Uncertainties in Risk Assessment

Dale Hattis[A] and John Froines[B]

[A]Center for Environment, Technology and Development, Clark University, Worcester, Massachusetts;
[B]School of Public Health, University of California at Los Angeles, Los Angeles, California

Current risk assessment practices largely reflect the need for a consistent set of relatively rapid, first-cut procedures to assess "plausible upper limits" of various risks. These practices have important roles to play in 1) screening candidate hazards for initial attention and 2) directing attention to cases where moderate-cost measures to control exposures are likely to be warranted, in the absence of further extensive (and expensive) data gathering and analysis.

A problem with the current practices, however, is that they have led assessors to do a generally poor job of analyzing and expressing uncertainties, fostering "One-Number Disease" (in which everything from one's social policy position on risk acceptance to one's technical judgment on the likelihood of different cancer dose–response relationships is rolled into a single quantity). At least for analyses that involve relatively important decisions for society (both relatively large potential health risks and relatively large potential economic costs or other disruptions), we can and should at least go one further step — and that is to assess and convey both a central tendency estimate of exposure and risk as well as our more conventional "conservative" upper-confidence-limit values.

To accomplish this, more sophisticated efforts are needed to appropriately represent the likely effects of various sources of uncertainty along the casual chain from the release of toxicants to the production of adverse effects. When the effects of individual sources of uncertainty are assessed (and any important interactions included), Monte Carlo simulation procedures can be used to produce an overall analysis of uncertainties and to highlight areas where uncertainties might be appreciably reduced by further study. Beyond the information yielded by such analyses for decision-making in a few important cases, the value of doing several exemplary risk assessments in this way is that a set of benchmarks can be defined that will help calibrate the assumptions used in the larger number of risk assessments that must be done by "default" procedures.

Introduction

To communicate with some degree of clarity on the subject of uncertainty, it is necessary to offer some basic definitions and distinctions from related concepts. In this paper, "uncertainty analysis" means an attempt to fairly assess and convey how likely it is that the estimated value of a particular parameter differs by various amounts from the "truth." A description of uncertainty, therefore, is a description of the imperfection in knowledge about something that is conceived of as having some "true" single value in some inaccessible reality. "Uncertainty analysis" must be distinguished from another concept that is also described with the aid of probability distributions. "Heterogeneity" or "interindividual variability" is the distribution of true values of a parameter that would be found in a population by perfectly accurate measurement techniques. For example, weighing a set of individuals with an excellent scale will not obtain the same results for different people. People really do differ in their weights (and other characteristics that affect individual risks) and no degree of improvement in the measuring instrument will make them all the same. The difficulties of knowing the degree of interindividual variability in susceptibility to toxicants in a population is one factor that contributes to uncertainties in assessing risks particularly for noncancer effects.[1]

First, some policy questions related to uncertainty analysis will be discussed, i.e., what is it potentially good for, and why is this such a touchy subject? This will be followed by some historical speculations on why many of those who were brought up in the 1940s to the 1960s and received technical training in the 1960s through the 1980s find the prospect of quantitatively assessing uncertainties foreign and troubling. An extended example will be offered which addresses the uncertainties in a practical risk assessment context; the example compares potential carcinogenic risks of drinking water from two sources: an "Advanced Water Treatment" system developed for the City of San Diego and a particular reservoir supplied from the Colorado River. Finally, some general caveats and warnings are offered about the "brave new world" of enhanced analysis of uncertainties in risk analysis.

Uncertainty Analysis: Policy Considerations

A good analysis of uncertainties serves policy goals of 1) "bounding the set of not clearly incorrect answers"[2] and 2) allowing the reader of a set of risk assessment results to "make as informed a decision on risk acceptance or control as if the reader him/herself had gone through the process of doing the risk assessment."[3] An appropriate analysis of uncertainties brings into the open the expected consequences

of standard summary treatments of uncertainties in different factors affecting risk, thus allowing greater scope in making risk control choices from among options that are expected to limit risk to various extents with different degrees of confidence.

The particular form of expression of uncertain results that is most relevant for policy-making depends on the legislative mandate of a particular program or type of decision. Changing the form for expressing uncertainty for specific regulatory purposes requires careful examination of whether one is changing the protective posture of the agency in implementing the intent of the law in question. Legislation that calls on an agency to "protect public health with an adequate margin of safety" may well be interpreted as directing the agency to pay attention to the level of risk that, in the judgment of agency analysts, will not be exceeded with relatively high confidence (99% or possibly more); in contrast, a statute calling for a cost/benefit balancing may well be interpreted as requiring the agency to pay attention to the mean "expected values" of the distributions of the estimates of risk and economic costs.

Ideally, an analysis of uncertainty is not an afterthought that one gets from the statistician after the rest of the work on a risk assessment has been completed. It should really be part of the "warp" and "woof" of each step in the analytical process that carries some uncertainty (i.e., everything but the arithmetic and sometimes the arithmetic, too). Thus, among the reasons this subject is sensitive are 1) it matters to the substantive protective posture for managing risks and 2) it has the potential to change the way a substantial number of technical analysts do their jobs.

Historical Attitudes Toward Uncertainties in Scientific Information

Many scientists involved in risk issues have some reluctance to attempt to quantify their uncertainties; this may derive from the cultural assumptions about science that prevailed when they were growing up and making the lifetime commitment to become scientists. Recent changes in these cultural assumptions are a likely source of a feeling among scientists that the understandings under which they chose science as a career have been altered, and not for the better.

One popular scientific attitude toward uncertainties present in the 1950s, when many of today's scientists were children, is typified by detective Joe Friday's saying, "Just the facts, ma'm, nothing but the facts." Joe Friday was not a scientist, of course, but we do not recall Mr. Wizard being uncertain about anything either. A favorite example, however, is from a classic early 1950s science fiction movie called *THEM* in which giant ants have been produced as the result of mutations of ordinary ants following exposure to radiation from the atomic bomb testing. This fact is not known at the beginning, of course. The local police and the FBI man only know that people are showing up horribly mutilated and smelling of formic acid; there is also a strange print cast from the desert sand near one of the victims that no one is able to identify. Sending the print to Washington brings, on the return flight, the stereotypical absent-minded professor who, with his lovely daughter (also a Ph.D.), proceeds to investigate. The father and daughter team assemble information but share absolutely nothing with the local police and the FBI agent who had called for help in the first place. The reason for their behavior is quickly explained. They do, indeed, have a theory. They are, in fact, specialists on ants, and formic acid stings are evidently the way ants kill. However, they are duty-bound as scientists not to share their theory until they are certain, even though giant, man-killing ants may be multiplying in the desert. Why? Because it might cause panic. One might think that at least some limited contingency planning might be done on the basis of tentative information, but apparently this was not the prescription for scientific behavior in the early 1950s.

In the 1990s, of course, we are expected to deal with hazards that are a great deal more subtle, with information that is less complete than was available in the movie about the giant ants. By today's standard of conduct, the two scientists in the movie should have said, "Incredible as it seems, there is a good chance that somehow a colony of giant ants has developed out there in the desert." Perhaps they would even state, "Based on the way that print looks and the circumstance of the formic acid, we think there may be an 80% chance that there are giant ants out there. What the other 20% possibility might be, we can't guess, but we give the ant possibility about 0.8."

On pain of usurping the autonomy of the civil authorities to exercise their authority in representing community value/policy preferences, the scientist is called upon to disclose information of potential social significance while it is still somewhat uncertain. Hopefully, this is done with as much attendant communication about the alternative possible states of the world, and the implications of these alternatives, as would be helpful in the decision-making process. The scientist must not arrogate the sole authority to make the relevant decisions. By withholding information on possible states of the world with potential implications for decision-making, the scientist would be doing just that.

A Practical Example of Uncertainty Analysis

Over the past few years, John Froines and others at UCLA have been measuring the concentrations of a number of contaminants in two potential sources of drinking water for the City of San Diego:

- The Miramar reservoir (MIRA), which contained chlorinated water derived from the Colorado River.
- The output of an advanced water treatment (AWT) system that recycles sewage.

We considered three different sources of uncertainty in our

analysis:

- Uncertainty as to which pollutants are actually present in the two water sources, after taking into account the measurement difficulties encountered in the study and the sensitivity with which different chemicals can be reliably detected.
- Uncertainty in the long-term average amounts of each carcinogenic pollutant present.
- Uncertainty in the potency of each carcinogenic pollutant for producing cancer in humans, i.e., how many cases of cancer should be expected per unit of lifetime average consumption?

The combined effects of uncertainties of the latter two types for all studied contaminants were calculated with the aid of a Monte Carlo simulation.

In a Monte Carlo simulation, the combined effects of different sources of uncertainty are assessed by randomly drawing values from distributions intended to represent different uncertain parameters that affect risk. For example, imagine that we are uncertain about both the concentration of a pollutant in water and the number of cases that might result per unit concentration (the carcinogenic potency). For each "trial" of a Monte Carlo simulation, a value is randomly selected from 1) a statistical distribution that represents our uncertainty in concentration and 2) another statistical distribution that represents our uncertainty in the cancer potency. Together with an assumption about the amount of water consumed per day or per lifetime, this leads to a prediction of the cancer risk from that chemical for that trial. By repeating this procedure for hundreds or thousands of trials, one obtains a picture of how likely it is that the cancer risk from each chemical might take on various values. Calculations of overall risk from multiple chemicals are added within each trial.

We chose not to include in our first-cut analysis one other uncertain parameter that directly bears upon the risk: the amount of water that people consume. All calculations were for the risk expected if exposed people consumed a standard 2 L of water per day for their entire lifetimes.

Which Pollutants Are Actually Present?

For the most part, the contaminants that were considered to be present in each water system were those that were reliably detected utilizing a criterion based on the standard error of the difference between the long-term average concentration in the sampled water and a set of concurrent "travel blanks." In our final analysis, we chose to depart from this procedure in one case. Arsenic, as it happens, dominates the overall risk for the MIRA water. In order to avoid overstating the difference in risk that the data could really show between MIRA and AWT water, we chose to retain arsenic in the AWT calculations. This allowed us to illustrate the level of risk that might be possible from AWT water if arsenic were, in fact, present at the levels indicated, even though those levels were not sufficient for reliable detection of arsenic in AWT water. Thus, the way we analyzed and presented uncertainties was affected by the type of intended risk comparison.

Representing Uncertainties in the Concentrations of Pollutants in the Waters

In many cases, the long-term average values of the travel blanks were similar in order of magnitude to the long-term average values of the AWT and/or MIRA water samples. Because of this, the best expected value for the long-term average concentrations had to be arrived at by subtracting the blank averages from the sample averages. When subtracting two numbers of roughly equal size, the resulting uncertainty is likely to be best described by a normal distribution. Therefore, our basic representation of uncertainties in the concentration values was a normal distribution with a standard deviation calculated from the standard deviations of the blanks and the sample means.

This needed to be modified somewhat because it made no sense to allow the distribution of concentrations to take on negative values. Therefore, in the course of the Monte Carlo simulations, wherever a negative value would otherwise be selected from the normal distribution, we instructed the computer to substitute zero.

Representing Uncertainties in the Carcinogenic Potency of Different Carcinogens

In work previously conducted for the National Institute for Occupational Safety and Health, we performed a series of three case studies (perchloroethylene, ethylene oxide, and butadiene) incorporating pharmacokinetic modeling of the delivered doses of putative, genetically acting agents or genetically active metabolites to improve the assessment of likely low-dose carcinogenic risks.[4–7] One of the innovative aspects of these studies was attempts to make "best" (or "least unlikely") estimates of risk in addition to more usual "plausible upper limit" estimates. Table I summarizes the basic approaches used to arrive at these different estimates of risk. Table II shows the basic comparison of the results of these more elaborate analyses to the results of more standard, upper-confidence-limit-only risk assessments by the U.S. Environmental Protection Agency's (EPA) Carcinogen Assessment Group.

As might have been expected, the estimates of the most likely values of the carcinogenic risk were considerably below the plausible upper limit values, although the upper confidence limit estimates of cancer potency were similar between our studies and those of EPA (Table II). The summary analysis in Table III indicates that, on average, the "least unlikely" estimates of cancer risk in these studies were about 7% of the "plausible upper limit" risk estimates. The lower portion of Table III shows how we reasoned further from this difference to obtain a likely distribution of cancer potency

TABLE I. Assumptions for Derivation of "Least Unlikely" and "Plausible Upper Limit" Cancer Potency Estimates for Putative Genetically Acting Agents

"Least Unlikely" Estimates
- "Best-estimate," physiologically based, pharmacokinetic models for estimation of metabolic activation and/or the persistence of active metabolites in the system.
- Expression of "delivered dose" as either the amount of active metabolite per (body weight)$^{3/4}$, or the internal concentration X time product of the direct-acting agent for interspecies projections.
- Calculation of "maximum likelihood estimates" (MLE) of the coefficients of the multistage model for individual tumor sites, summed for all sites. But an assumption of a modest amount of background interaction is introduced where the MLE linear terms (q_1s) for all sites are zero to produce a finite low-dose linear term.
- Calculation of the geometric mean of expected low-dose risk as predicted by data from experiments on different species and genders of animals.

"Plausible Upper Limit" Estimates
- Plausible high-risk-predicting pharmacokinetic models.
- Expression of delivered dose per (body weight)$^{2/3}$ for interspecies projections.
- 95% upper confidence limit estimates of the linear term from a multistage model fit in the most sensitive species and gender of animals tested.

estimates that we could use for Monte Carlo simulation analysis.

We began with a basic decision to treat our uncertainties in this parameter as lognormally distributed; i.e., the logarithms of the distribution of likely cancer potency values are normally distributed, as is illustrated in Figure 1. This was because in our judgement the different sources of likely error that contributed to our uncertainties in carcinogenic risk for individual chemicals (e.g., the relative amounts of absorption of toxicants in people and rodents; relative rates of elimination from the body; rates of metabolic activation to toxic metabolites; and inactivation to safe metabolites, cell replication, and DNA repair rates in animals and humans) all will tend to exert relatively independent, multiplicative effects on the level of likely human risk relative to the risk inferred directly from animal experiments (which are the source of most of the cancer potency estimates). The consequence of multiplying together a series of uncertain parameters is that the resulting overall uncertainty will tend to be lognormal. This is because 1) multiplying a series of uncertain parameters is the same as adding their logarithms and 2) by the "central limit theorem" of probability and statistics, the uncertainty in the sum of a large series of uncertain parameters takes on the standard "normal" or "Gaussian" form.

Given the choice of the lognormal form, we needed to determine how likely it was that the true cancer potency for a given chemical was equal to or above the EPA "plausible upper bound" cancer potency factor. It is important to understand that the EPA values, although calculated with "conservative" assumptions that are expected to overstate risk most of the time, cannot be expected to always overstate risk.

- For example, EPA routinely uses the most sensitive species tested for estimating human risk; however, in general, where only two other species are tested (rats and mice), there is nothing to prevent humans from being more sensitive than the more sensitive of the two rodents, at least for some modest proportion of all carcinogens.
- Standard risk assessment procedures that do not use

TABLE II. Comparisons of Results of Pharmacokinetic-Based Risk Analyses with EPA Projections of Low-Dose Risks (all data in lifetime risks for occupational exposure to 1 ppm, 8 hours/day, 5 days/week for 45 years)

Chemical	"Best (Least Unlikely) Estimate"	"Plausible Upper Limit Estimate"
Results from Pharmacokinetic-Based Risk Analyses:		
Ethylene oxide	0.0065	0.019
Butadiene	7.9E-4	0.032
Perchloroethylene	6.7E-4	0.013
Results from More Usual EPA/CAG Risk Analyses from Animal Data:		
Ethylene oxide	Not done*	0.028
Butadiene	Not done	0.098
Perchloroethylene	Not done	0.0033

*Implicitly, a best estimate equivalent to a lifetime risk of 0.104 was calculated from two observed human leukemias in the Hogstedt et al. (1986) study. This is a central tendency estimate because no statistical upper-confidence-limit procedure was used in computation.

TABLE III. Uncertainties in Carcinogenic Risk Estimates for Genetically Acting Agents, as Inferred from Three Case Studies of PBPK[A]-Based Risk Analyses

Compound	Hattis "Best Estimate"	Hattis "Plausible Upper Limit"	EPA UCL[B] CPF[C]	Hattis Best/ Hattis UCL Ratio	Hattis Best/ EPA UCL Ratio
Ethylene oxide	0.0065	0.019	2.80E-02	0.342	0.232
Butadiene	0.00079	0.032	9.80E-02	0.025	0.008
Perchloroethylene	0.00067	0.013	3.30E-03	0.052	0.203
			Geom. Mean	0.076	0.072
			Geom. Std. Dev.	3.881	6.703
			Geom. Std. Err.	2.188	2.999

If we take the EPA UCL estimate of risk as approximately a 95th percentile value (1.6449 standard deviations above the median), and if we represent our uncertainties as lognormally distributed about a median estimate at approximately 0.072 times the EPA UCL, then the geometric standard deviation of the lognormal distribution representing our uncertainties is $10^{[\log(0.724)/1.6449]} = 4.93$.

[A]PBPK = physiologically based pharmacokinetic model.
[B]UCL = upper confidence limit.
[C]CPF = cancer potency factor.

pharmacokinetic analysis may understate risk if, as in the case of vinyl chloride, there is a saturation at high doses in the metabolic activation of the carcinogen, leading to a plateau at high doses in the percentage of animals that develop tumors. If only the two highest dose points had been available for vinyl chloride (as would have been the case if the vinyl chloride data available for risk analysis had come solely from the usual National Toxicology Program for chronic animal bioassay), the low-dose slope of the cancer dose–response relationship would probably have been underestimated by about fivefold.

- If there are appreciable differences among humans in overall individual susceptibility, as there seem to be in many specific parameters that can be expected to affect susceptibility, this would be expected to increase the population risks of humans exposed at relatively low-dose levels relative to what would be expected for a completely uniform population of "median-susceptible" individuals.[8,9] This factor is not included in EPA's calculations, but it can routinely be expected to increase low-dose population risks to a diverse community of humans relative to the risks of the relatively uniform groups of rodents that are tested at high doses.

In the light of these various possibilities, we chose to treat the EPA upper confidence limit (UCL) cancer potency estimates derived from animal data as the 95th percentile of our distribution of uncertainties in cancer potency. As listed at the

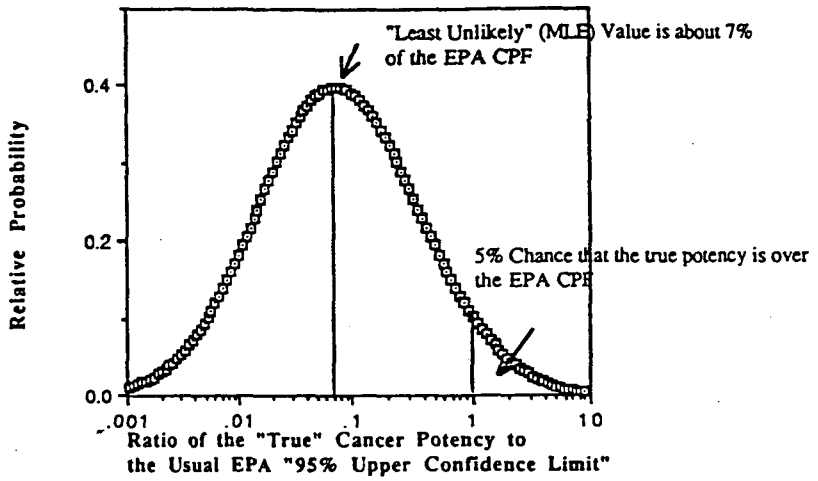

FIGURE 1. Log plot of the estimated likelihood distribution for cancer potency factors (CPF) for genetically acting carcinogens. (MLE = maximum likelihood estimate.)

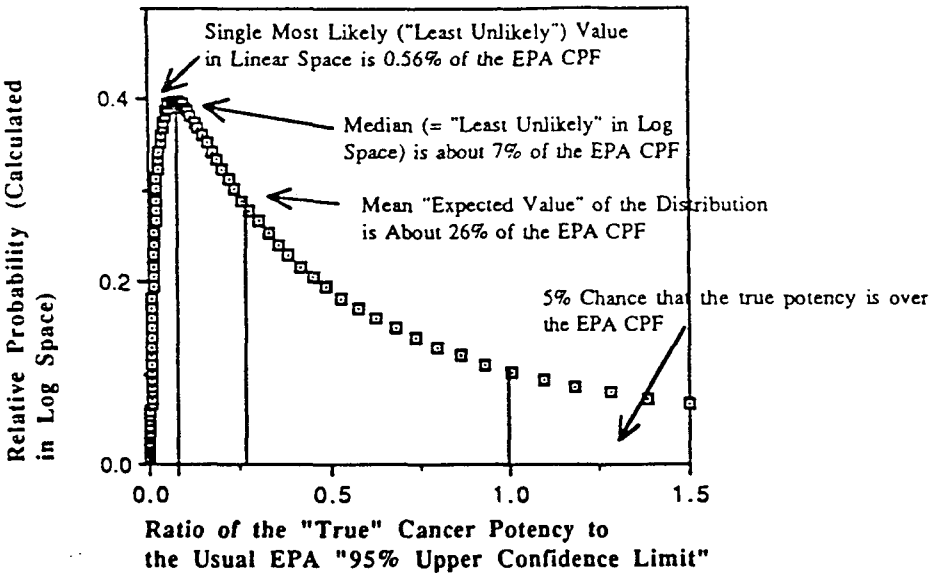

FIGURE 2. Linear plot of the estimated probability distribution for cancer potency factors (CPF) for genetically acting carcinogens. (MLE = maximum likelihood estimate.)

bottom of Table III, because the 95th percentile of a normal distribution is located about 1.64 standard deviations above the mean, it follows that a lognormal distribution with a 95th percentile value located 1/0.0724 = 13.8 times above the geometric mean must have a geometric standard deviation of about 4.93, or in other words, the standard deviation of the logarithms of our distribution of potency values is the log of 4.93, or about 0.6931.

Figure 2, which is a replot of Figure 1 on a linear axis rather than a log axis, illustrates some important properties of a lognormal distribution. It can be seen that the distribution is skewed (is asymmetrical with a long tail) to the right. Because of the asymmetry, the single most likely value (or as we prefer to term it, the "least unlikely" estimate) is not an unbiased estimate of the average or mean value of the distribution as a whole. If one were doing a classic cost–benefit assessment of various options for control of exposure, it is the mean, rather than the most likely value, that is the most relevant parameter for describing the "expected value" of the health improvements that might be obtained from the choice of one control option over another. The mean of the distribution is the average of all the potency values, weighted by their relative probability of being true (at least as represented by our lognormal assumption). This was completely missed in a recent, highly controversial critique of EPA risk assessment practices by the Federal Office of Management and Budget.[10]

The difference between the mean and the most likely value can be illustrated with a gambling analogy. Imagine that a person has the opportunity to participate in a lottery at a cost of $1.00 with a 1/10,000 chance of winning $100,000. The single most likely value of the return from this wager is zero, because there are 9999 chances of losing and only 1 chance of winning. However, the mean or expected value of the wager is $10.00 ($100,000 • 1/10,000+$0 • 9,999/10,000). *Caveat*: the "utility" of the wager to a specific person could be greater or less than $10.00, depending on the person's positive or negative enjoyment of the gamble itself, and whether $10,000 is worth exactly 10,000 times $1.00 in the person's own psychic calculus of value.

Overall, as seen in Figure 2, the mean of the lognormal distribution we have used as our best estimate of our uncertainty in cancer potencies is about 3.5 times greater than the most likely value, or, in other terms, a little more than 25% of the original EPA UCL. Thus, if our sparse set of three case studies is giving us an accurate picture of the general uncertainty in cancer potency estimates, and if the other assumptions we have made hold, the best expected value of cancer risk is only about fourfold less than the EPA UCL. What is our uncertainty in the original 7.2% ratio that was the result of our three case studies, and how would differences in this ratio affect our conclusion that a mean estimate of risk is only fourfold less than the EPA UCL? Conceivably, other carcinogens (perhaps some that do not act by direct genetic mechanisms, e.g., 2,3,7,8-tetrachlorodibenzodioxin) would have larger differences between a best estimate of cancer risk and the EPA UCL. Figure 3 shows the result of assuming a wide array of different ratios for the 7.2% best estimate from our three case studies, keeping all of the rest of the reasoning constant. It can be seen, surprisingly, that as one increases the distance between the UCL and the best estimate of potency much below about 6.7%, the ratio of the mean to the UCL

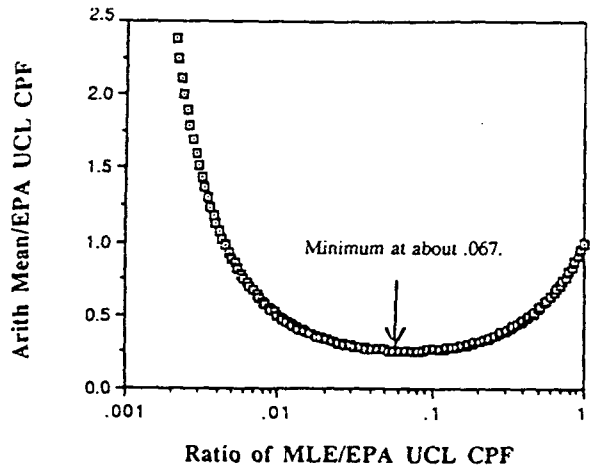

FIGURE 3. Influence of the MLE/CPF ratio on the "expected value" (mean) of a lognormal distribution of cancer potency factors (CPF).

starts to rise. This is because one necessarily increases the estimate of the geometric standard deviation, and below about the 6.7% point in the curve, the small but increasing likelihood of very large risks (far over the UCL) begins to outweigh the reduction of risk in the central/lower portion of the curve in the calculation of the overall mean.

A final and difficult point of methodology comes when we consider the small number of carcinogens (e.g., arsenic, benzene) whose cancer potency estimates are based on human epidemiological data by EPA, rather than on animal data. Because EPA estimates from human data are effectively maximum likelihood estimates, in these cases, we must take the EPA cancer potency value itself as our best estimate of risk. In the first simulations, however, we have elected to retain the same estimate of overall uncertainty as we use for the animal data. This is partly because, although the human-based estimates do not suffer from the difficulties and uncertainties inherent in animal-to-human projections, they have their own peculiar difficulties (especially related to the accuracy of the assessment of past individual exposures, the healthy worker effect, the effects of truncation of the period of observation following exposure, and the assessment of the interacting effects of potentially confounding exposures). These difficulties may often cause complications and even downward biases that are as serious as those produced in the extrapolation of animal data. Because of the importance of arsenic (whose cancer potency is based on human data) in our overall analysis, this particular judgment seriously influences the risk estimates below. We, therefore, also present below the results of a series of Monte Carlo simulation runs in which we assume that there is no uncertainty about the EPA estimate of the cancer risk from inorganic arsenic.

Tables IV and V give the Monte Carlo simulation results with and without the assumption of uncertainty in the arsenic cancer potency value. The Monte Carlo simulations were done using "Crystal Ball" simulation software (Version 1.04, Market Engineering Corporation, Denver, Colorado) for the Macintosh. Each run consisted of 5000 individual trials. Shown in these tables are the median, the mean, and the 95th percentiles of the calculated risk distributions for each chemical separately, for some aggregates of different chemicals, and for the total risk. The median is simply the middle value of each distribution of 5000 values which, if our assumptions are correct, has a 50% chance of being either larger or smaller than the true risk. The mean is the average of all 5000 values. Finally, the 95th percentile (the average of the 250th and the 251st highest of the 5000 values) is the level of risk which, if we have represented all of the uncertainties appropriately, has only a 5% chance of being smaller than the true risk. (The small differences in specific quantities that should be analogous in these two tables are due to statistical fluctuations in the results obtained in separate runs of 5000 trials each.) Comparing the results in the two tables, it can be seen that the assumption of uncertainty in the arsenic potency has a substantial influence both on the mean and high-percentile estimates of the risk and on the overall uncertainty of the risk results, i.e., the spread between the median and 95th percentile risks.

These tables give a glimpse of what might be expected if the "brave new world" of expanded use of Monte Carlo simulation for analysis of uncertainties ever arrives. One may well ask how one discusses the risk management implications of findings such as these. In presenting our work to the sponsors, we offered the following interpretive conclusions:

> "Overall, the risk analysis presented here gives some reason for concern for the long-term use of the MIRA (raw Colorado River) water supply, at least to the degree that the concentrations we measured in the reservoir samples accurately represent what is likely to be present in the finished water delivered to consumers. The overall mean estimate of lifetime risk from 2 L/day consumption is about 3 cancers per 10,000 people.[A] About 98% to 99% of this risk is derived from the presence of inorganic arsenic; trihalomethanes from chlorination represent the bulk of the remaining assessed carcinogenic risk. (The concentration of arsenic in MIRA water is well within applicable federal and state standards; however, at least at the state level, the arsenic water standards are in the process of being reevaluated.)
>
> "Our overall MIRA risk estimate is significantly

[A]We highlight the mean estimate of risk here because this is the measure of risk that would correspond to the "expected value" of risk, which might be used in a cost–benefit analysis. Higher percentiles of the risk distribution (e.g., the 95th percentile) would increase in importance for those people who wish to attach extra value to modest probabilities that the true risk might be considerably larger than the estimate of the mean risk.

TABLE IV. Results of a Monte Carlo Simulation Run (5000 Individual Trials) of Lifetime Cancer Risks from 2 L/Day Water Consumption — Assuming that the Uncertainty in the Cancer Potency of Arsenic is the Same as the Uncertainty in the Potency of Other Carcinogens

Data Set	50th Percentile (median) Risk	Arithmetic Mean Risk	95th Percentile Risk
MIRA Water			
Bromoform	1.5E-8	5.4E-8	2.2E-7
Chloroform	2.6E-8	9.5E-8	3.9E-7
Dibromochloromethane	5.4E-7	2.0E-6	7.7E-6
Bromodichloromethane	1.0E-6	3.6E-6	1.4E-5
Subtotal, all trihalomethanes	2.5E-6	5.7E-6	2.0E-5
DEHP	1.1E-7	4.4E-7	1.9E-6
Subtotal, all organics	3.0E-6	6.1E-6	2.1E-5
Arsenic	7.7E-5	3.0E-4	1.1E-3
Total	**8.4.E-5**	**3.1E-4**	**1.1E-3**
AWT Water			
DEHP	2.0E-7	7.8E-7	2.9E-6
*(Arsenic)**	*3.4E-6*	*2.3E-5*	*9.1E-5*
Total	**4.3E-6**	**2.4E-5**	**9.2E-5**

**Caveat.* Arsenic is not present in AWT water in amounts that are detectable with any confidence or reliability. These results are shown to illustrate the limitations, in terms of overall risk, of the chemical detection system for arsenic.

TABLE V. Results of a Monte Carlo Simulation Run (5000 Individual Trials) of Lifetime Cancer Risks from 2 L/Day Water Consumption — Assuming that there is No Uncertainty in the Cancer Potency of Arsenic[A]

Data Set	50th Percentile (median)	Arithmetic Mean	95th Percentile
MIRA Water			
Bromoform	1.6E-8	5.1E-8	2.0E-7
Chloroform	2.0E-8	1.0E-7	3.8E-7
Dibromochloromethane	5.2E-7	1.8E-6	7.0E-6
Bromodichloromethane	9.9E-7	3.8E-6	1.4E-5
Subtotal all trihalomethanes	2.4E-6	5.7E-6	2.0E-5
DEHP	1.0E-7	4.3E-7	1.8E-6
Subtotal, all organics	2.8E-6	6.1E-6	2.0E-5
Arsenic	8.1E-5	8.1E-5	9.6E-5
Total	**8.6.E-5**	**8.7E-5**	**1.1E-4**
AWT Water			
DEHP	2.0E-7	8.2E-7	3.2E-6
(Arsenic)[B]	*5.6E-6*	*6.3E-6*	*1.7E-5*
Total	**6.3E-6**	**7.2E-6**	**1.8E-5**

[A]Differences between Tables IV and V for chemicals other than arsenic reflect statistical sampling fluctuations between different simulation runs of 5000 trials each.

[B]*Caveat.* Arsenic is not present in AWT water in amounts that are detectable with any confidence or reliability. These results are shown to illustrate the limitations, in terms of overall risk, of the chemical detection system for arsenic.

dependent on an innovative procedure we used to assign uncertainties to cancer potency estimates from both human data (e.g., arsenic) and animal projections, as well as our analytical uncertainties in the average amounts of contaminants present. If the human-derived arsenic cancer potency value is treated as a point estimate — having no uncertainty — the estimate of the mean risk for the MIRA water falls to 0.9 cancers per 10,000 people.[B] (The mean risk estimate falls when uncertainty is removed from the calculation because the uncertainty distribution is assumed to be highly skewed (see Figure 2) with a mean that is considerably larger than the single most likely value.) Regardless of which of these estimates is deemed most appropriate for policy-making purposes by the City of San Diego, the indicated risk is not negligible (under California's Proposition 65 standards, the criterion for a de minimis risk is 0.1 conservatively estimated cancers per 10,000 people), although it is not as large as the largest environmental risks that have been identified to date (e.g., lung cancer from radon progeny in houses and from environmental tobacco smoke).

"By contrast, the water that we sampled from the AWT system does not have enough arsenic to be reliably detected. Nevertheless, our risk calculations indicate that even if we make the assumption that arsenic is, in fact, present at levels that, due to analytical insensitivities, might be present, the AWT water poses a mean risk that is tenfold lower than the MIRA water. If we base the risk assessment for the AWT water only on those chemicals that were reliably detected, the indicated mean estimate of lifetime risk is slightly less than 1 cancer per 1 million people."

Conclusions and Caveats

The extended example above illustrates that it is possible to provide some expanded insight into the range of "not clearly incorrect answers" for risks, taking into account multiple sources of uncertainty, and without a great deal more work than is usually done in more conventional "screening"-type risk analyses, such as those currently done for Superfund sites. Of course, the present analysis could have benefited from a far more systematic study of uncertainties for individual cancer potencies in the light of specific toxicological information for particular chemicals, a more adequate representation of the delivered dosage of water contaminants due to multiple routes of exposure from water (e.g., dermal, inhalation), and some other factors. Nevertheless, we feel that this level of analysis does offer some improvements in the information provided to decision-makers for a modest increment in analytical efforts.

Of course, no analysis of the combined effects of multiple sources of uncertainties can be better than the accuracy of the individual estimates of component uncertainties that are fed into the simulation. Arriving at appropriate descriptions of these component uncertainties is by no means a trivial task. In this regard, we have previously offered[11] three tongue-in-cheek "laws" of uncertainty/variability analysis that should inject a final note of the need for due skepticism here.

1. Nearly all distributions look lognormal, as long as you do not look too closely. (One is, therefore, well advised to reason carefully about the likely causes of uncertainty or variability in specific parameters and whether this suggests a particular distribution is appropriate.)
2. Any estimate of the uncertainty of a parameter value will always itself be more uncertain than the estimate of the parameter value. (For example, fluctuations in individual data points will generally have a greater influence on the estimate of a standard deviation than on the estimate of the mean of the parameter.)
3. The application of standard statistical techniques to a single data set will nearly always reveal only a trivial proportion of the overall uncertainty in the parameter value. (This is because systematic error among data sets is generally much larger than random error within data sets.)

[B] A more conventional calculation, utilizing simply the basic EPA cancer potency factors without modification, yields a lifetime risk of about 1.3 cancers per 10,000 people. These cancer potency factors are something of a mixed group, in that implicitly the arsenic and benzene values, derived from human data, are central tendency values, whereas the potency estimates derived from animal data (e.g., chloroform) are upper confidence limit values.

References

1. Hattis, D.; Silver, K.: Human Interindividual Variability — A Major Source of Uncertainty in Assessing Risks for Non-Cancer Health Effects. Presented at the Air and Waste Management Association Annual Meeting, Pittsburgh, PA (June 1990).
2. Ashford, N.A.: *Outcome vs. Process in Decision Making.* Ann. NY Acad. Sci. 572:76–78 (1989).
3. Hattis, D.; Smith, J.: What's Wrong with Quantitative Risk Assessment. *In:* Biomedical Ethics Reviews: 1986, Quantitative Risk Assessment. J.M. Humber and R.F. Almeder, Eds. Humanan Press, Clifton, NJ (1987).
4. Hattis, D.; Tuler, S.; Finkelstein, L.; Luo, Z.: A Pharmacokinetic/Mechanism-Based Analysis of the Carcinogenic Risk of Perchloroethylene. Report CTPID 86-7 to the National Institute for Occupational Safety and Health and the National Institute for Environmental Health Sciences by M.I.T. Center for Technology, Policy and Industrial Development. NTIS Pub. No. PB-88-163209. National Technical Information Service, Springfield, VA (September 1986).
5. Hattis, D.: A Pharmacokinetic/Mechanism-Based Analysis of the Carcinogenic Risk of Ethylene Oxide. Report No. CTPID 87-1. M.I.T. Center for Technology, Policy and Industrial Development, Cambridge, MA (August 1987).
6. Hattis, D.; Wasson, J.: A Pharmacokinetic/Mechanism-Based Analysis of the Carcinogenic Risk of Butadiene. Report No. CTPID 87-3. M.I.T. Center for Technology, Policy and Industrial Development, Cambridge, MA. NTIS Pub. No. PB-88-202817. National Technical Information Service, Springfield, VA (November 1987).
7. Hattis, D.: Use of Biological Markers and Pharmacokinetics in Human Health Risk Assessment. Environ. Health Perspect. 89:229–238 (1991).

8. Hattis, D.; Erdreich, L.; DiMauro, T.: Human Variability in Parameters that are Potentially Related to Susceptibility to Carcinogenesis — I. Preliminary Observations. Report No. CTPID 86-4 to the U.S. Environmental Protection Agency, by M.I.T. Center for Technology, Policy and Industrial Development, Cambridge, MA (May 1986).
9. Ballew, M.; Hattis, D.: Reproductive Effects of Glycol Ethers in Females — A Quantitative Analysis. Report No. CTPID 89-7. M.I.T. Center for Technology, Policy and Industrial Development, Cambridge, MA (July 1989).
10. Office of Management and Budget: Current Regulatory Issues in Risk Assessment and Risk Management. In: Regulatory Program of the United States Government, April 1, 1990–March 31, 1991, pp. 13–26. OMB, Executive Office of the President of the United States, Washington, DC (1990).
11. Hattis, D.: Letter to the Editor: Three Candidate "Laws" of Uncertainty Analysis. Risk Analysis 10:11 (1990).