## RISK EVALUATION: LEGAL REQUIREMENTS, CONCEPTUAL FOUNDATIONS, AND PRACTICAL EXPERIENCES IN THE UNITED STATES

# **Summary**

**Dale Hattis** 

William S. Minkowitz

Center for Technology, Environment, and Development (CENTED) George Perkins Marsh Institute, Clark University 950 Main Street Worcester, Mass 01610

<sup>\*</sup>Presented at the International Workshop on Risk Evaluation, Schloss Haegerloch, Germany, June 20-21, 1994. Proceedings in preparation under the working title, <u>What Does "Acceptable</u>"

<u>Risk" Mean? Theory and Practice of Risk Management in Six Countries</u>, O. Renn, G. Pfister, eds., Kluwer Academic Publishers, 1996.

The United States has substantially different institutional arrangements for managing health and safety risks than the rest of the world. This paper tells the story of how our distinctive approaches for risk evaluation and control have developed. We then propose some directions for the future--in the form of criteria that embody different values rooted in our heritage.

The War of Independence that separated us from rule by the British King and Parliament directly led to a system in which governmental power was subject to extensive "checks and balances." Authority was divided among three distinct and nearly coequal branches of government which, to this day, frequently overturn one another's decisions. Contentiousness and obstruction were deliberately built in as a bulwark against tyranny by either an authoritarian executive or a temporarily popular majority in the legislature.

Many other countries undoubtedly find this both unfamiliar and unattractive. We nevertheless think there is something of value that the rest of the world can draw from our distinctive experience. In part because of the critical challenge and public discussion our system produces, risk control policies in the U.S. often reflect greater input from a wider circle of citizens, interest groups, and technical specialists than is usually mobilized elsewhere. Moreover because of the need to describe and defend official regulations in court, our administrative agencies have created a substantial and accessible record of the bases of their decisions.\*

The need to formally evaluate the consequences of proposed governmental actions has been a major motivation for the development of quantitative risk assessment techniques. These techniques are not universally understood or admired, but they reflect attempts by agencies and others to make at least some very approximate quantitative estimates of the potential benefits of regulatory measures to protect public health. Through this means, agencies responsible for health protection have attempted to respond to challenges to the merits of alternative uses of societal attention and resources. To promote even-handed evaluations, there are strong attempts in the U.S. to limit the role of the experts assessing risks as much as possible to factual, rather than policy/value questions.\*\* The expert is idealized as a technical assistant to decision making by citizens and citizen representatives, not someone to be entrusted with autocratic power him- or herself.

<sup>\*</sup> See Figure 1 on page 6 of the main text.

<sup>\*\*</sup> This has benefits in fostering integrity in risk analysis, but it also carries difficulties. Both the structuring and the reporting of risk assessment results require pragmatic decisions about what features of a risk are likely to be relevant for decision-making and therefore merit analysis and reporting. Such decisions inevitably involve some policy and value issues.

Generating and making use of these risk numbers requires public representatives to confront painful choices about risk control options and ethical and social policies for additional control and/or acceptance of remaining risks. Out of the history of these confrontations, we have drawn four broad categories of risk control criteria:

- Fair process, open disclosure and, to the extent practicable, voluntary acceptance of risk;
- Equity (fairness) in the distribution of risks in relation to the benefits derived from accepting those risks;
- The greatest possible effectiveness of governmental agencies in using limited resources to achieve health and safety goals ("Do the very best you can"); and
- The principle of "first, do no harm" from medical ethics.

These are not mutually exclusive categories, and we make no attempt to weigh their relative importance. We do believe that, taken together, they form a useful set of dimensions for assessing and comparing various available risk management options. Greater collective wisdom on risk control choices may emerge as experience with such comparative assessments accumulates.

# Fair Process, Open Disclosure and, to the Extent Practicable, Voluntary Acceptance

A principal purpose of risk control statutes in our system is the resolution of disputes over the unfair imposition of risks by some groups on others. Risks may be imposed by either private parties or government agencies.

The generator of a hazard is usually in a privileged position with respect to information about the risk. Because of this, the most basic fair-process requirement is disclosure. Disclosure requirements as a class tend to be among the least intrusive of governmental actions and have been relatively common during conservative political periods. Since the late 1960's the Freedom of Information Act has provided broad citizen access to information of all kinds in the hands of the Federal government, although there are important exceptions covering trade secrets, etc.

Beyond disclosure, fair-process requirements depend on the decision making context. A high value is placed on giving individuals not only information, but the practical capability to decline to accept exposure to specific risks if they judge those risks to be unreasonable in relation to the benefits for them.

However all individuals cannot be given a veto on all activities that expose them to any level of risk. Some risk decisions must be made collectively. In the U.S., this usually means placing decisions in the hands of representative governmental authorities. In cases of pressing

necessity (e.g., wartime), it may be critical to society to impose very large risks involuntarily on individuals, and there is a corresponding individual obligation to accept those risks under such circumstances. But individuals may justly object to being so "drafted" in less extreme cases.

Unfortunately, there is as yet no clear theory for deciding exactly where to draw the boundary between individual and collective modes of decision making about risk. In practice, the detailed requirements for fair-process are defined at a variety of governmental levels in a myriad of contexts depending on the nature of the interested parties and the prominence of risk issues in the specific decision making field. Those entrusted with risk-related decisions include highway siting and construction authorities; local land use zoning boards; evaluators of the efficacy and safety of new drugs; licensing authorities for nuclear power plants; issuers of permits for air and water pollution emissions, etc.

Historically, significant problems with fair-process in risk control have arisen when a decision making body was perceived as having a conflict of interest in making risk decisions. As a result, the tendency in recent decades has been to separate technology-promotion from technology-regulation in the activities of governmental agencies. In addition, conflict of interest rules have been extended to a variety of specialized technical advisory groups.

A final backstop to all of these procedural rights is the U.S. system of tort liability law. People who can show by a "preponderance of the evidence" that they were injured through the fault of some other party can sometimes recover for their injuries via lawsuits. The climate for such suits is fostered by the system of compensation for legal fees. In general, the plaintiffs' lawyers are paid solely from the proceeds of successful litigation; unsuccessful plaintiffs usually do not have the added burden of large legal expenses.

#### **Equity (Fairness) in the Distribution of Risks**

As mentioned earlier, a major driving force for risk control legislation has been public "outrage" over the unfair imposition of risks on some people for the primary benefit of other people. This is a very different concept than the traditional welfare economist's concern to improve societal aggregate allocative efficiency by reducing "externalities." Extreme cases of non-consensual, secret and/or fraudulent\* imposition of risks can be considered to be a kind of theft. They are at least as worthy of societal suppression efforts as common burglary.

However, such qualitative judgments of inequity beg some quantitative questions for public policy. When does a hazard impose a large enough burden of risk on an individual or a community that state action is required to correct the inequity? When should measures to abate a

<sup>\*</sup> There have been unfortunate cases where parties responsible for a risk have not just kept a hazard secret, but have actively sought to deceive those upon whom the risk was imposed.

hazard be considered sufficient? How do we judge "How clean is clean?" in each specific social context?

A substantial portion of the prior literature on risk management is devoted to the concept of "significant risk" or, alternatively, "acceptable risk." Using quantitative risk estimates as a series of rough benchmarks, various agencies and authors have attempted to define the boundaries between (a) risks that are so large that action is mandatory, (b) risks that are large enough to be considered for preventive action but where action may not always be warranted, and (c) risks that are small enough that preventive action is generally not worthwhile. Where these boundaries should be set differs for different circumstances. Among the factors that should be considered are the degree of voluntary choice available to the risk bearers and the costs and hazards of the measures available to reduce the risk.

Other relevant factors include the confidence with which we know the risk, and the degree to which the risk is different for different people. In recent years the technical risk analysis community has become increasingly dissatisfied with the "single number" characterizations of risks that form the basis of the traditional "significant risk" judgments. A single number, (e.g., a 1/1,000,000 lifetime risk of death from a particular cause) often hides more than it reveals, and what is hidden is subject to manipulation. Risk assessors have moved toward more distribution-based approaches for expressing risks. Unfortunately, the risk management literature has not yet made use of these more sophisticated ideas and tools.

We believe that efforts to evaluate and compare risks can be improved by distinguishing two dimensions along which risks vary:

- <u>Interindividual variability</u> means real differences among people in the risk from a given hazard. These differences result from both differences in susceptibility and differences in exposure-related factors (e.g., breathing rates, dietary habits, location and duration of residence, occupation, etc.),
- <u>Uncertainty</u>--the imperfection of our knowledge about the magnitude of risk to the population as a whole or to any specific segment of the exposed population.

In general, uncertainty can be reduced by obtaining better information. Variability, however, cannot be reduced, although it may be more precisely defined by improved data and models. Other things being equal, large amounts of real variability (especially in exposure) mean that there are opportunities for more efficient targeting of societal resources to reduce risk where it is most intense. On the other hand, large amounts of uncertainty mean that there may be opportunities to improve the information base for decision making by additional research in areas where available information is most doubtful.

Minimally, if we are to retain numerical criteria for "significant risk" as guideposts for risk management, then these numerical criteria can be more fully defined. For example, let X be

a particular level of risk, such as 1/100,000. This risk will not be uniform in the population. People's real individual risks will be spread out in a dimension we can call "Y". To be clear in formulating the social objective, the risk manager needs to decide what fraction of the population (Y) must be kept below one or more X levels of individual risk. Of course there is also uncertainty. We can ask, "In order for this risk to be acceptable, how <u>confident</u> should we be that the risk is below X for the Yth percentile person?" We call this third dimension (confidence or its opposite, uncertainty) Z.

The table on the next page gives an example of what a risk might look like in these three dimensions. For this purpose, we have done some calculations that might represent the risk distribution for a particular type of carcinogen.<sup>\*</sup> The numbers in the cells of the table represent the X (risk) values in exponential notation.<sup>\*\*</sup> The Y (variability) dimension is traced from the top to the bottom of the table. The Z (uncertainty) dimension goes from left to right in the table.

The overall starting point for the table is the cell in the upper right hand corner. This cell says that, for a median person, we are 95% confident that the risk is below about 1 in a million. This starting point arises from the fact that, although conventional carcinogenesis risk assessments are intended to be "conservative" (there is supposed to be high confidence that the numbers overstate true risks) there is usually no consideration of the variation in susceptibility among a diverse human population. Variability in exposures is often the subject of analysis, but usually the assessors do not consider the implications of variability in individual biological susceptibility to the action of the carcinogen in question.\*\*\*

To the extent that risk managers wish to focus on a single point in this three-dimensional space to determine "acceptability", we would suggest an initial evaluation of the fairness of the X level of risk assessed at

• A relatively high value in the Y dimension--perhaps corresponding to the most atrisk individual expected to be present in the exposed population. The Y percentile evaluated would be larger for a larger exposed population, but in all cases the risk manager would be considering the maximum actual level of individual risk expected in the population, keeping in mind variation in both exposure and susceptibility.\*\*\*\*

<sup>\*</sup> I.e., a carcinogen that acts by reacting with DNA after being activated by metabolism in the liver.

<sup>\*\*1</sup>E-06 = 1 X  $10^{-6}$  = 1/1,000,000 lifetime risk of developing a cancer from the hypothetical agent in question. \*\*\*These results and the many underlying assumptions of the analysis are discussed in more detail on pp. 64-74. One surprising implication is that because of the positive skewness expected for both variability and uncertainty distributions (both of which are assumed to be lognormal), the mean "expected value" of this kind of risk for the average person may not be very much lower than the conventionally-assessed 95% upper confidence limit risk estimate. The lognormal shape of the distributions arises because many of the factors that contribute to both the variability and the uncertainty act multiplicatively in affecting risks.

<sup>\*\*\*\*</sup> The focus on the expected level of actual risk for the most at-risk individual has antecedents that go back at least to Rawls, whose philosophical analysis indicated that a just society will not let anyone fall below some minimal level of welfare.

• A moderately conservative value in the Z direction--usually no more than the 90th or 95th percent confidence level (see rationale below).

There are three reasons why we suggest that risk managers not attempt to focus on results at confidence levels (Z) higher than about 95%. First, we believe that, in our uncertain world, governmental administrators are not expected to be able to make decisions that are robust to greater levels of confidence than this. Second, we believe that such very-high-confidence

### **Example of a Combined Presentation of Variability and Uncertainty**

### Elaboration of Variability and Uncertainty for a Conventionally Assessed 10<sup>-6</sup>/Lifetime 95% Upper-Confidence-Limit Individual Cancer Risk for a Genetically Acting, Metabolically Activated Carcinogen--260 Million People Exposed

Uncertainty Dimens			imension (2	<u>(</u> )
Variability Dimension (Y)	50% Confidence Level	Arithmetic Mean ("Expected Value")	90% Confidenc e Level	95% Confidenc e Level
Median Sensitivity Individual (50th percentile)	7.2E-08	2.8E-07	5.3E-07	9.8E-07*
Arithmetic Mean Sensitive Individual**	1.5E-07	7.0E-07	1.3E-06	2.3E-06
1/10 Most Sensitive Individual (90th percentile)	3.4E-07	1.6E-06	2.8E-06	5.1E-06
1/20 Most Sensitive Individual (95th percentile)	5.3E-07	2.6E-06	4.7E-06	8.5E-06
1/100 Most Sensitive Individual (99th percentile)	1.2E-06	7.6E-06	1.2E-05	2.5E-05
1/1,000 Most Sensitive Individual (99.9th percentile)	3.0E-06	2.8E-05	3.9E-05	8.6E-05
1/10,000 Most Sensitive Individual (99.99th percentile)	6.3E-06	8.8E-05	1.0E-04	2.5E-04
1/Million Most Sensitive Individual (99.9999th percentile)	2.1E-05	6.7E-04	5.2E-04	1.5E-03

<sup>\*</sup> This value (approximately  $1E-06 = 10^{-6}$ ) is the starting point for the calculation of the rest of the numbers in the table.

<sup>\*\*</sup> The numbers in this row are an assessment of the uncertainty about the risk facing an individual who has no knowledge of where he or she sits in the overall distribution of sensitivities (this of course is the typical situation for individual people). The single most significant result is probably the second number in this row--the "expected value" of the risk for the average-risk person. It will be noted that this is just a little less than the "conservative" 95% upper confidence limit value of  $10^{-6}$  from which we began.

Number of People Expected to Be	0.57	2.6	4.7	8.6
Affected/Year Among 260 Million People				
Exposed (Calculated from the mean individual				
risk level listed in the second row above, and a				
standard 70-year lifetime)				

decisions are likely to be very large consumers of societal resources. Third, analysts have a tendency toward overconfidence that makes it difficult to produce good estimates that can realistically be assigned very high levels of confidence.

On the other hand, we caution against evaluating fairness with a completely "risk neutral" position by focusing solely on the "expected value" of X at the chosen Y percentile. This is because the burden of obtaining better information about the risk (thus narrowing the distribution in the Z dimension) is most fairly placed on the parties responsible for generating the risk rather than those intended to be protected by governmental intervention. Adopting a high percentile in the Z dimension creates an incentive for the responsible parties to improve the information base.

#### "Do the Very Best You Can" With Available Resources

The plain language of "do the very best you can" implies a somewhat different and, we think, less ethically objectionable decision rule than a traditional cost-benefit formulation. The use of economic, technical and management resources to abate one specific hazard may have the practical consequence that those resources are not used to abate another, perhaps similar, hazard. Recognizing those practical limits and trying to put limited resources to their most effective possible use has quite a different flavor than "trading lives for dollars" or making sure that "the benefits exceed the costs".

The agency would ask

- (1) Are we directing the attention and preventive efforts of the regulated parties to the most important hazards they have the opportunity to deal with, and
- (2) Considering the range of available opportunities for the use of our agency's specialized human standard setting and enforcement resources, does this standard (or this project) represent the very best that we can do?

Given this basic framework, priorities should be set in such a way as to maximize the "net benefit" (in whatever units this is evaluated) per unit of each separate limiting resource that a public agency has available. The table on the next page illustrates why this type of formula for the allocation of agency resources is generally superior to a societal benefit/societal cost formula for setting priorities. Agency resources have a catalytic role in bringing about desirable exchanges of societal cost for societal benefits, and need to be treated differently in a priority setting equation than general societal resources.

In this framework uncertainty and variability in priority scores have very different implications. Other things being equal,

• Greater variability (true heterogeneity in the actual results of allocating effort to different categories) will enhance the benefits of allocating resources preferentially to

relatively high-priority categories. Categories for evaluation should therefore be created which tend to maximize this variability.

## Rationale for Assessing Priorities in Terms of Net Benefit Per Unit of Agency Resources

# Consider a hypothetical choice among five different potential categories for the use of a particular limiting agency resource:

Project	% Agency Resource Required	Societal Benefits	Societal Costs	Net Societal Benefits	Societal Benefit/Cost Ratio	Net Societal Benefit/Agency Resource
А	25%	10	2	8	5.0	0.3
В	50%	1000	600	400	1.7	8.0
С	25%	100	40	60	2.5	2.4
D	50%	150	50	100	3.0	2.0
Е	25%	1200	800	400	1.5	16.0

# If we rank these options for resource expenditure in order of their societal benefit/cost ratios, and choose to undertake the highest three, we get:

Project	Agency	Societal	Societal	Net	Societal	Net Societal
	Resource	Benefits	Costs	Societal	Benefit/Cost	Benefit/Agency
	Required			Benefits	Ratio	Resource
А	25	10	2	8	5.0	0.3
D	50	150	50	100	3.0	2.0
С	25	100	40	60	2.5	2.4
Total	100	260	92	168	2.8	1.7
Chosen						
Not (	Chosen:					
В	50	1000	600	400	1.7	8.0
Е	25	1200	800	400	1.5	16.0

# By contrast, if we prioritize by societal net benefit per agency resource, and undertake the highest three as before, we get:

Project	Agency Resource	Societal Benefits	Societal Costs	Net Societal	Societal Benefit/Cost	Net Societal Benefit/Agency
	Required			Benefits	Ratio	Resource
E	25	1200	800	400	1.5	16.0
В	50	1000	600	400	1.7	8.0
С	25	100	40	60	2.5	2.4
Total	100	2300	1440	860	1.6	8.6
Chosen						
Not (	Chosen:					
D	50	150	50	100	3.0	2.0
А	25	10	2	8	5.0	0.3

Clearly, much greater net societal good is done if the second priority setting scheme is used to guide the selection of which projects receive the scarce agency resources. • By contrast, greater uncertainty (imperfection in knowledge of the actual results of allocating effort to different categories) will increase the desirability of (1) measures to obtain better information, and (2) some spreading of efforts to lower priority candidates for attention.

### The "First, Do No Harm" Principle

We borrow this last criterion from a traditional precept of medical ethics that dates back to Hippocrates. Simply put, there should be reasonable confidence that efforts undertaken in the name of risk control will not do more harm than good. Standards of due care in modern medical practice require some reasonable degree of evidence that remedies marketed for a specific medical condition are likely to have a favorable ratio of therapeutic benefits to risks for the patient. Advocates of regulatory control measures have an obligation to examine the potential consequences of their social prescriptions with enough breadth and care to be confident that they are not offering the policy equivalent of "snake oil."

### Use of the Criteria

These suggested criteria should not be seen as a kind of formula to be programmed into a computer in place of human decision making. Rather, we hope they will contribute to an evolving language. The vocabulary of this language needs to both

- accurately represent our advancing technical understanding of the facts about risks and risk control options, and
- frankly and compassionately convey our maturing understanding of the relevant value questions.

In this way, these criteria and their successors may help develop the capabilities of our democratic system to avoid the paralysis and dissembling that has sometimes characterized U.S. risk management decision making in the past.\*

<sup>\*</sup> Albert Gore, in the introduction to his book on environmental issues wrote,

<sup>&</sup>quot;The problem is not so much one of policy failures: much more worrisome are the failures of candor, evasions of responsibility, and timidity of vision that characterize too many of us in government. More than anything else, my study of the environment has led me to realize the extent to which our current public discourse is focused on the shortest of short-term values and encourages the American people to join us politicians in avoiding the most important issues and postponing the really difficult choices."

<sup>--</sup>Gore, A. Earth in the Balance Houghton Mifflin Co., Boston, 1992, p. 21.