Appearing in Risk Analysis, vol.28 no.1, 2008

Regulating under Uncertainty: Newsboy for Exposure Limits

Roger M. Cooke
Resources For the Future and
Department of Mathematics
Delft University of Technology

Margaret MacDonell
Argonne National Laboratories
Chicago, Ill

Nov. 2, 2007

ABSTRACT:

Setting action levels or limits for health protection is complicated by uncertainty in the dose-response relation across a range of hazards and exposures. To address this issue, we consider the classic newsboy problem. The principles used to manage uncertainty for that case are applied to two stylized exposure examples, one for high dose and high dose-rate radiation and the other for ammonia. Both incorporate expert judgment on uncertainty quantification in the dose-response relationship. The mathematical technique of probabilistic inversion also plays a key role. We propose a coupled approach, whereby scientists quantify the dose-response uncertainty using techniques such as structured expert judgment with performance weights and probabilistic inversion, and stakeholders quantify associated loss rates.

KEY WORDS: Dose-response uncertainty, health risk, action levels, probabilistic inversion, expert judgment, performance weights

1. INTRODUCTION

Setting action levels or exposure guides for hazardous substances involves balancing and distributing risks and benefits. Policy makers and regulators wish to make such decisions in a rational and fair manner. However, current science tells us that the relation between the dose of a hazardous substance and the adverse response is uncertain, sometimes *very* uncertain. Any choice the regulator makes can be perceived as favorable to the interests of some stakeholders and hostile to the interests of others, and any choice will be open to accusations of bias and favoritism. Existing tools can be applied to better define exposure limits in the face of large uncertainties, in a transparent and defensible manner, toward the goal of overall health protection.

Some believe uncertainty is not relevant to regulation because decisions will be taken on the basis of a “best estimate" without regard for underlying uncertainties. Others attempt to account for uncertainty by downwardly adjusting an empirically derived best estimate using deterministic

factors with a simple narrative. Standard toxicity values used to convert exposures to estimated cancer risks, or indices of non-cancer effects, have traditionally applied default factors of 10 or 3. These "uncertainty factors" are designed to address uncertainties in interspecies extrapolation, human variability, database deficiencies (such as lack of reproductive and developmental toxicity studies), and use of a lowest observed adverse effect level (LOAEL) rather than a no observed adverse effect level (NOAEL). Neither approach enjoys a firm scientific basis.

Decision theorists, on the other hand, have long had the tools to appropriately account for uncertainty. These tools can be applied to the development of exposure limits or action levels for environmental contaminants. Although the process has been illustrated in a number of other applications, for example, to guide water management in the Netherlands[(1)], these tools have not yet found their way into mainstream decision making for hazardous substances.

A simple example can help frame the issue of managing risk with uncertainty in the dose-response relation. You have a headache, and you have a box of pills. Instead of prescribing how many pills you should take, the box gives a probability distribution over the number of pills needed to cure a headache. How many pills do you take? The answer depends on the consequences of taking too many or too few. If one pill too many causes severe nausea but too few means you suffer a bit longer, you would choose a smaller number to be relatively certain of not becoming nauseous. However, if too many has no adverse affect other than wasting pills, but one pill shy has no positive effect, then you would likely choose a relatively high number. In deciding how many to take, you must balance the regret of taking too few with the regret of taking too many. The same applies when regulating under uncertainty. Greater or lesser uncertainty should guide the way risk is regulated to support overall public protection.

The principle for managing under uncertainty is elegantly captured by the classic *newsboy problem* in operations research. To demonstrate how the concept applies to exposure guides, this paper considers two examples, one for radiation and the other for an acute chemical exposure. The method requires from scientists the quantification of uncertainty in dose-response models, and from stakeholders the quantification of loss rate ratios. Mathematical techniques such as probabilistic inversion and iterative proportional fitting are also valuable for better quantifying dose-response uncertainty to support health-based exposure limits and action levels. Methods for quantifying loss rate ratios and dose-response uncertainty are available in the literature, and it is not the purpose of this paper to make specific recommendations in this regard. Rather, the goal is a proof of concept for using such quantifications to regulate under uncertainty.

2. THE NEWSBOY PROBLEM

A newsboy buys newspapers at the beginning of each day. If he buys more than he can sell in the day, he must dump the leftovers at a loss. If he buys fewer than he can sell, he loses profit. How many should he buy? To formulate this problem mathematically, let X be the number of newspapers that could be sold, with the cumulative distribution function *F*:

$$F(r) = PROB(X \leq r) \qquad (1)$$

The newsboy must adopt a value, x_o, for decision making. Suppose the costs of under- and over-estimating X can be characterized as:

Loss rate if realization is too low: *If* $x_o > X$, *pay* $L \times (x_o - X)$
Loss rate if realization is too high: *If* $x_o < X$, *pay* $H \times (X - x_o)$

The newsboy's expected loss if he orders x_o papers is:

$$E(\mathcal{L}(X,x_o)) = H \int_{X > xo} (X - x_o)\, dX + L \int_{X < xo} (x_o - X)\, dX \qquad (2)$$

Setting the derivative of (2) with respect to x_o equal to zero, the value that minimizes his expected loss is found to be:

$$x_o = F^{-1}(H/(H+L)) \qquad (3)$$

which depends only on the ratio of *H* and *L*. Thus, if the loss rate for too-high realizations is much greater than the loss rate for too-low realizations, then $H / (H + L)$ is close to *1* and our estimate would be a very high percentile of the distribution for *X*. (Note it is not necessary that the loss rates be constant. The technique applies to any loss function; however, the solution (3) would become more complicated if the loss rates are not constant.)

To determine the cumulative distribution function *F*, the newsboy can use historical data, if available. Otherwise, he relies on experts to assess this, according to protocols widely in use. The loss rates he must assess himself, since he is the one who suffers the losses. Suppose he buys each paper for $0.50 and sells it for $1.00, but dumps any unsold paper for $0.40. If the newsboy follows a narrow economic reasoning, his loss rate ratio would be 50 / (50 + 10) = 0.833. One missed sale would be just as painful as 5 dumped newspapers. He might also reflect that each missed sale is also a disgruntled customer who might defect to the competition. Taking these strategic factors into account, suppose he eventually decides that 10 missed sales is equivalent to 75 dumped newspapers. He then solves $H \times 10 = L \times 75$ to find the ratio of *H/L* (see Fig. 1). The number x_o that minimizes expected loss is the *H/(H+L) = 88.2 percentile* of his uncertainty distribution for how many papers could be sold.

To illustrate with some numbers, if the amount of sellable newspapers is normally distributed with mean 1,000 and standard deviation 200, the newsboy should order 1,237. However; if the standard deviation is not 200 but 400, he should order 1,475. Note that the *expected number of sellable newspapers is the same in both cases, only the uncertainty changes*.

The difference in price between the initial outlay for the expected number of sellable papers (1,000) and the optimal order quantity (1,475) may be called the *first uncertainty premium*. It is a price which the newsboy should pay initially to minimize his expected loss. A *second uncertainty premium* is discussed below.

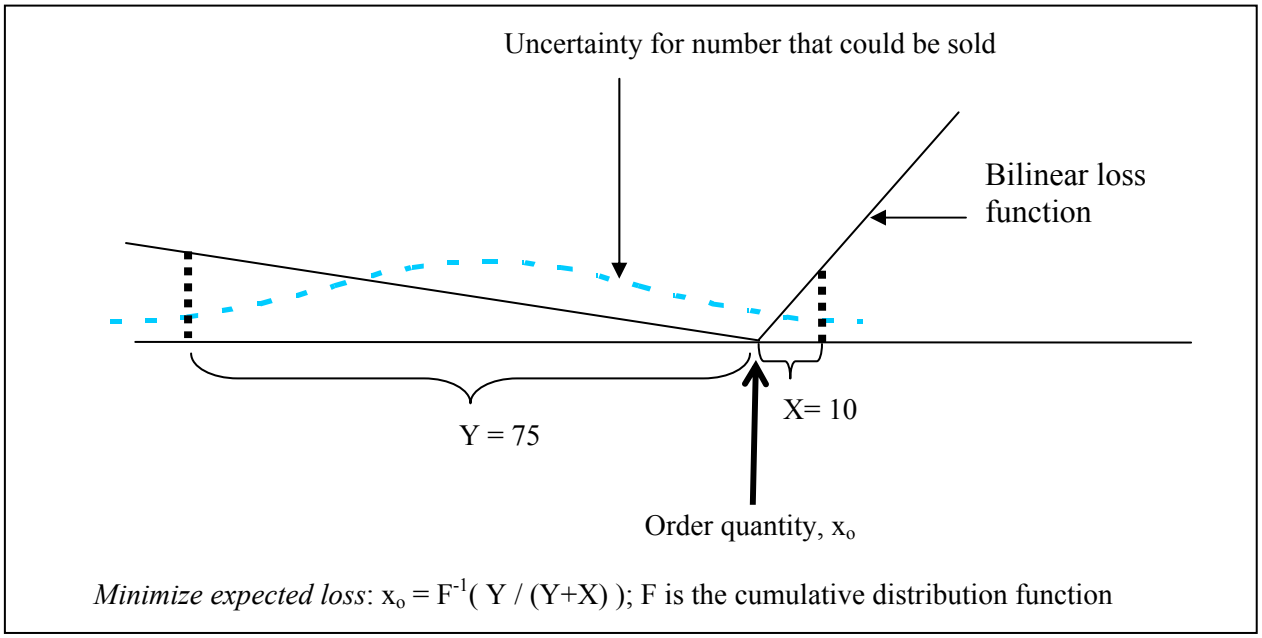

Fig. 1. Loss ratios for newsboy: H × 10 = L × 75.

An article in the Harvard Business Review illustrates how a savvy newsboy can improve his bottom line (Fisher et al.[(2)]). The authors remark:

> *The real problem, though, is that most companies do a poor job of incorporating demand uncertainty into their production-planning process. They are aware of demand uncertainty when they create a forecast – witness the widespread reliance on safety stocks – but they design their planning processes as if that initial forecast truly represented reality. They do this for two reasons. First, it's complicated to factor multiple demand scenarios into their planning, most companies simply don't know how to do it. Second, the dramatic increase in demand unpredictability is fairly recent, so most companies haven't yet changed their planning systems to adapt to it.*

The authors further report that correctly factoring uncertainty into the planning process resulted in a 67% increase in profits for their sportswear company.

A variation on this problem to be used in illustrating the concept here is as follows. Suppose we have expressed risk $R(c,Z)$ as a function of an exposure value c and a random variable Z, where R is monotonically increasing in c. Suppose that a level r has been set, and that expected loss is given by the bilinear function:

$$E(\mathcal{L}(R(c,Z))) = H \int_{R(c,Z)>r} (R(c,Z) - r)\, dZ + L \int_{R(c,Z)<r} (r - R(c,Z))\, dZ.$$

Our problem is to choose c to minimize this expression. Differentiating with respect to c, the solution is a value c^* satisfying:

$$F_{R(c^*)}(r) = \frac{H}{H + L(E(\partial_c R(c^*) \mid R(c^*) < r) / E(\partial_c R(c^*) \mid R(c^*) > r))}, \qquad (3.a)$$

where $F_{R(c^*)}$ is the cumulative distribution function of the risk when the exposure c^* is realized. The term

$$\frac{E(\partial_c R(c^*) \mid R(c^*) < r)}{E(\partial_c R(c^*) \mid R(c^*) > r)}$$

measures the non-linearity in risk, relative to r, as a function of exposure. It is approximately unity if R(c,Z) is approximately linear in the neighborhood of c^*. In this case *(3.a)* and *(3)* coincide, and we choose the value c^* which makes the pre-determined value r the newsboy solution. This approximation is used in the examples of Section 3. The difference in initial outlay between *(3.a)* and *(3)* may be called the *second uncertainty premium.* Limited experience to date suggests that the first uncertainty premium is generally larger than the second.

3. FOUR-STEP APPROACH AND EXAMPLES

In applying the newsboy concept to support the development of exposure guides, governmental agencies tasked with protecting public health do not measure their success by bottom-line profits. Nevertheless, they face similar problems in dealing with uncertainty, namely: uncertainty in the effects of regulatory decisions and unfamiliarity with planning under uncertainty. We suggest a four-step newsboy-based framework for determining an action level or exposure limit when the dose-response relationship is uncertain:

1. Decide what you would do if you knew the dose-response relation with certainty. (For example, you would regulate at a maximum allowable risk level and find the dose corresponding to that risk.)
2. Quantify the uncertainty in the dose-response relation.
3. Determine your loss rates for too-high and too-low realizations.
4. Choose the quantile of the dose distribution that minimizes expected loss.

These steps are illustrated with two fictive but realistic examples. The first addresses low-linear energy transfer (LET) radiation, for which the dose-response relation has been assumed to take a very simple form. The second addresses acute exposure to a chemical, for which sophisticated mathematical techniques are indicated for Step 2. Mathematical approaches are outlined following the example for this chemical.

3.1. Lifetime Cancer Risk from High Dose, High Dose-Rate, Low-LET Radiation

In situations like space travel and nuclear disasters, individuals can be exposed to high doses of low-LET radiation (e.g., gamma rays) at a high dose rate. Table I shows the estimated

lifetime cancer mortality per 100 people from an acute exposure of 1 Gray (Gy) as absorbed dose (or 100 rad), or 1 sievert (Sv) as equivalent dose (or 100 rem), for various cancer sites. Values in the first two columns come from the joint European Union-U.S. Nuclear Regulatory Commission (EU-NRC) expert judgment study for accidents at nuclear power plants;[3] confidence bounds are given in the third column. The next five columns reflect best estimates from other studies, including the Biological Effects of Ionizing Radiation (BEIR) V report.[4]

Cancer Site	Estimated Radiation Exposure-Induced Deaths (REID) from Cancer, per 100						
	EU-NRC Expert Judgment		BEIR V	BEIR VII	ICRP 60	UNSCEAR	COSYMA
	Median	90% Confidence Bound					
Bone	0.035	<0.001, 0.88	-	-	-	-	0.01
Colon	0.98	0.011, 3.35	-	0.61	3.24	0.6	2.24
Breast	0.78	0.11, 3.78	0.35	0.37	0.97	0.6	0.80
Leukemia	0.91	0.026, 2.33	0.95	0.61	0.95	1.0	0.52
Liver	0.086	<0.001, 2.02	-	0.16	-	0.9	-
Lung	2.76	0.59, 8.77	1.7	2.10	2.92	2.1	0.90
Pancreas	0.17	<0.001, 1.26	-	-	-	-	-
Skin	0.039	<0.001, 0.37	-	-	0.03	-	0.01
Stomach	0.30	<0.001, 4.01	-	0.22	0.51	1.2	-
Thyroid	0.059	0.001, 0.71	-	-	-	-	0.17
Other/solid	2.60	<0.001, 10.8	2.6	1.3/5.1	-	-	-
All cancers	10.2	3.47, 28.5	7.9	see note	12.05	9	5.02

Table I. Comparison of elicited high dose and high dose-rate lifetime low-LET radiation exposure-induced deaths (REID) from cancer at various sites for a general EU/U.S. population with those derived from other sources (10^{-2}Gy^{-1}).[3] BEIR V and VII calculate excess deaths for the extant (current) U.S. population and account for competing risks.[4,8] Values for ICRP 60[9] are as extracted in UNSCEAR.[10] The UNSCEAR REID reflects the extant Japanese population of both genders and all ages, using an attained-age model); the COSYMA REID is taken from Ehrhardt et al.[11]. The BEIR V and

VII values represent the average for males and females per the committee's preferred models. For breast cancer, values reflect half the estimates for females (as estimates per total population). "Other/solid" addresses additional sites or groups not captured by the preceding categories (further values exist but for other systems, e.g., respiratory); the BEIR V value is as summarized in BEIR VII; for BEIR VII, 1.3 is for other cancers and 5.1 is for solid cancers; to approach a total, these can be added with the leukemia value (summing to just above 7).

Recent analyses acknowledge limitations in current uncertainty characterization for radiological risk estimators.[5] The quantification underlying values in Table I column 3 represents perhaps the best effort to date to capture the scientific uncertainty in this area. Note that the spread of best estimates is much narrower than the 90% central confidence intervals from the structured expert judgment study. The spread of best estimates is not intended to (and indeed does not) characterize the uncertainty in cancer induction rates. From the EU-USNRC study data summarized in this table, the high dose, high dose rate, all-cancer mortality risk to an exposed individual has an uncertainty distribution that is adequately described as log normal with median 0.102, mean 0.126 and error factor 2.94. The four-step newsboy approach is applied as follows:

Step 1: What would we do if the dose-response relation were known with certainty?

Suppose it is decided that under the circumstances for which a given regulatory policy rule is being designed, the maximum allowable additional lifetime risk of radiation exposure-induced death (REID) is 0.05. This target is the result of a deliberation falling outside the purview of this study. (For comparison, it represents about 5 times the U.S. average risk estimated from natural background radiation[6] or about 2.5 times the general U.S. background rate including medical exposures, reflecting recent increases in the use of computed tomography [CT] scans[7]). For this example deliberation, presumably the losses of higher exposure limits in terms of risks to exposed individuals are balanced against the losses implied by the inability to accomplish the mission for the sake of which the risks are assessed.

Step 2: Quantify the dose-response uncertainty.

For high doses (the region considered here), the probability of death is assumed to be linear with dose.[8] The response in lifetime probability of REID at dose q [Gy/min] can be expressed with the simple formula:

$$Life\_Prob\_REID(q) = R(q) = Z \times q \qquad (4)$$

where Z is a random variable. Based on the EU-USNRC data summarized in Table I for all cancers, Z can be approximated as a log normal variable with median 0.102 and error factor 2.94.

Step 3: Determine losses for too-high and too-low realizations.

For simplicity, we assume that our loss is expressed as a bilinear function of the risk difference between the realized and target risk. In aiming to regulate at a dose, the risk will be uncertain. If the risk turns out to be *lower* than the target of 0.05 (i.e., we are more conservative than intended), then we suffer loss in terms of mission impairment at rate L. If the risk turns out

to be *higher* than the target, then we subject those exposed to a greater risk, with loss rate H. We do not need to assess these losses: here, it is sufficient to determine, say, that the loss of realizing a risk of 0.01 is just as bad as the additional loss of realizing a risk of 0.06; then $H/(H+L) = 0.8$. Whereas the newsboy's economic problem was quite simple, deriving loss rates from economic costs and benefits in a real situation would surely be difficult and fraught with uncertainty. Stakeholder preferences could also be elicited with tools economists have developed for this purpose; contingent valuation, willingness to pay, discrete choice, etc. Alternatively, the regulator may receive a default loss rate ratio from higher hand to be used if resources for case-specific loss rate quantification are not available. Our point is not to mandate how loss rate quantification should be done, but to indicate how it should be used, and to indicate which problem is whose."

Step 4: Choose the quantile of the dose distribution that minimizes expected loss.

From the preceding information, the expected loss can now be written just as in (2).

$$E(\mathcal{L}(R,q)) = H \int_{R > 0.05} (R - 0.05)\, dR + L \int_{R<0.05} (0.05 - R)\, dR \qquad (5)$$

$$= H \int_{Zq > 0.05} (Zq - 0.05)\, dZ + L \int_{Zq < 0.05} (0.05 - Zq)\, dZ$$

$$= Hq \int_{Z > 0.05/q} (Z - 0.05/q)\, dZ + Lq \int_{Z < 0.05/q} (0.05/q - Z)\, dZ.$$

From (3), the solution is the value q such that $P\{R < 0.05\} = P\{Z < 0.05/q\} = H/(H+L) = 0.8$. The 80th percentile Z is 0.181, so $0.05/q = 0.181$ yields $q = 0.28$. This dose is smaller than we would obtain using the median $(R(q) = 0.102 \times q;\ 0.05/0.102 = 0.49)$ or mean $(R(q) = 0.126 \times q;\ 0.05/0.126 = 0.40)$ as best estimate. Had we used (3.a) without the linear approximation, we would have found the expected loss minimized at dose 0.2 Gy/min instead of 0.28 Gy/min.

An equivalent formulation is: $P\{Z < 0.05/q\} = P\{q < 0.05/Z\} = 0.8$. From (4) we recognize 0.05/Z as the random variable whose distribution is the distribution of the dose when the risk level of 0.05 is stipulated. The value q that satisfies this equation is the 20th percentile of the distribution of the dose realizing risk 0.05. Fig. 2 shows densities for doses realizing risk levels from 0.01 to 0.1, while Fig. 3 presents the cumulative distribution of the dose at a risk of 0.05.

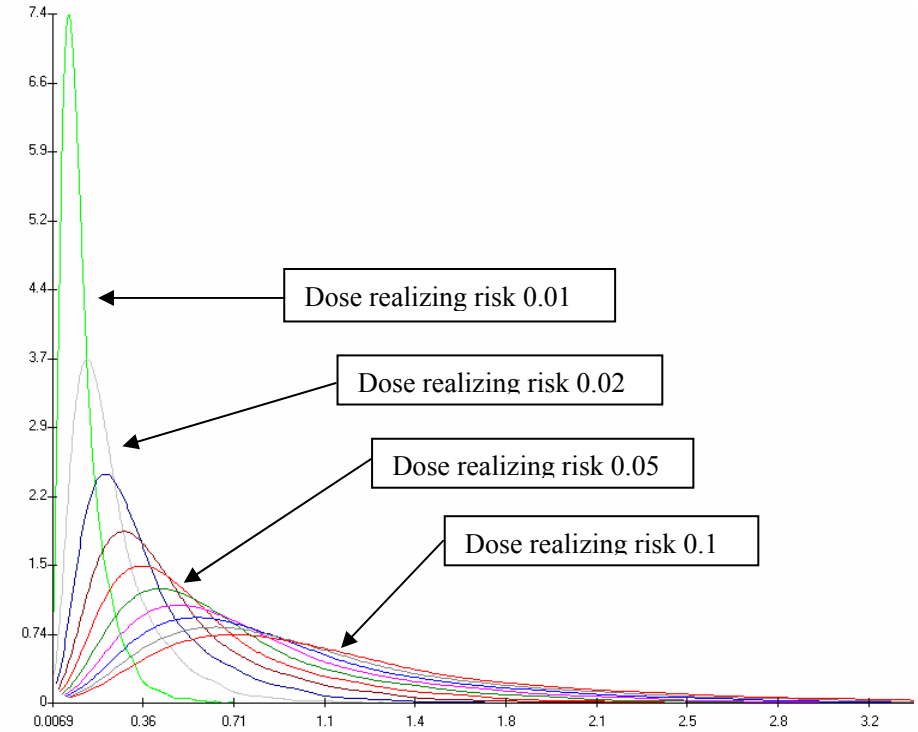

Fig. 2. Densities for dose (Gy/min) realizing risk levels of 0.01 to 0.1.

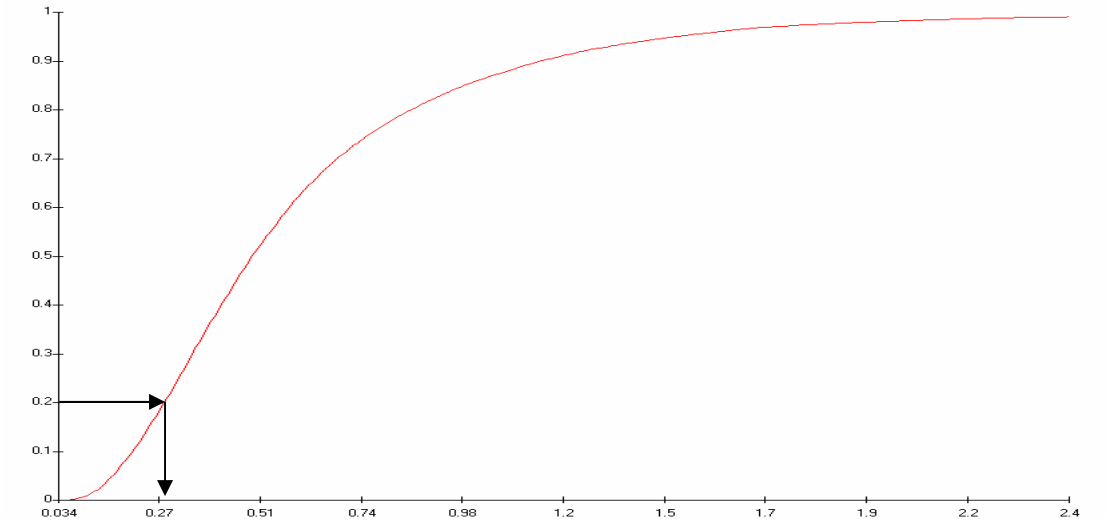

Fig. 3. Cumulative distribution function for the dose for R = 0.05; the 20th percentile is 0.28.

3.2. Acute Mortality Risk from Ammonia

This second example addresses the problem of setting an acute exposure level for the hazardous chemical ammonia. The same four steps are followed as above, but in this case the dose-response is not represented by a simple linear relation. The form of the relation is stipulated by tradition, but the parameters are uncertain. Uncertainty in the dose-response relation can be quantified via a joint distribution over these parameters.

Step 1: What would we do if the dose-response relation were known with certainty?

Suppose our regulatory policy rule is: The maximum allowable concentration is such that the probability of death within 24 hours due to a 60-minute exposure is 10^{-5}. (For context, 10^{-4} is the upper level of the U.S. Environmental Protection Agency (EPA) target range for incremental lifetime risk from contaminated sites, above which action is typically considered, while 10^{-6} represents the lower level or point of departure.[12] With the U.S. average lifetime cancer risk from natural radiation alone estimated at 10^{-2}, the target increment of 10^{-4} to 10^{-6} means cancer risk is being managed to a low fraction of natural background.)

If we knew the dose-response relation for ammonia with certainty, we would simply compute the parts per million (ppm) that implemented the regulatory rule. The Dutch government, for example, stipulated a dose-response relation for ammonia under the above conditions as the probit function:

$$\Phi^{-1}(r) + 5 = A + B \times \ln(C^n \times t) = -9.35 + 0.71 \times \ln(C^2 \times t) \qquad (6)$$

where Φ is the standard normal cumulative distribution function, r is the probability of death, C is the concentration in ppm, and t is the exposure time in minutes.[13,14,15] (See Finney[16] for the basis of why dose-response functions are often expressed in this form.)

Ten published probit relations as applied to acute ammonia exposure are collected in Table II. The values that realize risk levels of 10^{-5}, 10^{-4}, 10^{-2}, and 10^{-1} are shown in the four right columns. The NOAEL identified in the EPA Integrated Risk Information System (IRIS) database as the point of departure for the inhalation reference concentration is 9.2 ppm (reflecting repeated daily exposures;[17] no dose-response relation is given). Note that the spread of concentrations in Table II at a risk level of 10^{-5} intersects the spread for the 10^{-1} risk level.

Source	Ammonia Probit Parameters			Fatal Concentration (ppm) from a 60-Min Exposure in 1 Day			
	A	B	n	0.00001	0.0001	0.01	0.1
Lees[18]	-9.82	0.71	2	218	321	855	1784
Goossens et al.[15]	-35.02	2.01	2	941	1078	1525	1977
CCPS[19]	-35.9	1.85	2	2576	2986	4351	5770
TNO[20]	-16.5	1	2	714	938	1881	3171
Canvey Island[21]	-46.95	2.205	2.75	587	642	808	960

IchE[22]	-16.14	1	2	596	783	1572	2649
Purple Book[23]	-14.92	1	2	324	426	854	1439
Perry and Articola[24]	-28.33	2.27	1.36	630	753	1185	1665
Green Book[25]	-15.12	1	2	358	471	944	1591
Rijnmond[26]	-30.57	1.385	2.5	1642	1923	2874	3885

Table II. Probit values for acute exposure to ammonia: $\Phi^{-1}(r) + 5 = A + B \times \ln(C^n \times t)$, where Φ is the standard normal cumulative distribution function, r is the probability of death, C is the concentration in ppm, and t is the exposure time in minutes (min).

Step 2: Quantify the dose-response uncertainty

Evidently the dose-response relation is *not* known with certainty. When experts are asked *What concentration over 60 minutes would lead to death for 10% (or 50%, or 90%) of the exposed reference population within 24 hours?,* they are not certain of the answer. Table III presents the median assessments and the 5% to 95% confidence range from a structured expert judgment study conducted in the Netherlands.[15] The performance-based weighted combinations of the experts' judgments are shown; the equal-weight combinations gave substantially wider distributions.

Percent of reference population affected	Percentiles of concentration (ppm) producing effect		
	5%	50%	95%
10%	2295	2700	3105
50%	3213	4158	4782
90%	4208	4950	5693

Table III. Percentiles for ammonia concentration (ppm) causing death within 24 hours in given percentages of the reference population, from a 60-minute exposure.[15]

To capture this uncertainty in a tractable form, the probit equation is written with unknown constants *A* and *B* (the exponent of the concentration *C* could also be regarded as unknown, but that component is not considered here):

$$\Phi^{1}(r) + 5 = A + B \times \ln(C^2 \times t) \qquad (7)$$

Now the question is: Which joint distribution over the unknowns *(A,B)* would reproduce the probabilities in Table III when we substitute r = 0.10, r = 0.50, and r = 0.90 in Equation (7)? This is called a *probabilistic inversion* problem. Techniques for finding an optimal distribution given

a starting distribution are briefly summarized in Section 3.3. Any number of different mathematical models could have been used, but tradition here favors the probit model. Solving (7) for C, we obtain

$$\frac{[\exp((\Phi^{-1}(r) + 5 - A)/B)]^{0.5}}{t^{0.5}} = C(r). \quad (8)$$

Sampling from the distribution over *(A,B)*, we find the uncertainty distributions on the *C(r)* values realizing various given risk levels. These are shown in Table IV.

Risk Level	Percentiles		
	5%	50%	95%
0.00001	963	1550	1640
0.01	1740	2220	2430
0.1	2360	2820	3090

Table IV. Uncertainties in ppm values realizing given risk levels, from the structured expert judgment study in Goossens et al.[13, 15]. The reference population was a "general population of a typical industrialized and developed nation like the Netherlands" with an assumed breathing rate of 15 m^3/day (appropriately transformed for children), 20% of the population under 10 yrs or over 70 yrs, and 5% having illness.

To provide context for the numbers in Table IV, the National Academies'[27] acute exposure guideline level (AEGL)-3 is "the airborne concentration above which it is predicted that the general population (including sensitive individuals) could experience life-threatening effects or death". The current AEGL-3 value for 60-minute exposure to ammonia (derived by probit analysis) is 1,100 ppm. The expert judgment study gave 5%. 50% and 95% values for Emergency Response Planning Guidelines (ERPGs) 'above which there would be an unacceptable likelihood of observing life-threatening health effects' of 500, 1000, and 1500 ppm (Goossens et al[13]). These numbers are consistent with the experts' assessments in Table IV, which give a *5%* chance that the concentration value imposing a risk of 10^{-5} of death within *24* hours following a *60* minute exposure is below *963ppm*.

Step 3: Determine losses for too-high and too-low realizations.

For this step, the same approach is applied as in the radiation example. Risk is written as a function of concentration C, remembering that this depends on the random vector *(A,B)* as well:

$$R(C) = \Phi(A - 5 + B \times \ln(C^2 \times 60)). \quad (9)$$

$$E(\mathcal{L}(R(C))) = H \int_{R(C) > 0.00001} (R - 10^{-5})\, dR + L \int_{R(C)<0.00001} (10^{-5} - R)\, dR. \quad (10)$$

We assume as before that the loss rate ratios satisfy H/(H+L) = 0.8.

Step 4: Choose the quantile of the dose distribution that minimizes expected loss.

We solve as before, this time for concentration, by substituting from (9):

$$0.8 = P\{\Phi(A - 5 + B \times \ln(C^2 \times 60)) \leq 10^{-5}\} = P\{A + B \times \ln(C^2 \times 60) \leq 0.735\}. \qquad (11)$$

Unlike the radiation example, we cannot solve this equation by finding a quantile of a random variable. Rather, we must sample the distribution for *(A,B)* for different values of *C* until we find one for which (11) holds. The solution is *C = 1,236 ppm.* Had we used (3.a) without the linear approximation, we would have found the expected loss minimized at dose 1,217 ppm instead of 1,236 ppm.

This solution strategy, although valid, is rather clunky. Using (8), we apply the same trick we used in the first example; namely, the concentration that makes 10^{-5} the 80th percentile of the risk distribution is the 20th percentile of the concentration distribution realizing a risk of 10^{-5}:

$$0.8 = P\{\Phi(A - 5 + B \times \ln(C^2 \times 60)) \leq 10^{-5}\} \Leftrightarrow \qquad (12)$$

$$0.2 = P\{\exp[(\Phi^{-1}(10^{-5}) + 5 - A)/B)]^{0.5} / t^{0.5}] \leq C\}.$$

Fig. 4 shows the cumulative distribution function of the dose corresponding to a 10^{-5} risk, with the 20th percentile indicated.

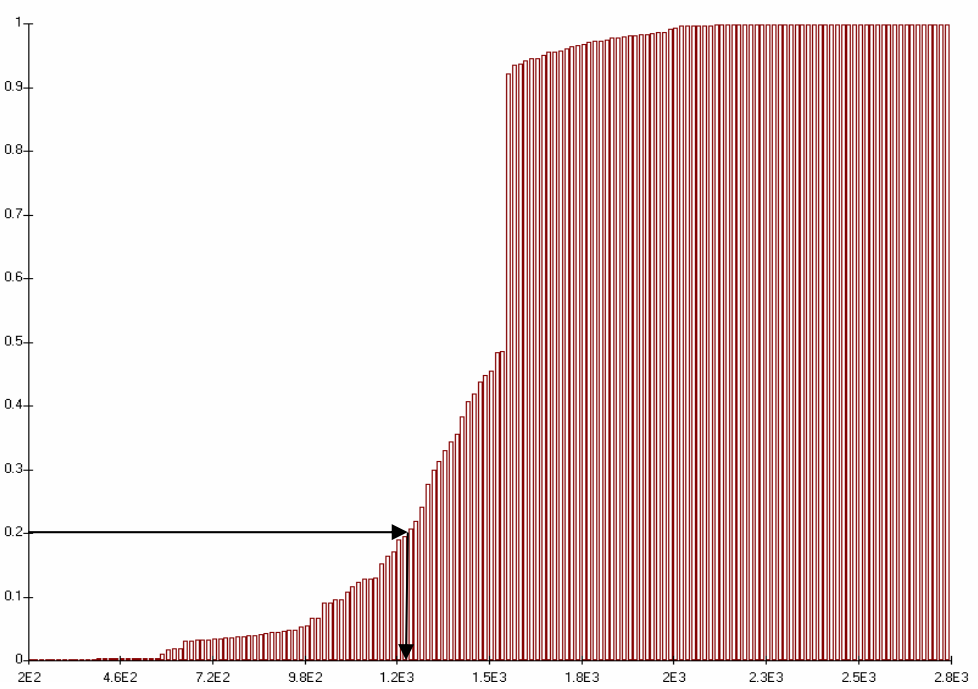

Fig. 4. Cumulative distribution of the ammonia exposure level in ppm for the risk level of 10^{-5}; the 20th percentile is about 1,240 ppm.

3.3. Finding a Distribution over the Dose-Response Relation with Probabilistic Inversion

Finding a distribution over *(A,B)* for situations like ammonia in the example is actually a rather simple problem of probabilistic inversion. Techniques developed at the Delft University of Technology have been applied many times for similar cases.[28] In fact, the first problem of this kind involved obtaining distributions of coefficients in dose-response relations similar to those described above.[15]

The process involves the following elements. Experts are asked to state selected percentiles (e.g., their 5, 50, and 95 percentiles) for a given exposure (in appropriate units) that would cause responses $r_1, r_2, \ldots r_n$ in a reference population. These expert assessments are combined into one set of percentile assessments, as in Table III. Consider then the problem of finding a joint distribution on *(A,B)* which, when applied in (8), optimally reproduces the experts' probabilities. In the ammonia example, the experts' percentiles are given in Table III. The problem can be solved with a simple re-weighting Monte Carlo technique. The idea is as follows.

First:

(a) Choose a wide distribution for A and B encompassing all plausible values.

Then:

(b) Sample this distribution a large number of times.

(c) For each sample, i.e., each combination of A and B values, compute (8), but with the responses $r_1, \ldots r_n$.

(d) Now re-weight these samples such that the weighted Monte Carlo sample complies with the experts' quantile specifications.

The joint distribution of A and B in this re-weighted distribution is a probabilistic inverse of the experts' distributions through the model (8).

A well-known statistical algorithm called iterative proportional fitting (IPF) can be used to find the weights for re-weighting the original sample. If this algorithm converges, then it converges to the distribution over *(A,B)*, satisfying the constraints, which is minimally informative given the starting distribution.[29] If IPF does not converge, then a recently discovered variation on this technique can be used for which the stationary points are solutions, if solutions exist, and which minimizes an entropy-based measure of fit.[28,30] The latter reference contains many examples and applications of probabilistic inversion.

4. CONCLUSIONS

A decision theoretic framework for deciding under uncertainty does not solve problems. Instead, it clarifies the elements required for a solution, and it suggests roles well-suited for different players. To set action levels or exposure limits under uncertainty, we must first decide what levels we would choose or limits we would set without uncertainty. This is not an easy problem, but it is a *different* problem from deciding under uncertainty. Confusing these two problems is perhaps the biggest obstacle to dealing with uncertainty. The key to overcoming this lies in defining the target risk level at which we aim, then applying systematic techniques to address the uncertainty in the underlying dose-response relation.

The type of uncertainty that is quantified when toxicological information is limited or unknown is typically that of knowledgeable experts. Several different expert judgment methods have been applied and have passed muster in peer review. The current state of the art is that a number of methods are accepted, but none has emerged as a universal standard. Assessing loss

rates among stakeholders has not yet been conducted in the context of regulating under uncertainty as described here. However, many approaches could be applied to this task, including contingent valuation, willingness to pay, and consumer preference theory. Research could profitably target this area.

In short, we know how to quantify uncertainty, and considerable experience exists in this arena. We also know how to quantify losses and loss rates, and again considerable experience exists, although not in the particular area of dose-response for exposure guides. Finally, we know how to use these elements in a transparent and rational process of decision making under uncertainty. Given the importance of a transparent, systematic characterization of uncertainty to support exposure guides for health protection, the only remaining question is: why aren't we doing it?

REFERENCES

1. van Noortwijk, J. M., & van Gelder, P. H. A. J. M. (1999). Bayesian estimation of quantiles for the purpose of flood prevention. In B.L. Edge (Ed.), *Proceedings of the 26th International Conference on Coastal Engineering,* Copenhagen, Denmark, 1998; pp. 3529-3541. New York: American Society of Civil Engineers.

2. Fisher, M. L., Hammond, J. H., Obermeyer, W. R., & Raman, A. (1994). Making supply meet demand in an uncertain world. *Harvard Business Review,* May-June, 83-93.

3. Little, M., Muirhead, C. M., Goossens, L. H. J., Harper, F. T., Kraan, B. C. P., Cooke, R. M., & Hora, S. C. (1997). *Probabilistic Accident Consequence Uncertainty Analysis: Late Health Uncertainty Assessment (Vols. 1 and 2).* Prepared for U.S. Nuclear Regulatory Commission and Commission of European Communities, NUREG/CR-6555, EUR 16774, Washington, DC, and Brussels-Luxembourg.

4. NAS (National Academy of Sciences) (1990). *Health Effects of Exposure to Low Levels of Ionizing Radiation: BEIR V.* Committee on the Biological Effects of Radiation, Board on Radiation Effects Research, Commission on Life Sciences, National Research Council. Washington, DC: National Academies Press.

5. EPA (U.S. Environmental Protection Agency) (2006). *Modifying EPA Radiation Risk Models Based on BEIR VII*, Draft White Paper, prepared by the Office of Radiation and Indoor Air, Washington, DC (August 1). Available at http://www.epa.gov/sab/pdf/rac_oria_white_paper_08-01-06.pdf.

6. NCRP (National Council on Radiation Protection and Measurements) (1997). *Uncertainties in Fatal Cancer Risk Estimates Used in Radiation Protection*. Report No. 126. Bethesda, MD.

7. Mettler, F. A., Jr. (2007). *Magnitude of radiation uses and doses in the United States: NCRP Scientific Committee 6-2 analysis of medical exposures*, Program of the 43rd Annual Meeting of the National Council of Radiation Protection and Measurements

(Advances in Radiation Protection in Medicine), April 16-17, Arlington, VA; http://www.ncrponline.org/Annual_Mtgs/2007_Ann_Meet_Prog.pdf.

8. NAS (National Academy of Sciences) (2006). *Health Effects of Exposure to Low Levels of Ionizing Radiation: BEIR VII Phase 2.* Board on Radiation Effects Research, Commission on Life Sciences, National Research Council. Washington, DC: National Academies Press.

9. ICRP (International Commission on Radiological Protection) (1991). *1990 Recommendations of the International Commission on Radiological Protection*, Publication 60, Annals of the ICRP, *21*(1-3). Oxford: Pergamon Press.

10. UNSCEAR (United Nations Scientific Committee on the Effects of Atomic Radiation) (2000). *Sources and Effects of Ionizing Radiation. Report to the General Assembly, with Scientific Annexes. Volume II: Effects.* New York: United Nations.

11. Ehrhardt, J., Hasemann, I., Matzerath, C., Raicevic, J., & Seinhauer, C. (1995). *COSYMA: Health Effects Models*. Report FZKA 5567, Forschungszentrum Karlsruhe, Germany.

12. EPA (2007). *Integrated Risk Information System*, online database, Office of Research and Development, Washington, DC. Available at http://www.epa.gov/iris/subst/0422.htm (accessed October 2007; ammonia reference concentration established in 1991).

13. Goossens L. H. J., Cooke, R. M., Woudenberg, F., & van den Torn, P. (1992). *Probit Functions and Expert Judgment: Protocol and Results*. Report prepared for the Ministry of Housing, Physical Planning and Environment, the Netherlands; Delft University of Technology, Safety Science Group and Department of Mathematics, Municipal Health Service, Rotterdam, Section Environmental Health, the Netherlands.

14. L.H.J. Goossens, R.M. Cooke, F. Woudenberg and P. van der Torn (1993) Protocol for using expert judgement in establishing probit relations of hazardous materials Proceedings of the SRA-Europe Third Conference on Risk analysis: Underlying rationales (eds. Ph. Hubert and M. Poumadere), Paris, France, 16-18 December 1991, published 1993, p.181-189.

1. L.H.J. Goossens, R.M. Cooke, F. Woudenberg and P. van der Torn (1993) Protocol for using expert judgement in establishing probit relations of hazardous materials Proceedings of the SRA-Europe Third Conference on Risk analysis: Underlying rationales (eds. Ph. Hubert and M. Poumadere), Paris, France, 16-18 December 1991, published 1993, p.181-189.

15. Goossens, L. H. J., Cooke, R., Woudenberg, F., & van der Torn, P. (1998). Expert judgement and lethal toxicity of inhaled chemicals. *Journal of Risk Research*, *1*(2), 117-133.

16. Finney, D. J. (1980). *Probit Analysis*, 4th ed. Cambridge University Press.

17. EPA (1990), *National Oil and Hazardous Substances Pollution Contingency Plan; Final Rule (40 CFR Part 300),* Federal Register, *55*(46):8666-8865, March 8.

18. Lees, F. P. (1996). *Loss Prevention in the Process Industries.* 2nd ed. London: Butterworth.

19. CCPS (Center for Chemical Process Safety) (2000). *Guidelines for Chemical Process Quantitative Risk Analysis.* 2nd ed. American Institute of Chemical Engineers. New York: Wiley.

20. Aarts, J. H. E, Mommers, C., & Muijser, H. (2000). *Toxic Effects from Accidental Releases of Hazardous Substances (TEARHS) - Lethal and Non-Lethal Effects in Rats upon Exposure during Short Periods of Time*. Vol. V99.1136, TNO Nutrition and Food Research, Zeist, the Netherlands.

21. HSE (Health and Safety Executive) (1978). *Canvey: An Investigation of Potential Hazards from Operations in the Canvey Island/Thurrock Area.* Her Majesty's Stationary Office, London, UK.

22. IChE (Institution of Chemical Engineers) (1987). *Ammonia Toxicity Monograph*. United Kingdom.

23. Uijt de Haag, P. A. M., & Ale, B. J. M. (2005). *Guidelines for Quantitative Risk Assessments, Purple Book, Part 1: Establishment.* Ministerie van Verkeer en Waterstraat, CPR 18E, Publication Series on Dangerous Substances (PGS 3).

24. Perry, W. W., & Articola, A. P. (1980). *Study to Modify the Vulnerability Model of the Risk Management System*. Prepared for U.S. Department of Transportation, U.S. Coast Guard, Washington, DC.

25. CPR (Commissie voor Preventie van Rampen, or Committee for the Prevention of Disasters) (1992). *Methods for the Determination of Possible Damage to People and Objects Resulting from Releases of Hazardous Materials. Green Book.* SZW The Hague: Directorate-General of Labour of the Ministry of Social Affairs and Employment.

26. Rijnmond (Rijnmond Public Authority) (1982). *Risk Analysis of Six Potentially Hazardous Industrial Objects in the Rijnmond Area, A Pilot Study*. Presented by COVO Steering Committee. D. Reidel Publishing Co.: Dordrecht, Holland.

27. NAS (National Academy of Sciences) (2007). Acute Exposure Guideline Levels for Selected Airborne Chemicals: Volume 6 (2007), prepared by the Committee on Acute Exposure Guideline Levels, Committee on Toxicology, Board on Environmental Studies and Toxicology (BEST), National Research Council of the National Academies; National Academies Press, Washington, DC (prepublication copy); http://books.nap.edu/openbook.php?record_id=12018&page=47.

28. Kurowicka, D., & Cooke, R. M. (2006). *Uncertainty Analysis with High Dimensional Dependence Modeling.* New York: Wiley.

29. Csiszar, I. (1975.) *I*-divergence geometry of probability distributions and minimization problems. *Annals of Probability*, *3*(1), 146–158.

30. Du, C., Kurowicka, D., & Cooke, R. M. (2006). Techniques for generic probabilistic inversion. *Computational Statistics & Data Analysis 50*(5), 1164-1187.