Methodological Rules for Four Classes of Scientific Uncertainty

Kristin Shrader-Frechette

In the face of inadequate data or unclear theories, is a hypothesis confirmed or falsified? Is a potential scientific threat innocent till proved guilty or guilty till proved innocent? When potential harm is unknown, does one assume that an average level of danger will occur or that a worst case is possible? All of these questions raise the issue of how to deal with scientific uncertainty, particularly in cases that could have serious consequences.

Among the many classes of cases of scientific uncertainty that confront contemporary problem solvers, at least four stand out because of their special relevance to environmental decision making. The four classes involve (1) framing uncertainty, (2) modeling uncertainty, (3) statistical uncertainty, and (4) decision-theoretic uncertainty. Using examples from quantitative risk assessment, especially hazards associated with nuclear wastes, this essay argues that there are a number of epistemologic and ethical rules that scientists ought to follow in each of these cases of uncertainty and that, indeed, many of the preferred rules are contrary to established principles that members of the scientific community actually follow. Consider first the uncertainties surrounding how scientists frame their questions.

Framing Uncertainty and Using Science for Policy[1]

How scientists frame their questions often controls their answers. Newton, for example, framed his questions about mechanics in terms of the assumption that he needed to explain what caused uniform rectilinear motion *to stop*. As a consequence, he affirmed the first law of motion (1). Aristotle, however, framed his questions about mechanics in terms of the assumption that he needed to explain what caused uniform rectilinear motion *to begin*. As a result, he denied

[1]Much of the discussion of framing uncertainty is based on Shrader-Frechette (15, 18).

that bodies remain either at rest or in motion, unless compelled by impressed forces to change their state (2).

Like Newton and Aristotle, different scientists frequently have alternative "frames," different sets of theoretical assumptions for structuring their data and their problem solving (3, 4). Some of the most basic scientific uncertainties concern how to frame a particular question, for example, when and how to interpret data as providing grounds for accepting particular hypotheses. In framing their questions, scientists frequently use a two-value frame (falsification/provisional acceptance) specifying that, in a situation of uncertainty, when rigorous testing fails to falsify some testable hypothesis (such as "organic molecules can come in right- and left-handed mirror-image versions"), then it is reasonable to accept the hypothesis provisionally.

Why Scientists Often Use the Two-Value Frame

Scientists sometimes employ the two-value frame because, from an empirical point of view, rigorously attempting (and failing) to falsify a precise testable hypothesis provides one of the strongest criteria for accepting it (5). Scientists often attempt to devise "crucial experiments," tests for which two mutually exclusive, exhaustive hypotheses predict conflicting outcomes. Classic examples of "crucial experiments" are Millikan's (6) attempt, in the early 1900s, to show whether electric charges are integral multiples of the charge of the electron and Lenard's 1903 test of two conflicting implications concerning the light energy that a radiating point can transmit (7). A second reason that scientists sometimes use the two-value frame is pragmatic. As Duhem (8) recognized, they provisionally accept a precise hypothesis (that has survived rigorous attempts to falsify it), even one with obvious deficiencies, if there is no better (more probable) hypothesis available. Because hypotheses have an infinite number of observational consequences that can never be "verified" conclusively (9), scientists sometimes opt—in a situation of uncertainty—for provisional acceptance of the best available nonfalsified hypothesis.

A third situation in which scientists use the two-value frame is when they give provisional acceptance to null (no-effect) hypotheses that survive rigorous attempts at falsification. Because they are more interested in avoiding false positives (type I errors) rather than false negatives (type II errors) in situations of uncertainty, scientists place the greater burden of proof on the person who postulates some, rather than no, effect. For example, a geologist might postulate the effect that, because of tectonic activity, the water table will rise at a given location by at least 500 meters over the next 10,000 years. Although "no-effect" results run the risk of type II errors, scientists usually assume, as in criminal law, that null hypotheses are provisionally acceptable (innocent) until they are rigorously falsified (proved guilty). In the third section, on statistical uncertainty, we discuss type I and type II errors in more detail.

The Two-Value Frame in Quantitative Risk Assessment

Because policy makers often need immediate decisions about particular risks, scientists have used the two-value frame in many risk assessments, from studies of hazardous landfills to childhood exposure to lead (10). To illustrate potential problems with assessors using the two-value frame, consider the 1992 *Early Site Suitability Evaluation* (ESSE) completed by the U.S. Department of Energy (DOE) for the proposed Yucca Mountain (Nevada) nuclear waste repository (11). Reporting site-suitability findings for every condition specified, the 1992 ESSE used a two-value frame to assess the site: "conclusions about the site can be either that current information supports an unsuitability finding or that current information supports a suitability finding. . . . If . . . current information does not indicate that the site is unsuitable, then the consensus position was that at least a lower-level suitability finding could be supported" (11, pp. E-5, E-11). To understand why the two-value frame may be problematic here, recall that scientists typically use it for at least one of three reasons: (1) The attempted falsification is rigorous and precise. (2) The hypothesis that has survived precise, rigorous attempts at falsification is the "best" of candidate hypotheses. (3) The surviving hypothesis is a null hypothesis.

Sometimes societal risk assessments meet none of these three conditions for use of the two-value frame. At Yucca Mountain, for example, the long time period (10,000 years of site suitability) precludes the precise predictions, specified in condition (1), necessary for rigorous attempts at falsification. Indeed, the ESSE peer reviewers unanimously warned: "many aspects of site suitability are not well suited for quantitative risk assessment. . . . Any projections of the rates of tectonic activity and volcanism, as well as natural resource occurrence and value, will be fraught with substantial uncertainties that cannot be quantified" (12, p. B-2). They cautioned that although "there is . . . currently not enough defensible, site-specific information available to warrant acceptance or rejection of this site" (12, pp. 460, 257, 40–51), nevertheless they used the two-value frame (site suitable/site unsuitable) that they "were given" and agreed with the "site-suitable" conclusions of the ESSE: "The DOE General Siting Guidelines (10 CFR Part 960) do not allow a 'no decision' finding. . . . Thus the ESSE Core Team followed the intent of the guidelines" (12, p. 460). The peer reviewers' warnings suggest that, when rigorous and precise testing is impossible, using a two-value frame might beg some questions of risk evaluation and that, in such situations, a three-value frame (site suitable/site unsuitable/site suitability uncertain at present) might be preferable.

Condition (2), that the nonfalsified hypothesis be the best available, likewise appears problematic for assessments that compare neither alternative risk sites nor different risk hypotheses. Complete comparative analyses of proposed repository sites, for example, were precluded by the 1987 amendment to the Nuclear Waste Policy Act that named Yucca Mountain, Nevada, as the only candidate location for the nation's first permanent repository for commercial nuclear waste and spent fuel. Using the two-value frame to give provisional acceptance to a particular site-suitability hypothesis as "best," however, appears problematic to the degree that alternative sites are not compared and to the degree that different hypotheses have "substantial uncertainties that cannot be quantified" (12, p. B-12).

Another justification for use of the two-value frame—that (3) surviving hypotheses be provisionally accepted if they are null—likewise seems inapplicable in many societal risk assessments. Theoretical science, of course, often places the burden of proof on those arguing against the null hypothesis because science must be *epistemologically conservative* (avoid false positives). Science applied to risk assessment, however, also must be *ethically conservative*, as the National Academy of Sciences points out (13, 14), in the sense of taking account of social consequences affecting the needs, rights, and welfare of the public. To the extent that the public has limited financial resources and information or bears inequitable or involuntary risk impositions, it may need more risk protection than the proponents of a particular null hypothesis regarding risk (15). For example, because more than 80% of Nevadans say they would vote against the Yucca Mountain repository proposed for their state (16), they may need greater risk protection. Future generations, in particular, may have special needs regarding Yucca Mountain both because they cannot exercise their consent and because current regulations require no monitoring beyond the first 50 years, after which waste migration is more likely. Likewise, in cancer risk assessment, potential victims of a false null hypothesis may need special protection because many epidemiologic studies are too insensitive—owing to small samples and their dealing with rare diseases—to detect positive effects. Also, field studies of populations exposed to hazardous substances often involve more uncertainties than those based on theoretical models (10).

Framing Uncertainties: Hypotheses Versus Decisions

Other disanalogies between theoretical science and science applied to risk assessment also argue against using the two-value frame for evaluating societal risk in a situation of uncertainty. Theoretical scientists usually evaluate the *truth or falsity* of *hypotheses* (such as "convection currents have moved this geological plate"). Risk assessors, however, also evaluate the *acceptability* of risk *decisions* (such as "this site is suitable for permanent waste disposal"). As the National Academy of Sciences put it: "risk assessment must always include policy as well as science" (13, p. 76). Because the acceptability of risk decisions includes non-epistemic factors—such as decision-theoretic, social, economic, and ethical considerations—risk assessment may be more suited to a three-value frame that explicitly takes account of uncertainty. For example, classical methods of Bayesian decision making typically employ a three-value frame, in the sense of including a category for events that are "uncertain" or about which we have inadequate information to make a decision (17). Decision theorists also recognize that even a high *probability* that a site is suitable for some activity may not be "high enough" if the activity could pose serious *consequences* for public welfare.

Because of the disanalogies between theoretical science and science applied to societal risk evaluation, assessors confronting framing uncertainties may need to consider using three-value, rather than two-value, frames for scientific decisions involving both significant uncertainty and potentially serious public consequences. Regardless of the frames they choose, however, scientists and policy makers may need to recognize that uncertainty gives framers significant power. As happened with Newton, Aristotle, and risk assessors, whoever frames the questions may control the answers (18).

Modeling Uncertainty, Verification, and Validation[2]

Another type of uncertainty that occurs frequently in science arises when modelers assume that their constructs have been verified or validated because they are consistent with other computer models. This second section of the chapter argues that, in cases of modeling uncertainty created by incomplete data or failure to employ the available data, scientists ought not claim that their models have been verified or validated.

Modeling Uncertainty and Affirming the Consequent

Problems with verification and validation of models are part of a larger set of difficulties associated with the inference known as "affirming the consequent." This inference occurs whenever one postulates that a hypothesis is true or accurate merely because some test result, predicted to follow from the hypothesis, actually occurs. In fact, however, failure of predictions can only falsify theories, but success of predictions can only confirm (but not verify) theories. All that can be validly inferred from a test is that the results are consistent with the hypothesis or that the results have falsified the hypothesis. In other words, from "*h* entails *r*," one can infer "not *r* entails not *h*." To assume that one can infer "*r* entails *h*" from "*h* entails *r*" is to affirm the consequent. Of course, it is very important to test one's hypotheses in order to determine whether the data falsify them or tend to confirm them. Moreover, the greater the number of tests, and the more representative they are, the greater is the assurance that the data are consistent with the hypotheses. Indeed, one of the repeatedly acknowledged failures of the assessments of the proposed Yucca Mountain nuclear waste repository is that the models often are not tested (19). It is important to test the models, to attempt to falsify them and to determine the degree to which they are consistent with the data. If the models turn out to be consistent with the data, however, it is wrong to assume that they have been absolutely "verified" or "validated" because, short of affirming the consequent, it is impossible to verify or validate any model. It is possible merely to know—through testing—that the hypothesis or model has been confirmed to this or that degree.

At the proposed Yucca Mountain repository, risk assessors have repeatedly proposed to test some *h*, some hypothesis, such as that the number of calculated groundwater travel times is less than 10,000 years. When the calculations, data, and models are shown to be *consistent* with the hypothesis, then the assessors have erroneously assumed, in the face of modeling uncertainties, that the hypothesis has predictive power or has been "verified." For example, one group of assessors, studying groundwater travel time, concluded: "this evidence indicates that the Yucca Mountain repository site would be in compliance with regulatory requirements" (20). Many other risk assessors speak of "verifying" their models and "validating" them. For instance, one group of assessors concluded that the tools they used demonstrated "verification of engineering software used to solve thermomechanical problems" (21, p. i) at Yucca Mountain (22).

Admittedly, software and systems engineers speak of computer models being "validated" and "verified." Yet, such "validation" language obscures the fact that the alleged validation really only guarantees that certain test results are consistent with a model or hypothesis; it does not validate or verify the model or hypothesis because affirming the consequent prevents legitimate validation or verification. Hence, when computer scientists speak of "program verification" (23–25), at best they are making a problematic inference by affirming the consequent in the face of modeling uncertainty. At worst, they are trading on an equivocation between "algorithms" and "programs." As Fetzer argues (26, 27), algorithms, as logical structures, are appropriate subjects for deductive verification. As such, *algorithms* occur in pure mathematics and pure logic. They are subject to demonstration or verification because they characterize claims that are always true as a function of the meanings assigned to the specific symbols used to express them. *Programs*, however, as causal models of logical structures, are not verifiable because the premises are not true merely as a function of their meaning. As Einstein put it, insofar as the laws of mathematics refer to reality, they are not certain; insofar as they are certain, they do not refer to reality.

Modeling Uncertainties and Misleading Language

In using "verification" and "validation" language, both official U.S. Department of Energy (DOE) documents and individual risk assessments for repositories like

[2]Much of the discussion of modeling uncertainty is based on Shrader-Frechette (15) and Oreskes, Shrader-Frechette, and Belitz (1994).

Yucca Mountain are systematically misleading both about the modeling uncertainties and about whether the studies are reliable. For example, explicitly affirming the consequent, the DOE affirmed (28, p. 3-11)

> Validation . . . is a demonstration that a model as embodied in a computer code is an adequate representation of the process or system for which it is intended. The most common method of validation involves a comparison of the measured response from in-situ testing, lab testing, or natural analogs with the results of computational models that embody the model assumptions that are being tested.

Authors of the same official DOE document, used to provide standards for Yucca Mountain risk assessments, also talk about the need to verify computational models of the waste site. They say (28, p. 3-7)

> Verification, according to the guidelines in NUREG-0856 . . . is the provision of assurance that a code correctly performs the operations it specifies. A common method of verification is the comparison of a code's results with solutions obtained analytically. . . . Benchmarking is a useful method that consists of using two or more codes to solve related problems and then comparing the results.

Although the term "verification," as used by DOE assessors, suggests that the computer models or codes accurately represent the phenomena they seek to predict, it is merely a misleading euphemism for "benchmarking," comparing the results of two different codes (computer models) for simulating an identical problem. On this scheme, one "verifies" a model of Yucca Mountain against another model. What is required in the real world, however, is validating a model against reality. This validation or confirmation can be accomplished only by repeated testing of the code or model against the real world, against field conditions.

Even with repeated field testing, however, modeling uncertainties remain. Compliance can never be confirmed, short of full testing of all cases throughout all time periods. Classic studies of the problem of induction show that complete testing is impossible. Therefore, the shorter the time of testing and the fewer the cases considered, the less reliable and the less confirmed are allegedly "validated" computer models or codes. The tests can only falsify or confirm a hypothesis, not validate it. To assume otherwise is to affirm the consequent. Hence, every conclusion of compliance with government regulations, or every conclusion of repository safety, on the basis of "verified" or "validated" test or simulation results, is an example of affirming the consequent. Program *verification*, in other words, "is not even a theoretical possibility" (29). One cannot *prove* safety. One can only demonstrate that one has attempted to falsify one's results and either has failed to do so or has done so.

Responses to Modeling Uncertainties

Because of the problems associated with modeling uncertainties, scientists need to be wary of claiming that they have verified or validated a model on the basis of limited data. As both the DOE risk documents and the risk assessors at Yucca Mountain illustrate, they are misleading in speaking of "validation" and "verification" of models used at Yucca Mountain. *First*, real validation and verification are impossible because of the problems of induction and affirming the consequent. Only falsification of a hypothesis, or determining that the data are consistent with it, is possible. In the latter case, when one obtains repeated results indicating that the data are consistent with the model or hypothesis, one is able merely to increase the probability that the model or hypothesis has been confirmed. *Second*, the DOE's and assessors' use of the terms "verification" and "validation" misleads the public about the reliability of models allegedly guaranteeing repository safety. *Third*, use of the term "verification" by DOE assessors is, in particular, misleading because they typically only compare different computer codes or models, with no reference to the real world, and because any model can be tuned or calibrated to fit any pattern of data, even when the model is not well confirmed. *Fourth*, it is arguable that most useful programs are not merely unverifiable but incorrect, that even programs that function correctly in isolation may not do so in combination, and that most of the important requirements of real programs are not formalizable (30).

Given these four difficulties with "verifying" programs used in real-world situations, such as repository modeling, there are both prudential and ethical problems with risk assessors' continuing to use the language of "program verification" in connection with modeling causal relationships in situations of uncertainty. The *prudential* problem is that aiming at "verification" does not tell us what we most want to know—something about complex relationships in the physical world. The more complex the system, the less likely it is to perform as desired and the less reliable is inductive testing of it. Moreover, by emphasizing verification, theorists have increased the expense of achieving "transparent software upgrades," and they have decreased software reliability because of their emphasis on "misplaced advocacy of formal analysis" (31, p. 792). The *ethical* problem is that, by encouraging confidence in the operational performance of a complex causal system, claims of "verification" oversell the reliability of software and undersell the importance of design failures in safety-critical applications like waste repositories. Such overselling and underselling not only expose the safety of the public to the dangerous consequences of risk assessors' "groupthink" (32, p. 422), but also risk misunderstanding of software in cases where the risks are greatest. To avoid affirming the consequent, invalid inferences such as that repository safety models can be "verified," scientists need to refrain from the claim that their results "indicate" or "show" or "prove" compliance with government regulations or with some standard of safety. Scientists also would do well, when they face modeling uncertainties and have not checked the models against field data, to avoid misleading claims that they have "verified" or "validated" mathematical models at Yucca Mountain or anywhere else (33). Such terms suggest a level of reliability and predictive power which, in the face of many cases of modeling uncertainty, is impossible in practice. Instead, assessors might do better to speak in terms of *probabilities* that a given model or hypothesis has been *confirmed* and to avoid misleading claims about verification.

Statistical Uncertainty, False Positives, and False Negatives[3]

Yet another class of uncertainties scientists face is statistical. In the face of such uncertainties, they often must decide whether to limit false positives or false negatives, because they cannot do both. For example, such decisions arise because scientists performing environmental risk assessments typically face uncertainties of six orders of magnitude (34, 35). These uncertainties mean that typically we do not know, for example, whether an Indian's chance of dying in a Bhopal accident is 1 in 1 million per year or 1 in 1 per year. Our ignorance of such events is astounding and potentially catastrophic.

How ought scientists make environment-related decisions when they are ignorant of basic data and probabilities? How should they behave in a situation of statistical uncertainty? Should they assume that a particular environmental condition is safe or acceptable until it is proved unsafe or unacceptable? Or should they assume that it is unsafe or unacceptable until it is proved safe or acceptable? Where ought they to place the burden of proof? Do they place the burden of proof on polluters or on potential victims of polluters? Do they place the burden of proof on developers of the rain forest or on environmentalists who protest such development? Where they decide to place the burden of proof will determine who bears enormous risks and who receives great benefits. What is fair, equitable, and ethical in a situation of statistical uncertainty?

In this section we argue that the typical scientific norm dictating behavior under uncertainty is wrong. It is wrong to be reluctant to posit effects such as serious environmental consequences in a situation of uncertainty. Therefore, it is wrong, in a situation of uncertainty in which we cannot adequately assess effects, to place the burden of proof on possible victims of pollution or development. Instead we argue that, in situations of statistical uncertainty affecting human and environmental well-being, we should be reluctant not to posit effects such as serious harm. Therefore, in a situation of statistical uncertainty in which we cannot adequately assess effects, we should place the burden of proof on the persons who create these potentially adverse effects—that is, on polluters and developers.

[3]This discussion of statistical uncertainty is based on Shrader-Frechette (35, 15, 18) and Shrader-Frechette and McCoy (1992).

Statistical Uncertainty and the Burden of Proof

In a situation of uncertainty, errors of type I occur when one posits some possible effect and thereby rejects a null hypothesis that is true; errors of type II occur when one decides not to posit some possible effect and thereby fails to reject a null hypothesis that is false. (One null hypothesis might be, for example, "the pesticide benzene hexachloride will cause no deaths among pesticide applicators during 5 years of using 100,000 pounds per year of the chemical, provided that the applicators follow the manufacturer's instructions.") Given a situation of uncertainty about the pesticide, which is the more serious error, type I or type II? An analogous issue arises in law. Is the more serious error to acquit a guilty person or to convict an innocent person? In a situation of uncertainty, ought one to run the risk of rejecting a true null hypothesis, of not using the benzene hexachloride technology that is really acceptable and safe? Or, in a situation of uncertainty, ought one to run the risk of not rejecting a false null hypothesis, of employing the benzene hexachloride pesticide technology that is really unacceptable and unsafe? The basic problem is that to decrease type I risk might hurt the public, especially workers in developing nations where approximately 50,000 persons per year are killed by pesticides (36, 37). Yet, to decrease type II risk might hurt those who are economically dependent on this particular pesticide industry.

Pure Science, Applied Science, and Statistical Uncertainty

In the area of pure science and statistics, most persons believe that in a situation of uncertainty one ought to minimize type-I risks so as to limit false positives, assertions of effects where there are none. Pure scientists often attach a greater loss to accepting a falsehood than to failing to acknowledge a truth (35, 38). Societal decision making under uncertainty, as in cases involving sustainable energy or agricultural technologies, however, is arguably not analogous to decision making in pure science. Societal decision making involves rights, duties, and ethical consequences that affect the welfare of persons, whereas purely scientific decision making involves largely epistemologic consequences. For this reason, it is not clear that in societal cases under uncertainty, one ought to minimize type I risks. Instead, there are a number of in-principle reasons for minimizing type II errors. For one thing, it is arguably more important to protect the public from harm (from dangerous pesticides, for example) than to provide, in some positive sense, for welfare (creating jobs as pesticide applicators, for example), because protecting from harm seems to be a necessary condition for enjoying other freedoms (39, 40). Admittedly, it is difficult to draw the line between providing benefits and protecting from harm, between positive and negative laws or duties. Nevertheless, just as there is a basic distinction between welfare rights and negative rights (41), so there is an analogous distinction between welfare policies (that provide some good) and protective policies (that prohibit some infringement). Moral philosophers continue to honor related distinctions, such as that between letting a person die and killing someone. It therefore seems more important to protect citizens from public hazards, like those created by particular pesticides, than to attempt to enhance their welfare, over the short term, by implementing a potentially dangerous technology such as use of benzene hexachloride (35).

A second reason for minimizing type-II errors under uncertainty is that the public typically needs more risk protection than do the industry or government proponents of the risky technology, like particular pesticides. The public usually has fewer financial resources and less information to deal with societal hazards that affect it, and laypersons are often faced with bureaucratic denials of public danger, as in the 1973 case of the Michigan PBB contamination of cattle feed, or the 1976 dioxin poisoning at Seveso, Italy, or the 1953 Minimata poisoning in Japan. As these and other cases illustrate, public needs for protection seem larger than those of developers or manufacturers, and the importance of minimizing type-II errors appears greater than that of minimizing type-I errors (35).

Third, it is more important to minimize type-II error, especially in cases of great uncertainty, because laypersons ought to be accorded legal rights to protection against technological decisions that could threaten their health and physical security. These legal rights arise out of the considerations that everyone has both due-process rights and rights to bodily security. In cases where those responsible or liable cannot redress the harm done to others by their faulty decisions—as they often cannot in the case of dangerous technologies—there are strong arguments for minimizing the public risk. Industrial and technological decision makers cannot adequately compensate their potential victims for the bad consequences of many pesticides, for example, because the risks involve death. Therefore, the risks are what Judith Jarvis Thomson calls "incompensable" (42, p. 158). Surely incompensable risks ought to be minimized for those who fail to give free, informed consent to them. Whenever risks are incompensable (that is, imposing a significant probability of death on another), failure to minimize the risks is typically morally unjustifiable without the free, informed consent of the victim (35). And, in cases of uncertainty, it is impossible to obtain free, informed consent of potential victims because, by definition, the risks are uncertain and we have inadequate information about them.

A final reason for minimizing type II error in cases of uncertainty is that failure to do so would result in using some persons (such as pesticide applicators) as means to the ends of other persons (such as pesticide manufacturers). It would result in their bearing a significantly higher risk from toxic chemicals than other persons, despite the fact that some of those other persons (the pesticide manufacturers) have received most of the benefits associated with benzene hexachloride, for example. Such discrimination (in this case, against

pesticide applicators, particularly those in developing nations), as Frankena has pointed out, is justified only if it would work to the advantage of everyone, including those discriminated against. Any other attempt to justify discrimination fails because it would amount to sanctioning the use of some humans as means to the ends of other humans (43). Hence, in situations of uncertainty, the morally desirable position is to place the burden of proof on those who can most bear it, developers and manufacturers, rather than on the persons who are potential victims of either development or some technology.

Dealing with Statistical Uncertainty

If the arguments in this section are correct, then, in situations of statistical uncertainty, one ought not assume that potential environmental hazards are "innocent until proved guilty." Although this is the typical position adopted by most scientists and courts of law, it does not presuppose innocence. In matters of potential global harm and human catastrophe, this position places the burden of proof on those who are least able to bear it. To change this burden of proof, in cases of statistical uncertainty, environmental effects should be assumed "guilty until proved innocent." The burden of proof ought not be on the most vulnerable—potential victims of environmental hazards. Even, and especially, victims ought not bear the burden of proof.

Decision-Theoretic Uncertainty and the Maximin Rule[4]

In addition to statistical uncertainty, many scientists, especially those who work in applied areas, face decision-theoretic uncertainty. One of the crucial questions they must address is when to use a maximin decision rule in a situation of probabilistic uncertainty and when to use an expected-utility rule. This section argues that there are a number of criteria for using maximin, versus expected utility, rules in situations of scientific and decision-theoretic uncertainty. The criteria for using maximin focus on potentially catastrophic consequences and probabilistic uncertainty.

What are the consequences of using different decision-theoretic rules in situations of uncertainty? A recent U.S. government study pointed out that estimates of saccharin-caused increase in bladder cancer differed by seven orders of magnitude. Despite these uncertainties, U.S. officials have sanctioned use of saccharin, justifying their decision on the basis of a liberal risk assessment. As a result of a very conservative risk analysis, however, they banned cyclamates. Were the two risk-assessment methodologies consistent, experts have argued that cyclamates could easily have been shown to present a lower relative risk than saccharin. In Canada, for example, cyclamates are permitted and saccharin

[4]The discussion of decision-theoretic uncertainty, in this section of the chapter, relies heavily on Shrader-Frechette (35).

is banned, making Canadian regulations in this area exactly the reverse of those in the United States (35, 44–46). As the saccharin-cyclamates controversy illustrates, scientific conclusions may be uncertain, not only because of the wide range of predicted hazard values, but also because scientists use different decision-theoretic rules of evaluation. This section of the chapter will assess several of the prominent decision-theoretic rules for evaluating situations of uncertainty. To see how alternative rules can generate different conclusions consider the following case.

For U.S. reactors, the core-melt probability is about one in four for their lifetimes (47, 48). Risk assessments done by both the Ford Foundation-Mitre Corporation and by the Union of Concerned Scientists (UCS) *agree* on the probability and consequence estimates associated with the risk from commercial nuclear fission, but *disagree* in their recommendations regarding the advisability of using atomic energy to generate electricity. The UCS risk analysis decided against use of the technology; the Ford-Mitre study advised in favor of it (49–51). The two studies reached different conclusions because they used quite different decision-theoretic rules to evaluate the same data. The Ford-Mitre research was based on the widely accepted Bayesian decision criterion that it is rational to choose the action with the best expected value or utility, where "expected value" or "expected utility" is defined as the weighted sum of all possible consequences of the action and where the weights are given by the probability associated with each consequence. The UCS recommendation followed the maximin decision rule that it is rational to choose the action that avoids the worst possible consequence of all options (51). Ought we to be technocratic liberals and choose a Bayesian rule? Or ought we to be cautious conservatives and follow a maximin strategy (52)? The "prevailing opinion" among scholars, according to John Harsanyi, is to use the Bayesian rule (53, 54) even in conditions of uncertainty (55–57). This last section argues that scientists have compelling reasons for rejecting the Bayesian or utilitarian strategy when they face a situation of decision-theoretic uncertainty having potentially catastrophic consequences, and that it is often more rational to prefer the maximin strategy.

Utilitarians Versus Egalitarians

Perhaps the most famous contemporary debate over which decision rules ought to be followed in situations of risk and uncertainty is that between Harvard philosopher Rawls and Berkeley economist Harsanyi (53, 56, 57). Harsanyi believes that under conditions of uncertainty, we should maximize expected utility, where the expected utility of an act for a two-state problem is

$$u_1 p + u_2(1 - p)$$

where u_1 and u_2 are outcome utilities, p is the probability of S_1, $1 - p$ is the probability of S_2, and p represents the decision maker's own subjective probability estimate (54, 56, 58). More generally, members of the dominant Bayesian school claim that expected-utility maximization is the appropriate decision rule under uncertainty (53, 59–64). They claim that we should value outcomes, or societies, in terms of the average amounts of utility (subjective determinations of welfare) realized in them (52, 54, 56, 66–67).

Proponents of using the maximin rule (like Rawls) maintain that one ought to maximize the minimum—that is, avoid the policy having the worst possible consequences (17, 53, 54), which harms the worst-off persons (68). The obvious problem is that often the maximin and the Bayesian/utilitarian principles recommend different responses to uncertainty. To illustrate these different responses, consider an easy case involving two societies. The first consists of 1000 people, with 100 being workers (workers who are exposed to numerous occupational risks) and the rest being free to do whatever they wish. We can assume that, because of technology, the workers are easily able to provide for the needs of the rest of society. Also assume that the workers are miserable and unhappy, in part because of the work and in part because of the great risks that they face. Likewise, assume that the rest of society is quite happy, in part because they are free not to work, because they face none of the great occupational risks imposed on the 100 workers and because the nonworkers' happiness is not disturbed by any feeling of responsibility for the workers. With all these (perhaps implausible) assumptions in mind, let us suppose that, using a utility scale of 1 to 100, the workers each receive 1 unit of utility, whereas the others in society each receive 90 units each. Thus the average utility in this first society is 81.1. Now consider a second society, similar to the first, but in which, under some reasonable rotation scheme, everyone takes a turn at being a worker. In this society everyone has a utility of 35 units. Bayesian utilitarians would count the first society as more just and rational, whereas proponents of maximin and the difference principle would count the second society as more just and rational (17, 54).

Although this simplistic example is meant merely to illustrate how proponents of Bayesian utilitarianism and maximin would sanction different responses to decision-theoretic uncertainty, its specific assumptions make maximin (in this case) appear the more reasonable position. Often, however, the reverse is true. In this section we attempt to determine the better decision rule for cases of scientific decision making involving both uncertainty and great consequences to welfare (69). A reasonable way to determine whether the Bayesian/utilitarian or maximin position is superior in such cases of decision-theoretic uncertainty is to examine carefully the best contemporary defenses, respectively, of these rules. The best defenses are probably provided by Harsanyi, a utilitarian, and Rawls, an egalitarian.

Utilitarian Arguments

Harsanyi's main arguments in favor of the utilitarian, and against the maximin, strategy under decision-theoretic uncertainty are as follows: (1) Those who do not follow the utilitarian strategy are irrational and ignore probabilities. (2) They cause unacceptable ethical consequences. (3) Using the utilitarian rule, with the equiprobability assumption, promotes equal treatment.

Do Nonutilitarians Ignore Probabilities?

Choosing the maximin strategy, claims Harsanyi, is wrong because "it is extremely irrational to make your behavior wholly dependent on some highly unlikely unfavorable contingencies, regardless of how little probability you are willing to assign to them" (45, 53, p. 595). To substantiate his argument, Harsanyi gives an example of maximin decision making and alleges that it leads to paradoxes. The example is this. Suppose you live in New York City and are offered two jobs, in different cities, at the same time. The New York City job is tedious and badly paid, but the Chicago job is interesting and well paid. However, to take the Chicago job, which begins immediately, you have to take a plane, and the plane travel has a small, positive, associated probability of fatality. This means, says Harsanyi, that following the maximin principle would cause you to accept the New York job. In this example, Harsanyi assumes that your chances of dying in the near future from reasons other than a plane crash are zero. Hence, he concludes that maximin, because it directs choosing so as to avoid the worst possibility, forces one to ignore both the low probability of the plane crash and the desirability of the Chicago job and to choose the New York job. However, Harsanyi claims that a rational person, using the expected-utility criterion, would choose the Chicago job for those very reasons—its desirability and the low probability of a plane crash on the way to Chicago.

How successful is Harsanyi's first argument in employing the counterexample of the New York and Chicago jobs? For one thing, the example is highly counterintuitive; even if the example were plausible, it would prove nothing about the undesirability of using maximin in situations of *societal* risk under uncertainty, such as deciding whether to open a liquefied natural gas facility. Harsanyi makes the questionable assumption in this example that the situation of uncertainty regarding *one* individual's death, caused by the same person's decision to fly to Chicago, is no different than a situation of uncertainty regarding *many* individuals' deaths, caused by a societal decision to employ a hazardous technology.

Objecting to Harsanyi's example, John Rawls claimed that the example failed because it was of a small-scale, rather than a large-scale, situation (46, 70). My claim is similar, but more specific: situations of individual risk caused by scientific or technological decisions are *voluntarily chosen*, whereas situations

of societal risk are typically *involuntarily imposed;* hence they are not analogous. This means that, to convince us that societal decisions in situations of uncertainty are best made by following a utilitarian rule, Harsanyi cannot merely provide an example of an individual decision. In the individual case, one has the right to use expected utility so as to make efficient, economic decisions regarding oneself. In the societal case, scientists and policy makers do not always have the right to use expected utility so as to make efficient, economic decisions regarding others in society, since maximizing utility or even average utility might violate rights or duties. On the individual level, scientists' rules under uncertainty must be theoretically justifiable. On the societal level, because the rules have consequences for welfare, they must be democratically justifiable in terms of ethical procedure. Decision-theoretic rules under uncertainty require scientists to take account of the fairness of the allocational *process*, not merely the *outcomes* (71). Democratic process is probably more important in cases where probabilities are unknown than in cases of scientific uncertainty where they are certain, since it would be more difficult to ensure informed consent in the former cases. This, in turn, suggests that the individual case of decision making under uncertainty, involving pure science, requires merely a *substantive* concept of rationality. However, the societal case of decision making under uncertainty, involving applied science, requires a *procedural* or "process" concept of rationality (72, 73), because it must take account of conflicting points of view, possible consequences to welfare, as well as various ethical and legal obligations, such as those involving free, informed consent and due process. For example, if I use a decision-theoretic rule affecting my own risk, I can ask "how safe is rational enough?" and I can be termed "irrational" if I have a fear of flying. But if I use a decision-theoretic rule affecting risks to others in society, I do not have the right to ask, where their interests are concerned, "how safe is rational enough?" In the societal case, I must ask, because I am bound by moral obligation to others, "how safe is free enough?" or "how safe is fair enough?" or "how safe is voluntary enough?" (53, 55, 65, 66, 74).

When they discuss decision-theoretic uncertainty, many risk assessors, like Bruce Ames, assume that risk aversion ought to be a linear function of probability, and they criticize laypersons for being more averse to industrial chemicals than to natural toxins (like the mold in foods) that have a higher probability of causing injury or death. Invoking the concept of "relative risk," they fault laypersons for their "chemophobia," for greater aversion to lower-probability risks than to higher ones (76, 77). Probability, however, is neither the only, nor the most important, factor determining risk aversion. Risks that threaten consent, equity, or other values might also cause extreme aversion. Moreover, if subjective probabilities are frequently prone to error (78–81), then, contrary to Harsanyi's first argument, rational people might well avoid them in deciding how to handle uncertainty.

Harsanyi's first argument is also problematic because he assumes that it is irrational to base decisions on *consequences* and to ignore either a small or uncertain probability associated with them (45, 53). However, it is not irrational to avoid a possibly catastrophic risk (e.g., nuclear winter) even if it is small.

Does Maximin Lead to Unethical Consequences?

Harsanyi next claims that following maximin rules in situations of decision-theoretic uncertainty would lead to unacceptable *moral* consequences: benefiting the least-well-off individuals, even when they do not deserve it and even when doing so will not help society. To establish this point, Harsanyi gives two examples (53). In the first example, there are two patients, critically ill with pneumonia, but there is only enough antibiotic to treat one of them, one of whom has terminal cancer. Harsanyi says that Bayesians or utilitarians would give the antibiotic to the victim who did not have cancer, whereas maximin strategists would give it to the cancer victim, since he is the worse off. In the second example, there are only two citizens, one severely retarded and the other with superior mathematical ability. The problem is whether to use society's surplus money to help educate the mathematician or provide remedial training for the retarded person. The Bayesian utilitarian would spend the surplus money on the mathematician, says Harsanyi, whereas the maximin strategist would spend it on the retarded person, since he is the less well off.

The problem with Harsanyi's examples is that they are not cases of *societal* decision making under *uncertainty.* The risk is of fatality, in the pneumonia example, but one knows, with certainty, that the cancer victim is soon to die, since Harsanyi defines his state as "terminal." Likewise, in the second case, the risk is of improving the lot of two persons, one retarded and one gifted mathematically. Hence, one is not in a state of uncertainty about the probability of success in spending the monies for education in the two cases. But if so, then Harsanyi has not argued for using Bayesian/utilitarian rules under uncertainty.

A second difficulty with these examples is that Harsanyi defines the retarded person as "less well off" and therefore deserving of funds for remedial education under the maximin strategy. However, being "less well off" is not merely a matter of intelligence. It is also a matter of financial well-being and of having equal political and social opportunities. If society has given equal consideration to the needs and interests of both the mathematician and the retarded person, if the retarded person is happy and incapable of being made better off, regardless of what society spends on him, then it is not clear that he is less well off than the mathematician. If the mathematician could be made better off, with greater societal expenditures, then he may be less well off than the retarded person who has reached his potential, who is as happy as he is capable of being.

Does Using Expected Utility Treat People Equally?

Having given general, utilitarian justifications for his position, Harsanyi provides a final argument for using expected-utility rules in situations of decision-theoretic uncertainty. It focuses on what Harsanyi calls "the equiprobability assumption" (53, p. 598). Decision makers ought to subscribe to this assumption as part of the expected-utility rule, says Harsanyi, because doing so enables them to treat all individuals' a priori interests as equal (34, 53, 81) to give everyone an equal chance of being better off or worse off. In a situation of decision-theoretic uncertainty, Harsanyi claims that the rational person would always make the decision that assumes everyone's interests are equal and that yields the highest "average utility level" (52, 53, pp. 598, 67).

The most basic difficulty with the equiprobability assumption is that if there is no justification for assigning a set of probabilities, because one is in a situation of uncertainty, then there is no justification for assuming that the states are equally probable (82–84). Other difficulties are that to assign the states equal probabilities is to contradict the stipulation that the situation is one of uncertainty (54), that it is often impossible to specify a list of possible states that are mutually exclusive and exhaustive (17), and hence that different ways of defining states could conceivably result in different decision results, different accounts of how best to maximize average utility (17). The equiprobability assumption is also ethically questionable because using it does not assign equal a priori weight to every individual's interests, as Harsanyi claims. It merely *postulates* that in a situation of uncertainty, in different social systems or states of affairs, every individual has the same probability of being the best-off individual, or the second-best-off, and so on. Reality, however, is quite different from this postulate. Different states of affairs are rarely equally probable. To assume that they are, when one is in a situation of uncertainty, is problematic in part because equally probable states often affect different individuals' interests unequally.

Using *averages* also affects individuals unequally. This is why, even if one granted that it is rational to maximize expected utility in individual decisions, it would not necessarily be rational to choose the average of the expected utilities of different persons. Such a procedure would not maximize *my* expected utility, but only the average of the expected utilities of members of society (52, 67, 85). This means that the concepts of "average utility" and "equiprobability" could hide the very problems of discrimination and inequality that most need addressing. Moreover, even though the equiprobability assumption assigns every individual the *same* probability (in every state of affairs) of being the best off, second best off, and so forth, this does not guarantee that every individual's interests receive *equal* weight. Because Bayesian utilitarianism focuses on *average* utility, it dictates that decisions be made on the basis of highest average utility. This rule guarantees that the minority, with less-than-average utility, can receive a disproportionate risk burden. In such cases, one would not be treating the interests of each person in the minority as equal to those of each person in the majority. Thus, in at least one important sense, Harsanyi does not treat people the same, as he claims to do through his equiprobability assumption (53). Genuinely *equal* treatment requires that we treat people differently, so as to take account of different degrees of merit, need, rights to compensation or reparation, and so on. Treating people the same, in a situation in which existing relationships of economic and political power are already established, merely reinforces those relationships, apart from whether they are ethically defensible. Treating people the same, as most persons wish to do in situations of uncertainty, also ignores the fact that duties and obligations almost always require that people's interests *not* be treated the same. For example, suppose that Mr. X builds a pesticide manufacturing plant in Houston. Also suppose that Mr. Y, who lives next door, has demonstrably damaging health effects from the emissions of the pesticide facility. To say that Mr. X's and Mr. Y's interests in stopping the harmful emissions ought to be given the same weight is to skew the relevant ethical obligations. It would give license to anyone wishing to put others at risk for his own financial gain (67, 86, 87). Hence, there are rarely grounds for treating persons' interests the *same*, since they are almost always structured by preexisting obligations that determine whose interests ought to have more weight. This means that equity of treatment can only be achieved after *ethical analysis*, not after an appeal to treating everyone the same, in the name of the "equiprobability assumption."

Egalitarian Arguments

Admittedly, discovering difficulties with Harsanyi's arguments for Bayesian and utilitarian rules is not a sufficient condition for rejecting them. We also need to assess maximin, perhaps the best alternative rule for certain classes of cases under uncertainty. To assess this option, we evaluate Rawls' analysis. He has two main arguments to support the maximin strategy in situations of uncertainty: (1) It would lead to giving the interests of the least advantaged the highest priority. (2) The maximin strategy would avoid using a utility function, designed for risk taking, in the area of morals, where it does not belong.

Giving Priority to the Least Advantaged

Consider the first argument in favor of using maximin rules in situations of decision-theoretic uncertainty: it would lead to a concept of justice based on "the difference principle," which evaluates every possible societal or policy arrangement in terms of the interests of the least-advantaged or worst-off persons (52, 68). Rawls believes that this is an advantage of maximin, because he argues that the "first virtue" of social institutions is justice or fairness. We could arrive at just or fair social institutions, according to Rawls, if we were all rational individuals caring only about our own interests, and if we negotiated with each other (about the nature of these institutions) behind the "veil of ignorance" (i.e., without anyone knowing her own social or economic positions, special interests, talents, or abilities). Not knowing what our own situation would be, Rawls claims that we would arrange society so that even the least-well-off persons would not be seriously disadvantaged (68). This means choosing the risk distribution where the least well off are least disadvantaged (85, 88, 89).

The main objection to this argument is that we ought not use maximin because it might not increase the average utility of society, and average utility is more important than helping a subset of persons. Therefore, goes the argument, in the situation of scientific or technological decision making under uncertainty, one ought not try to protect those who are most at risk, since this would take away resources from society. Instead, one ought to use a Bayesian/utilitarian strategy to employ expected utility so as to maximize the average well-being of each member of the group (53, 90).

The main problem with this objection is that it could sanction using members of a minority who are most at risk so as to benefit the majority, namely, using some persons as means to the ends of other persons, something condemned by most moral philosophers. Presumably, however, every person ought to be treated as an end in her own right, not merely as a way to satisfy the desires of someone else, not merely as an object. Moreover, there are good grounds for believing that everyone ought to receive equal treatment, equal consideration of interests: (1) The comparison class is all humans, and all humans have the same capacity for a happy life (91). (2) Free, informed rational people would likely agree to principles of equal rights or equal protection (92, 93). (3) These principles provide the basic justifications for other important concepts of ethics and are presuppositions of all schemes involving consistency, justice, fairness, rights, and autonomy (93–98). (4) Equality of rights is presupposed by the idea of law; "law itself embodies an ideal of equal treatment for persons similarly situated" (99). If all members of society have an equal, prima facie right to life, and therefore to bodily security, as the most basic of human rights, then allowing one group of persons to be put at greater risk, without compensation and for no good reason, amounts to violating their rights to life and to bodily security. Indeed, if there were no obligation to equalize the burden of technological risk imposed on one segment of the population for the benefit of another segment, then there could be no authentic bodily security and no legal rights at all. The majority could simply do whatever they wished to any victimized minority. This is why John Rawls called his notion of justice "fairness" and why he spoke about maximin under the rubric of fairness (68, 100). Of course, sanctioning *equal* treatment, in the name of fairness, does not mean guaranteeing the *same* treatment (101). Establishing the prima facie duty to treat persons equally, so far as possible, does require that we use maximin in situations of societal risk under uncertainty (68, 102) unless we have relevant moral reasons for treating people differently (43, 103).

Efficiency, or increasing overall average utility, does not appear to provide relevant moral grounds for discrimination, especially discrimination against the least well off, for several reasons. First, discrimination against persons on grounds of efficiency is something that would have to be justified for each situation in which it occurs. The reason is that to argue (as we just have) that a principle of equal rights and equal treatment under the law is desirable, but that there may be morally relevant grounds for discrimination, is to argue for a principle of prima facie political equality (101). On this view, sameness of treatment of persons and communities needs no justification; it is presumed defensible, whereas only unequal (different) treatment requires defense (34, 96, 101). This means that the burden of proof is on the person who wishes to discriminate, who wishes not to give equal protection to some minority that is exposed to societal risk. But if the burden of proof is on the discriminator and if, by definition, we are dealing with a situation of decision making under uncertainty, then it is difficult to believe that the discriminator (the person who does not want to use maximin) could argue that efficiency provides *morally relevant* grounds for discrimination (43, 103). The reason is that the potential grounds justifying the discrimination (e.g., empirical factors about merit, compensation, or efficiency) would be, by definition, unknown in a situation of uncertainty.

Efficiency also does not appear to serve any higher interest (68, 104–108). Admittedly many risk assessors and policy makers claim that efficiency (i.e., disregarding maximin) serves the interests of everyone; they say that "the economy needs" particular hazardous technologies (75, 90, 109). They also claim that certain scientific or technological decisions (made in situations of decision-theoretic uncertainty) are not cost-effective and efficient and therefore beneficial to our national well-being (34, 48, 90, 100, 111, 112). However, for efficiency to serve the overall interest of everyone would mean that it was "required for the promotion of equality in the long run"; any other interpretation of "serving the overall interest" would be open to the charge that it was built upon using humans as means to the ends of other persons rather than treating them as ends in themselves (43, p. 15). But does efficiency per se (e.g., avoiding pollution controls and therefore equal distribution of risk) lead to the promotion of equality in the long run? The problem with answering this question in the affirmative, as Harsanyi would do, is that such an answer would contain a highly questionable *factual assumption*, that promoting technology, without also seeking equal risk distribution, will lead to greater equality of treatment in the long run. This is false. Historically, there is little basis for believing that efficiency will help promote a more equitable distribution of wealth and, therefore, more political equality (97, 107, 113, 114). In the United States, for

example, in the past 35 years, although there has been an absolute increase in the standard of living, the relative shares of U.S. wealth held by various groups have not changed. The poorest 20% of persons still receive 5% of the wealth, while the richest 20% still hold 41%; the share of the middle three quintiles has remained just as constant (105, 107, 115, 116). These data suggest that economic and technological growth, coupled with efficiency in the form of inequity of risk abatement, have not promoted economic equality. Because of the close relationship between wealth and the ability to use equal opportunities (101, 105, 106, 117–119), it is unlikely that this efficiency and economic expansion has promoted equal political treatment (48, 120, 121). If anything, it has probably made inequities even wider (107, 115, 117, 121).

Technological expansion (achieved through economic efficiency and through failure to abate technological risks) also does not ordinarily help to create a more egalitarian society because technology generally eliminates jobs; it does not create them (122). But if so, then there are not necessarily grounds for arguing that efficiency and Bayesian/utilitarian risk strategies help to equalize opportunities (101, 120). If anything, the plight of the least advantaged, whether the poor or those who bear a heavier burden of technological risk, is exacerbated by technological progress because they must compete more frantically for scarcer jobs. Moreover, because a larger portion of the indigent are unemployable, progress makes little immediate impact on the problem of hardcore poverty (121). Scientific and technological progress, without a commitment to equal distribution of societal risks, typically fails to remove distributive inequities because the poor usually bear the brunt of technological hazards. Most environmental policies, including risk policies, "distribute the costs of controls in a regressive pattern while providing disproportionate benefits for the educated and wealthy, who can better afford to indulge an acquired taste for environmental quality [and risk mitigation]" (123, p. 274; 124, 125). This means that, for the poor, whatever risk abatement and environmental quality cannot be paid for cannot be had. For example, a number of studies have shown that "those square miles populated by nonwhites and by all low socioeconomic groups were the areas of highest pollution levels" (126–132). In fact, various adverse environmental impacts, like higher risk burdens, are visited disproportionately upon the poor, while the rich receive the bulk of the benefits (52, 120, 123, 133). This all suggests that Bayesian/utilitarian strategies, in allowing the poor (persons who are least advantaged economically and therefore most helpless politically) to be further burdened with disproportionate technological risks, are especially questionable.

Do Egalitarians Avoid Utility Functions?

What about another argument of maximin proponents, that maximin would avoid using a von Neumann–Morgenstern utility function, designed for risk taking, in the area of morals, where it does not belong? This argument is that utility functions express the subjective importance people *do attribute* to their needs and interests, not the importance that they *ought to attribute*. Harsanyi wishes to make moral judgments on the basis of subjective utility functions rather than on the basis of unchanging moral principles, such as "grant equal justice to equal beings." For him, weighting the subjective importance attached to things is more important than guaranteeing adherence to moral principles, because people's preferences are *different*. But if people's preferences are different, then their utility functions may operate according to different psychological laws. But this conclusion contradicts two of Harsanyi's claims: (1) that "preferences and utility functions of all human individuals are governed by the same basic psychological laws" (53, p. 602); (2) that interpersonal utility comparisons are theoretically capable of being specified completely because they "have a completely specific theoretical meaning" (53, p. 602).

If the reasoning in the previous arguments is correct, then Harsanyi cannot coherently claim *both* that (A) preferences are needed as measures of welfare, because people's preferences/utility functions are *different*, *and* that (B) interpersonal comparisons of utility are possible because people's utility functions "are governed by the *same* basic psychological laws" (53, p. 602).

Conclusion

Because all four classes of uncertainty—framing uncertainty, modeling uncertainty, statistical uncertainty, and decision-theoretic uncertainty—are common to many environmental sciences, arriving at our four sets of rules for dealing with these cases should be useful to a variety of scientists. (1) In cases of framing uncertainty, scientists ought not make question-begging use of two-valued frames for assessing hypothesis suitability in cases of radically incomplete data. (2) In cases of modeling uncertainty, scientists ought not claim to have verified or validated their results when they have merely determined their consistency with computer models. (3) In cases of statistical uncertainty in which they are forced to choose between maximizing false positives (type I errors) or false negatives (type II errors), when scientists are faced with potentially catastrophic consequences affecting welfare, they ought to maximize false positives. (4) In cases of decision-theoretic uncertainty involving potentially catastrophic consequences, scientists ought to use maximin rather than expected-utility rules.

This chapter has argued that many decisions in situations of scientific uncertainty have been inappropriate in precisely the four ways just outlined. They have overestimated the epistemologic errors likely to result in bad science and underestimated the ethical errors likely to result in bad science policy. Doing science well thus requires us to understand the environmental contexts of its applications.

References

1. Newton I. (A Motte trans.). Philosophiae naturalis principia. New York: Greenwood Press, 1969.
2. Aristotle. (W Ross trans.). De Caelo. Oxford: Oxford University Press, 1928.
3. Minsky M. A framework for representing knowledge. In: Haugeland J, ed. Mind design. Cambridge, MA: MIT Press, 1981:95–128.
4. Fetzer JH. Artificial intelligence. Dordrecht: Kluwer, 1990.
5. Popper K. Conjectures and refutations. New York: Harper and Row, 1968.
6. Millikan R. The electron. Chicago: University of Chicago Press, 1917.
7. Frank P. Philosophy of science. Englewood Cliffs, NJ: Prentice-Hall, 1962.
8. Duhem P. (P Wiener trans.). The aim and structure of physical theory. Princeton, NJ: Princeton University Press, 1982.
9. Carnap R. Philosophical foundations of physics. New York: Basic Books, 1966.
10. Cranor C. Regulating toxic substances. New York: Oxford University Press, 1993.
11. Younker JL, et al. Report of early site-suitability evaluation of the potential repository site at Yucca Mountain, Nevada. SAIC-91/8000. Washington, DC: US Department of Energy, 1992.
12. Younker JL, et al. Report of the peer review panel on the early site suitability evaluation of the potential repository site at Yucca Mountain, Nevada. SAIC-91/8001. Washington, DC: US Department of Energy, 1992.
13. National Research Council. Risk assessment in the federal government. Washington, DC: National Academy Press, 1983.
14. National Research Council. Pharmacokinetics in risk assessment. Washington, DC: National Academy Press, 1987.
15. Shrader-Frechette K. Burying uncertainty: Geological disposal and the risks of nuclear waste. Berkeley: University of California Press, 1993.
16. Slovic P, Flynn J, Layman M. Perceived risk, trust, and the politics of nuclear waste. Science 1991;254:1603–1604.
17. Luce RD, Raiffa H. Games and decisions. New York: Wiley, 1957.
18. Shrader-Frechette, K. Environmental risk assessment and the frame problem. BioScience 1994;44:8:548–551.
19. Loux R. State of Nevada comments on the US Department of Energy site characterization plan, Yucca Mountain site, Nevada. Carson City, NV: Nevada Agency for Nuclear Projects, 1989.
20. Sinnock S, et al. Preliminary estimates of groundwater travel time and radionuclide transport at the Yucca Mountain repository site. SAND85-2701. Albuquerque, NM: Sandia National Labs, 1986.
21. Costin L, Bauer S. Thermal and mechanical codes first benchmark exercise, Part I: Thermal analysis. SAND88-1221 UC-814. Albuquerque, NM: Sandia National Labs, 1990.
22. Hayden N. Benchmarking: NNMSI flow and transport codes: Cove 1 results. SAND84-0996. Albuquerque, NM: Sandia National Labs, 1985.
23. Dijstra E. A discipline of programming. Englewood Cliffs, NJ: Prentice-Hall, 1976.
24. Hoare C. An axiomatic basis for computer programming. Comm ACM 1969;12:576–580.
25. Hoare C. Mathematics of programming. BYTE 1986 August:115–149.
26. Fetzer JH. Program verification: The very idea. Comm ACM 1988;31:9:1048–1063.
27. Fetzer JH. Mathematical proofs of computer system correctness. N Am Math Soc 1989; 36:10:1352–1353.
28. Office of Civilian Radioactive Waste Management. Performance assessment strategy plan for the geologic repository program. DOE/RW-0266P. Washington, DC: Government Printing Office, 1990.
29. Fetzer JH. Another point of view. Comm ACM 1989;32:8:921.
30. Savitzky S. Letters. Comm ACM 1989;32:3:377.
31. Nelson D. Letters. Comm ACM 1989;32:7:792.

32. Dobson J, Randell B. Program verification. Comm ACM 1989;32:4:422.
33. Hopkins P. Cone 2A Benchmarking calculations using LLUVIA. SAND88-2511-UC-814. Albuquerque, NM: Sandia National Labs, 1990.
34. Cox L, Ricci P. Legal and philosophical aspects of risk analysis. In: Paustenbach D, ed. The risk assessment of environmental and human health hazards. New York: Wiley, 1989: 1025–1038.
35. Shrader-Frechette K. Risk and rationality: philosophical arguments for populist reforms. Berkeley: University of California Press, 1991.
36. Mathews JT. World resources 1986. New York: Basic Books, 1986.
37. Repetto R. Paying the price: pesticide subsidies in developing countries. Washington, DC: World Resources Institute, 1985.
38. Axinn S. The fallacy of the single risk. Philos Sci 1966;33:1/2:154–162.
39. Shue H. Exporting hazards. In: Brown P, Shue H, eds. Boundaries: national autonomy and its limits. Totowa, NJ: Rowman and Littlefield, 1981:107–145.
40. Lichtenberg J. National boundaries and moral boundaries. In: Brown P, Shue H, eds. Boundaries: national autonomy and its limits. Totowa, NJ: Rowman and Littlefield, 1981: 79–100.
41. Becker L. Rights. In: Becker L, Kipnis K, eds. Property. Englewood Cliffs, NJ: Prentice-Hall, 1984:76–77.
42. Thomson JJ. Rights, restitution, and risk: essays, in moral theory. Cambridge, MA: Harvard University Press, 1986.
43. Frankena WK. The concept of social justice. In: Brandt R, ed. Social justice. Englewood Cliffs, NJ: Prentice-Hall, 1962:10–14.
44. Ames B, Magaw R, Gold L. Ranking possible carcinogens. In: Paustenbach DJ, ed. The risk assessment of environmental and human health hazards. New York: Wiley, 1989: 1090–1093.
45. Whipple C. Nonpessimistic risk assessment. In: Paustenbach DJ, ed. The risk assessment of environmental and human health hazards. New York: Wiley, 1989:1109–1110.
46. MacLean D. Introduction. In: MacLean D, ed. Values at risk. Totowa, NJ: Rowman and Littlefield, 1986:1–2.
47. US Nuclear Regulatory Commission. Reactor safety study: an assessment of accident risks in US commercial nuclear power plants. NUREG-75/014, WASH 1400. Washington, DC: US Government Printing Office, 1975.
48. Shrader-Frechette K. Nuclear power and public policy. Boston: Reidel, 1983.
49. Union of Concerned Scientists. The risks of nuclear power reactors: a review of the NRC reactor safety study, WASH 1400. Cambridge, MA: Union of Concerned Scientists, 1977.
50. Nuclear Energy Policy Study Group. Nuclear power: issues and choices. Cambridge, MA: Ballinger, 1977.
51. Cooke RM. Risk assessment and rational decision. Dialec 36;4:334–335. 1982.
52. Samuels S. The arrogance of intellectual power. In: Woodhead A, Bender M, Leonard R, eds. Phenotypic variation in populations. New York: Plenum, 1988:113–120.
53. Harsanyi J. Can the maximin principle serve as a basis for morality? A critique of John Rawls's theory. Am Polit Sci R 1975;69:2:596–606.
54. Resnick M. Choices. Minneapolis: University of Minnesota Press, 1987.
55. Harsanyi J. Advances in understanding rational behavior. In: Elster J, ed. Rational choice. New York: New York University Press, 1986:88–89.
56. Harsanyi J. Understanding rational behavior. In: Butts RE, Hintikka J, eds. Foundational problems in the special sciences. Vol 2. Boston: Reidel, 1977:320–323.
57. Tversky A, Kahneman D. The framing of decisions and the psychology of choice. In: Elster J, ed. Rational choice. New York: New York University Press, 1986:123–125.
58. Otway H, Peltu M. Regulating industrial risks. London: Butterworths, 1985.
59. Davis R. Introduction. In: Thrall R, Coombs C, Davis R, eds. Decision processes. London: Wiley, 1954:1–14.
60. Ellsworth L. Decision-theoretic analysis of Rawls' original position. In: Hooker CA, Leach

JJ, McClennen EF, eds. Foundations and applications of decision theory. Dordrecht: Reidel, 1978:39–40.
61. Jeffrey RC. The logic of decision. Chicago: University of Chicago Press, 1983.
62. Marschak J. Towards an economic theory of organization and information. In: Thrall R, Coombs C, Davis R, eds. Decision processes. London: Wiley, 1954:187–197.
63. McClennen EF. The minimax theory and expected utility reasoning. In: Hooker CA, Leach JJ, McClennen EF, eds. Foundations and applications of decision theory. Dordrecht: Reidel, 1978:337–338.
64. Milnor J. Games against nature. In: Thrall R, Coombs C, Davis R, eds. Decision processes. London: Wiley, 1954:49–50.
65. Harsanyi J. Understanding rational behavior. In: Butts RE, Hintikka J, eds. Foundational problems in the special sciences. Vol 2. Boston: Reidel, 1977:320–323.
66. Harsanyi J. On the rationale of the Bayesian approach. In: Butts RE, Hintikka J, eds. Foundational problems in the special sciences. Vol 2. Boston: Reidel, 1977:382–384.
67. MacLean D. Social values and the distribution of risk. In: MacLean D, ed. Values at risk. Totowa, NJ: Rowman and Littlefield, 1986.
68. Rawls J. A theory of justice. Cambridge: Harvard University Press, 1971.
69. Sen A. Rawls versus Bentham: An axiomatic examination of the pure distribution problem. In: Daniels N, ed. Reading Rawls. New York: Basic Books, 1981:283–292.
70. Rawls J. Some reasons for the maximin criterion. Am Econ Rev 1974;64:1:141–146.
71. MacLean D. Risk and consent. In: MacLean D, ed. Values at risk. Totowa: Rowman and Littlefield, NJ, 1986:17–30.
72. Andrews RN. Environmental impact assessment and risk assessment. In: Wathern P, ed. Environmental impact assessment. London: Unwin Hyman, 1988:85–97.
73. March JG. Bounded rationality. In: Elster J, ed. Rational choice. New York: New York University Press, 1986:148–149.
74. Clifford WK. Lectures and essays. London: Macmillan, 1986.
75. Michalos AC. Foundations of decisionmaking. Ottowa: Canadian Library of Philosophy, 1987.
76. Ames B, Magaw R, Gold L. Ranking possible carcinogenic hazards. Science 1987;236: 271–280.
77. Shrader-Frechette K. Risk analysis and scientific method. Boston: Reidel, 1985.
78. Hacking I. Culpable ignorance of interference effects. In: MacLean D, ed. Values at risk. Totowa, NJ: Rowman and Littlefield, 1986:136–154.
79. Kahneman D, Tversky A. Subjective probability. In: Kahneman D, Tversky A, Slovic P, eds. Judgment under uncertainty: heuristics and biases. Cambridge: Cambridge University Press, 1981:32–46.
80. Kahneman D, Tversky A. On the psychology of prediction. In: Kahneman D, Tversky A, Slovic P, eds. Judgment under uncertainty: heuristics and biases. Cambridge: Cambridge University Press, 1981:68–69.
81. MacLean D. Philosophical issues for centralized decisions. In: MacLean D, ed. Values at risk. Totowa, NJ: Rowman and Littlefield, 1986:17–30.
82. Smith VK. Benefit analysis for natural hazards. Risk Anal 1986;6:3:325–326.
83. Levi I. Newcomb's many problems. In: Hooker CA, Leach JJ, McClennen EF, eds. Foundations and applications of decision theory. Vol 1. Dordrecht: Reidel, 1978:369–383.
84. Vail S. Alternative calculi of subjective probabilities. In: Thrall R, Coombs C, Davis R, eds. Decision processes. London: Wiley, 1954:87–88.
85. Gauthier D. The social contract. In: Hooker CA, Leach JJ, McClennen EF, eds. Foundations and applications of decision theory. Vol 2. Dordrecht: Reidel, 1978:47–53.
86. Borch K. Ethics, institutions, and optimality. In: Gottinger HW, Leinfellner W, eds. Decision theory and social ethics. Dordrecht: Reidel, 1978:242–243.
87. Gersuny C. Work hazards and industrial conflict. Hanover, NH: University of New England Press, 1981.
88. Gardenfors P. Fairness without interpersonal comparisons. Theoria 1978;44:part 2:57–58.

89. Zagare F. Game theory. Beverly Hills: Sage, 1984.
90. Lave L, Leonard B. Regulating coke oven emissions. In: Paustenbach DJ, ed. The risk assessment of environmental and human health hazards. New York: Wiley, 1989:1064–1081.
91. Blackstone WT. On the meaning and justification of the equality principle. In: Blackstone WT, ed. The concept of equality. Minneapolis: Burgess, 1969:121–122.
92. Baier A. Poisoning the wells. In: MacLean D, ed. Values at risk. Totowa, NJ: Rowman and Littlefield, 1986:49–74.
93. Rawls J. Justice as fairness. In: Feinberg J, Gross H, eds. Philosophy of law. Encino, CA: Dickenson, 1975:284–285.
94. Beardsley MC. Equality and obedience to law. In: Hook S, ed. Law and philosophy. New York: New York University Press, 1964:35–193.
95. Berlin I. Equality. In: Hook S, ed. Law and philosophy. New York: New York University Press, 1964:33–34.
96. Frankena WK. Some beliefs about justice. In: Feinberg J, Gross H, eds. Philosophy of law. Encino, CA: Dickenson, 1975:250–251.
97. Markovic M. The relationship between equality and local autonomy. In: Feinberg W, ed. Equality and social policy. Urbana: University of Illinois, 1978:93–94.
98. Vlastos G. Justice and equality. In: Brandt RB, ed. Social justice. Englewood Cliffs, NJ: Prentice-Hall, 1962:50–56.
99. Pennock JR. Introduction. In: Pennock JR, Chapman JW, eds. The limits of the law. Nomos 15. The Yearbook of the American Society for Political and Legal Philosophy. New York: Lieber-Atherton, 1974:2–6.
100. Kates R, et al., eds. Hazards: technology and fairness. Washington, DC: National Academy Press, 1986.
101. Shrader-Frechette K. Science policy, ethics, and economic methodology. Boston: Reidel, 1985.
102. Sen A. Welfare inequalities and Rawlsian axiomatics. In: Butts RE, Hintikka J, eds. Foundational problems in the special sciences. Boston: Reidel, 1977:271–288.
103. Taylor R. Justice and the common good. In: Blackstone WT, ed. The concept of equality. Minneapolis: Burgess, 1969:94–97.
104. Benn SI. Egalitarianism and the equal consideration of interests. In: Pennock JR, Chapman JW, eds. Equality. Nomos 9, Yearbook of the American Society for Political and Legal Philosophy. New York: Atherton Press, 1967:75–76.
105. Patterson O. Inequality, freedom and the equal opportunity doctrine. In: Feinberg W, ed. Equality and social policy. Urbana: University of Illinois Press, 1978:31–32.
106. Pennock JR. Democratic political theory. Princeton: Princeton University Press, 1979.
107. Rees J. Equality. New York: Praeger, 1971.
108. Stewart RB. Paradoxes of liberty, integrity, and fraternity: the collective nature of environmental quality and judicial review of administrative action. Environ Law 1977;7:3:474–476.
109. Bethe H. The necessity of fission power. Sci Am 1976;234:1:26–27.
110. Drucker P. Saving the crusade. In: Shrader-Frechette K, ed. Environmental ethics. Pacific Grove, CA: Boxwood Press, 1980:200–207.
111. Maddox J. The doomsday syndrome. London: Macmillan, 1972.
112. Maxey MM. Radwastes and public ethics. Health Phys 1978;34:2:129–135.
113. Gans HJ. The costs of inequality. In: Mooney M, Stuber F, eds. Small comforts for hard times. New York: Columbia University Press, 1977:50–51.
114. Laski H. Liberty and equality. In: Blackstone WT, ed. The concept of equality. Minneapolis: Burgess, 1969:170–173.
115. Larkin A. The ethical problem of economic growth vs. environmental degradation. In: Shrader-Frechette K, ed. Environmental ethics. Pacific Grove, CA: Boxwood Press, 1980: 208–220.
116. North DC, Miller RL. The economics of public issues. New York: Harper and Row, 1971.

117. Plamenatz JP. Equality of opportunity. In: Blackstone WT, ed. The concept of equality. Minneapolis: Burgess, 1969:88–89.
118. Scharr JH. Equality of opportunity and beyond. In: Pennock JR, Chapman JW, eds. Equality. Nomos 9. Yearbook of the American Society for Political and Legal Philosophy. New York: Atherton Press, 1967.
119. Williams B. The idea of equality. In: Blackstone WT, ed. The concept of equality. Minneapolis: Burgess, 1969:49–53.
120. Gibbard A. Risk and value. In: MacLean D, ed. Values at risk. Totowa, NJ: Rowman and Littlefield, 1986:97–99.
121. Mishan EJ. 21 popular economic fallacies. New York: Praeger, 1969.
122. Grossman R, Daneker G. Jobs and energy. Washington, DC: Environmentalists for Full Employment, 1977.
123. Freeman AM. Distribution of environmental quality. In: Kneese AV, Bower BT, eds. Environmental quality analysis. Baltimore, MD: Johns Hopkins, 1972:271–275.
124. Kneese AV, Schultze CL. Pollution, prices and public policy. Washington, DC: Brookings Institution, 1975.
125. Stewart RB. Pyramids of sacrifice? Problems of federalism in mandating state implementation of natural environmental policy. In: Strom FA, ed. Land use and environment law review—1978. New York: Clark Boardman, 1978:172–173.
126. Asch P, Seneca JJ. Some evidence on the distribution of air quality. Land Econ 1978;54:3:278–297.
127. Brodine V. A special burden. Envir 1971;13:2:24–25.
128. Dane DN. Bad air for children. Envir 1976;18:9:26–34.
129. Freeman AM. Income distribution and environmental quality. In: Enthoven AC, Freeman AM, eds. Pollution, resources, and the environment. New York: WW Norton, 1973:101–102.
130. Freeman AM, Haveman RH, Kneese AV. The economics of environmental policy. New York: Wiley, 1973.
131. Kneese AV. Economics and the quality of the environment. In: Enthoven AC, Freeman AM, eds. Pollution, resources, and the environment. New York: WW Norton, 1973:74–79.
132. Ramsey DD. A note on air pollution, property values, and fiscal variables. Land Econ 1976;52:2:230–234.
133. Stein J. Water for the wealthy. Envir 1977;19:4:6–14.

Further Reading

Oreskes N, Shrader-Frechette K, Belitz K. Verification, validation, and confirmation of numerical models in the earth sciences. Science 1994;263:641–646.
Shrader-Frechette K, McCoy E. Statistics, costs, and rationality in ecological inference. Tr Ecol Evo 1992;7(3):96–99.