... except insofar as it involves parameters [such as...] pulse shape, frequency, field strength, current.'
17. If the term was not in such disfavor, 'paradigm shift' is another description (Kuhn, 1970). In the philosophy of science, rigorous uses of 'paradigm' have been largely given up, but the intention is to indicate major changes in the structure and content of science. Such changes may be indicated by new fundamental entities or processes (e.g., atoms or bioelectromagnetic health effects) or meaning shifts in basic theoretical terms (e.g., from 'mass' naming a substance to naming a relativistic relation). As a suggestion of the potential

relevance to EMF risk (*which may just as well disappear like a phantom*) see Jauchem (1990: emphasis added) who notes that 'At this time, the variability and complicated nature of EMF characteristics do not allow researchers to even *design* definitive studies of EMF health effects. As R. A. Cartwright has stated, "The criticisms of surrogate measures [of fields] mean that *no proposed study will ever directly address the issue.*"' The possibility that EMF risk largely outstrips current understanding is also suggested by Morgan (1990: p. 118).

18. This assumption would be part of what Lakatos calls the 'hard core' of ionizing radiation research programs, since giving it up entails halting a significant research direction, for example by no longer invoking (and perhaps revising past applications of) the assumption as part of the program's 'negative heuristic,' or rules which defend the research program from anomalous results or competing programs. For an example, see the quotation in note 16.
19. M. G. Morgan, quoted in *Science Research Reports* (1990).
20. An exception is so-called 'value of information' problems (i.e., the comparative value of acting now or first obtaining more information at some cost) which does not affect what follows.
21. See text and note 14 *above* on Koopman and De Finetti.
22. If numerical outcomes  $x_1, \dots, x_n$  occur with probabilities  $p_1, \dots, p_n$ , the expected value is  $\sum x_i \cdot p_i$ . Typically  $x_i$  is a measure of value or utility so that if  $o_i$  is the  $i$ th 'physical' outcome,  $x_i = U(o_i)$  for a utility function  $U$ .
23. The informal question and others like it may be expressed as a mathematical condition involving algebraic preference models (von Winterfeldt and Edwards, 1986: ch. 9).
24. A subtlety not addressed by Ramsey in such tradeoff questions is whether there must be values of mathematical variables (the 1/10 probability, the dollar value, or the number of acres of wetlands) such that the decision-maker explicitly prefers one of the options provided, or considers the options to be equally preferable. That is, can we *always* vary the parameters in A and B such that either 'A is preferred to B,' 'B is preferred to A,' or 'B is equivalent to A' is a proposition or statement accepted or willingly asserted by the decision-maker or her proxy? For, to say that you cannot judge between two options, that you are *undecided*, is not to say you are *indifferent* between them. We'll return to this issue of the *completeness* of a set of probability or preference judgments below.
25. The application also requires what is known as multiattribute utility theory to assign numerical values to outcomes, but the relevant issue here is the role of uncertainty.
26. Taken from Tables VI and VII of Merkhofer and Keeney (1987: pp. 180–181). Fifteen scenarios are considered here, and geology panels provided high and low likelihood estimates which were included as part of the study; shown here only are base-case values. Zero values are interpreted from 'NC,' taken to mean 'not credible' and 'NA' means 'not applicable,' and hence also is zero. 'Expected conditions' likelihood was calculated as 1-(sum of likelihoods of other assessed scenarios), rather than being assessed directly.
27. The DOE has been roundly criticized for not considering options other than a permanent repository, such as monitored and retrievable storage (Shrader-Frechette, 1993; Keeney and von Winterfeldt, 1994). The latter provides a multiattribute decision analysis like the one discussed in the text showing that the present Yucca Mountain strategy is likely wasting billions of dollars compared to non-permanent storage alternatives.
28. 'Neither Ramsay, nor Savage, nor de Finetti, to name three leading figures in the personalistic movement, can find it in his heart to detect any logical shortcomings in anyone, or to find anyone logically culpable, whose degrees of belief in various propositions satisfy the laws of the probability calculus, however odd those degrees of belief may otherwise be. Reasonableness, in which Ramsey was also much interested, he considered quite another matter ... one can still be wildly unreasonable without sinning against either logic or probability' (Kyburg, 1988: p. 104).
29. Spetzler and von Holstein (1975: p. 340), emphasis added. There's a clear bias toward a mechanistic, information-processing view that is wholly arbitrary and repeated in much of



- the decision-analytic literature. For an alternative approach, see Watson and Buede (1987; pp. 271–274).
30. Morgan and Henrion (1990: pp. 161–162), emphasis added to text following. I expect that the emphasized text is meant facetiously, but that attitude masks, I think, some basic methodological problems.
  31. The standard objects of study in mathematical logic include languages, sentences, formulas, terms, and theories. Propositions can be defined in complicated ways (e.g., as equivalence classes with respect to truth in a model or provability within a theory) but they have never gained the fundamental status envisioned for them by Russell, the early Wittgenstein, or many logical positivists; propositions represent a poor theoretical construct. To the best of my knowledge, probability assessment has never been defined, as in the text below, using logical categories.
  32. On some of the ideas which follow, see Hacking (1988). This paper (first published in 1967) is remarkable for its consistency with the methodological views contained in Hacking's later work on the history of probability and statistics.
  33. See, e.g., von Winterfeldt and Edwards (1986: p. 209) on 'strength of preference' being no different from psychophysical judgments of 'loudness, pitch, brightness, and the like...'
  34. See Shafer (1988) and Shafer and Tversky (1988: pp. 238, 263): 'Thought of in this way, a theory of subjective probability is very much like a formal language ... expressing probability judgments rather than as psychological models, however idealized ... . Probability judgment is a process of construction rather than elicitation.'
  35. Ramsey (1988: p. 33, n. 2). Even experts stumble their way through propositions, for example Lindley (1972: p. 4) who confuses Ramsey's concept of an 'ethically neutral proposition'  $p$  with its associated event, as shown by Lindley's taking the set-theoretic complement of  $p$ , which only makes sense for events.
  36. My use of Hacking's phrase should not be taken as representative of his detailed study (see note 3), which at any rate is specific to the last century.
  37. 'Technology' is not intended to be pejorative, but taken, e.g., in the sense of von Winterfeldt and Edwards (1987: frontispiece), who consider decision analysis an intellectual technology and practice for making decisions.
  38. See Landy et al. (1990) on how EPA turns policy problems into pseudo-technical problems; here is yet another instance.
  39. For an example of how the media frames the assessment issue toward economists, see *New York Times* (1993).

## References

- Apostolakis, G. (1990). 'The concept of probability in safety assessments of technological systems,' *Science* 250: 1359–1364.
- Arrow, K. (1988). 'Behavior under uncertainty and its implications for policy,' in David Bell, Howard Raiffa, and Amos Tversky, eds., *Decision Making: Descriptive, Normative and Prescriptive Interactions*. New York: Cambridge University Press.
- Cohrssen, J. and V. Covello (1989). *Risk Analysis: A Guide to Principles and Methods for Analyzing Health and Environmental Risks*. Washington D.C.: Council on Environmental Quality.
- Definetti, B. (1980). 'Probability: Beware of falsifications,' (first published 1977) in Henry Kyburg, Jr. and H. Smokler, eds., *Studies in Subjective Probability*. New York: John Wiley and Sons, 2nd ed.
- Douglas, M. (1966). *Purity and Danger*. London: Penguin.
- Fischhoff, B. (1991). 'Value elicitation: Is there anything in there?' *American Psychologist* 46: 835–847.
- Fischhoff, B. (1982). 'Debiasing,' in D. Kahneman et al., eds., *Judgment under Uncertainty: Heuristics and Biases*. New York: Cambridge University Press.

- Fischhoff, B. and L. Furby (1988). 'Measuring values: A conceptual framework for interpreting transactions with special reference to contingent valuation of visibility,' *Journal of Risk and Uncertainty* 1: 147–184.
- Funtowicz, S. and J. Ravetz (1990). *Uncertainty and Quality in Science for Policy*. Dordrecht: Kluwer Academic Publishers.
- Gardenfors, P. and N. Sahlin, eds. (1988). *Decision, Probability, and Utility*. New York: Cambridge University Press.
- Gregory, R., S. Lichtenstein, and P. Slovic (1993). 'Valuing environmental resources: A constructive approach,' *Journal of Risk and Uncertainty* 7: 177–197.
- Hacking, I. (1988). 'Slightly more realistic personal probability' (first published 1967), in P. Gardenfors and N. Sahlin, eds., *Decision, Probability, and Utility*. New York: Cambridge University Press.
- Hacking, I. (1990). *The Taming of Chance*. New York: Cambridge University Press.
- Jacob, G. (1990). *Site Unseen: The Politics of Sitting a Nuclear Waste Repository*. Pittsburgh: University of Pittsburgh Press.
- Jasanoff, S. (1986). *Risk Management and Political Culture*. New York: Russell Sage Foundation.
- Jasanoff, S. (1990). *The Fifth Branch: Science Advisers as Policymakers*. Cambridge: Harvard University Press.
- Jauchem, J. (1990). Letter to *Science* 249: 739.
- Kahneman, D., P. Slovic, and A. Tversky, eds. (1982). *Judgment under Uncertainty: Heuristics and Biases*. New York: Cambridge University Press.
- Keeney, R. (1977). 'The art of assessing multiattribute utility functions,' *Organizational Behavior and Human Performance* 19: 267–310.
- Keeney, R. (1982). 'Decision analysis: An overview,' *Operations Research* 30: 803–838.
- Keeney, R. and D. von Winterfeldt (1994). 'Managing nuclear waste from power plants,' *Risk Analysis* 14: 107–130.
- Koopman, B. (1980). 'The bases of probability' (first published 1940) in H. Kyburg, Jr., and H. Smokler, eds., *Studies in Subjective Probability*. New York: John Wiley and Sons, 2nd ed.
- Krimsky, S. and A. Plough (1988). *Environmental Hazards: Communicating Risks as a Social Process*. Dover, MA: Auburn House.
- Kuhn, T. (1970). *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press, 2nd ed.
- Kyburg, H. (1988). 'Bets and beliefs' (first published 1968), in P. Gardenfors and N. Sahlin, eds., *Decision, Probability, and Utility*. New York: Cambridge University Press.
- Kyburg, H. Jr. and H. Smokler, eds. (1980) *Studies in Subjective Probability*. New York: John Wiley and Sons, 2nd ed.
- Lakatos, I. (1978). *The Methodology of Scientific Research Programmes: Philosophical Papers I*, J. Worrall and G. Currie, eds., New York: Cambridge University Press.
- Landy, M., M. Roberts, and S. Thomas (1990). *The Environmental Protection Agency: Asking the Wrong Questions*. New York: Oxford University Press.
- Lave, L. (1982). *Quantitative Risk Assessment in Regulation*. Washington: The Brookings Institute.
- Liboff, A. (1987). Review of F. Carstensen's *Biological Effects of Transmission Line fields in Microwave News* 1 (July/August): 11–12.
- Lindley, D. (1972). *Bayesian Statistics: A Review*. Philadelphia: SIAM.
- Merkhofer, M. and R. Keeney (1987). 'A multiattribute utility analysis of alternative sites for the disposal of nuclear waste,' *Risk Analysis* 7: 173–194.
- Merkhofer, M. and V. Covello (1994). *Risk Assessment Methods*. New York: Plenum Press.
- Morgan, M. G. (1990). Review of P. Brodeur's *Currents of Death*, in *Scientific American* 267: 118.
- Morgan, M. G. and M. Henrion (1990). *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*. New York: Cambridge University Press.

- Nair, J. et al. (1989). *Power-Frequency Electric and Magnetic Fields: Exposure, Effects, Research, and Regulation*. Washington, D.C.: U.S. Congressional Office of Technology Assessment.
- National Research Council (1989). *Improving Risk Communication*. Washington D.C.: National Academy Press, 1989.
- New York Times* (1993). 'Polls may help government decide the worth of nature,' September 6.
- New York Times* (1994). 'Cigarette makers debated risks they denied,' June 16.
- O'Brien, D. (1987). *What Process is Due?: Courts and Science-Policy Disputes*. New York: Russell Sage Foundation.
- Ramsey, F. (1988). 'Truth and probability' (first published 1931), in P. Gardenfors and N. Sahlin, eds., *Decision, Probability, and Utility*. New York: Cambridge University Press.
- Rodricks, J. (1992). *Calculated Risks*. New York: Cambridge University Press.
- Schneider, S. (1989). *Global Warming: Are We In the Greenhouse Century?* San Francisco: Sierra Club Books, 1989.
- Science Research Reports* (1989). *Science* 243-244: 28-29, 1041-1047, 1119.
- Science Research Reports* (1990). *Science* 249: 23.
- Shafer, G. (1988). 'Savage revisited,' (first published 1986) in D. Bell, H. Raiffa and A. Tversky, eds., *Decision Making: Descriptive, Normative and Prescriptive Interactions*. New York: Cambridge University Press.
- Shafer, G. and A. Tversky (1988). 'Languages and designs for probability judgment,' (first published 1985) in D. Bell, H. Raiffa and A. Tversky, *Decision Making: Descriptive, Normative and Prescriptive Interactions*. New York: Cambridge University Press.
- Shrader-Frechette, K. S. (1993). *Burying Uncertainty: Risk and the Case Against Geological Disposal of Nuclear Waste*. Berkeley: University of California Press.
- Slovic, P., B. Fischhoff, and S. Lichtenstein (1990). 'Rating the risks,' in T. Glickman and M. Gough, eds., *Readings in Risk*. Washington D.C.: Resources for the Future, pp. 61-75.
- Spetzler, C. and C. Stael von Holstein (1975). 'Probability encoding in decision analysis,' *Management Science* 22: 340-358.
- Turco, R. P. et al. (1990). 'Climate and smoke: An appraisal of nuclear winter,' *Science* 247: 166-174.
- von Winterfeldt, D. and W. Edwards (1986). *Decision Analysis and Behavioral Research*. New York: Cambridge University Press.
- Watson, S. and D. Buede (1987). *Decision Synthesis*. New York: Cambridge University Press.
- Whipple, C. ed. (1987). *De Minimus Risk*. New York: Plenum Press.
- Wilson, R. (1987). 'Estimating risks of known and unknown carcinogens,' in L. Lave, ed., *Risk Assessment and Management*. New York: Plenum Press.