Internal Dose, Uncertainty Analysis, and Complexity of Risk Models

Louis Anthony Cox, Jr.
Cox Associates

503 Franklin Street

Denver, CO, 80218

April, 1999

(Forthcoming in *Environment International*, 1999.)

ABSTRACT: Practitioners and consumers of risk assessments often wonder whether the trend toward more complex risk models, incorporating increasing amounts of biological knowledge and increasing numbers of biologically interpretable parameters, actually lead to better risk estimates. A contrary view might be that the need to estimate more uncertain quantities undermines the advantages of greater descriptive realism so much that the final risk estimates are less certain than the ones traditionally obtained from simpler, less realistic, statistical curve-fitting models. In opposition to this pessimistic view is the widespread common-sense notion that including more information in a risk model can never worsen (and will usually improve) the resulting risk estimates. This paper appeals to mathematical arguments to resolve these conflicting intuitions. First, it emphasizes the fact that risk depends on multiple inputs only through a small number of "reduced quantities" -- perhaps on only one, which would then be defined as internal dose. Thus, uncertainty about risk may have limited sensitivity to uncertainties in the original input quantities. The concept of internal dose and its possible algebraic relations to the original input quantities are clarified using concepts from dimensional analysis. Then, the question of whether greater model complexity leads to better or worse risk estimates is addressed in an information-theoretic framework, using entropies of probability distributions to quantify uncertainties. Within this framework, it is shown that models with greater intrinsic or "structural" complexity (meaning complexity that can not be eliminated by reformulating the model in terms of its reduced quantities) lead to better-informed, and

hence more certain (lower-entropy) risk estimates. The compatibility of this result with results from decision theory, in which expected loss rather than entropy is used as a criterion, is discussed.

Key Words: Uncertainty analysis, derivation graphs, Monte Carlo simulation, dose-response modeling, complexity

 503 Franklin St., Denver, CO, 80218. Ph 303-388-1778; Fax –303-388-0609. tony@cox-associates.com

INTRODUCTION

There has been much thoughtful debate about the most useful definition of "internal dose" for predicting and extrapolating risks. This paper attempts to clarify the problem and to show that, when formulated carefully, the concept of internal dose has strong implications both for the possible algebraic forms of dose-response functions and for managing uncertainties in complex risk models. Section 1 examines the concept of "internal dose", and proposes to define it as a numerical quantity containing the minimal sufficient information needed to determine risk from administered dose. Section 2 investigates the algebraic consequences of this definition, applying the "product theorem" from dimensional analysis (Bhaskar and Nigam, 1990, p. 79) to conclude that internal dose should be expressible as a product of powers of input quantities. This result leads to a novel derivation of the Weibull dose-response function and some generalizations.

Since internal dose typically depends on more than just administered dose, using internal doses in risk models generally requires an increase in model complexity. "Biologically-based" risk assessment (BBRA) models that predict risk based on calculations of internal (e.g., delivered or biologically effective) doses include additional biological parameters (e.g., metabolic rate constants, metabolite elimination rates and partition coefficients) beyond those required in simpler statistical risk models that relate probability of response directly to administered dose. The additional parameters may be very uncertain, or even unknown. Thus, the use of BBRA models raises the question of whether their added complexity leads to risk estimates that are more uncertain, or less accurate on average, than simpler models that give a cruder conceptual approximation to biological reality but that represent fewer uncertain input quantities.

Section 3 examines this question in detail and answers it negatively. The reasoning is roughly as follows. A complex model's output quantities depend on its input quantities only through certain "reduced quantities" (e.g., internal dose) derived from the input quantities. Uncertainty about the reduced quantities is never greater (and is usually less) than uncertainty about the raw inputs -- a result that can be proved formally within an information-theory framework in which entropy is used to quantify uncertainty about random variables (Cover and Thomas, 1991). Moreover, even poor information that is relevant to estimating internal dose (e.g., "noisy" measurements containing large measurement errors, or measurements of dose surrogates that are only weakly correlated with the true quantities of interest) leads to better estimates when it is included than when

it is ignored. Therefore, more complex models that allow (even poor) additional observations to be brought to bear in estimating a dose-response relation tend to produce better risk estimates. Section 4 shows that this conclusion, based on entropy measures, is compatible with conventional "value of information" analysis from Bayesian decision theory. A simple example is presented to illustrate the Bayesian approach.

On the other hand, a model that is more complex than another only in the sense that it contains more uncertain quantities (variables and/or parameters) to estimate, but that does not allow any additional observations, contains what we term "spurious complexity". It can be simplified, e.g., by algebraically combining unobserved quantities, without impairing ability to estimate the observable ones. This simplification process, referred to in this paper as "model reduction", prunes away spurious uncertainties. Failing to use reduced models can lead to increased uncertainties about outputs (e.g., predicted risk) when there are more input quantities to estimate, but this is better regarded as an artifact of improper modeling methodology than as a true relation between model complexity and uncertainty of model predictions.

1. THE CONCEPT OF INTERNAL DOSE

Recall how internal doses are used by risk assessors. Suppose that administered dose creates a probability of response in a particular sex, strain, and species. This causal relationship might be diagrammed as

$$x \rightarrow P$$

or written in algebraic notation as

$$P = f(x),$$

where

x = administered dose (e.g., a time series of concentrations of a chemical in inhaled air)

P = lifetime probability of response (typically estimated from the observed fraction of animals responding in a bioassay experiment)

f denotes a dose-response mathematical function mapping values of x to corresponding values of P

$\rightarrow$ means "affects the value of". More specifically, the value of any variable is modeled as being determined by the values of the variables that point into it. This "derivation graph" notation is explained and used further below.

In practice, it is unlikely that the administered dose is sufficient to completely and uniquely determine the value of P. Given the value of x and additional relevant information (e.g., about the body weight, breathing rate, enzyme and hormone status, individual metabolism, etc.), one should be better able to predict the value of P than would be possible from knowledge of the value of x alone.

To indicate that factors other than the administered dose affect the probability of response, the conceptual model of the dose-response relation might be redrawn as

$$x \rightarrow P \leftarrow z$$

and the algebraic dose-response model rewritten as

$$P = f(x ; z)$$

where z is a vector of individual covariates.

The vector z may be partitioned further, e.g., into observed covariates (such as body weight, age, duration of follow-up period, etc.), unobserved or "latent" variables (such as individual "susceptibility", "frailty", or "tolerance threshold"), and parameters (covariates whose values are the same for all individuals in a population, e.g., the sex and strain of the mice in a particular exposure group). The symbol z will denote any and all quantities other than x that affect the value of P.

A central problem of applied risk assessment is to predict the value of P for any pair of input values, (x, z) after observing the responses of animals i = 1, 2, ..., N satisfying exposure conditions (x_1, z_1), ..., (x_N, z_N), respectively. If the vector z contains many components, then the input-output relation between (x, z) and P may not depend on all the details of the whole vector (x, z), but only on a smaller number of values that can be calculated from (x, z). The technique of "nondimensionalization" (Lin and Segel, 1988, p. 198) reduces the number of parameters and variables in a model to a minimum by re-expressing its equations in terms of dimensionless algebraic combinations of the original

parameters and variables. Group-theory methods can be used to find the minimum number of quantities needed to describe a model and to construct (usually several possible) reduced models that achieve this minimum (Doucet and Sloep, 1992, 314-318).

To illustrate model reduction, suppose that z contains the animal's body weight, BW, and a potency coefficient, β, defined as the average number of interactions of metabolite molecules with target cells per unit of administered dose. Assume that β is a constant (independent of dose) and let A(x) denote the total units of administered dose in x. [If x is an administered dose time series, then A(x) is the area-under-curve obtained by integrating x(t) from the start to the end of the dosing period.] Then the probability of response might depend on the factors that affect it through a formula such as

$$P = f[\beta A(x)/BW],$$

where f is an algebraic function, perhaps $f(y) = 1 - \exp(-y)$. The covariate vector in this example is

$$z = [\beta , BW].$$

For conceptual and notational economy, it is useful to introduce a new, derived variable, y, defined as

$$y = \beta A(x)/BW.$$

The value of y can be calculated from the information in (x, z). But many of the details in (x, z) are irrelevant for the purposes of determining P. Any two pairs of (x, z) values that give the same value of y will determine the same value of P. In other words, P depends on x, β, and BW only *through* the quantity y. Diagrammatically, we have

$$\begin{array}{l} x \rightarrow\ y \rightarrow P \\ \qquad\ \uparrow \\ \qquad\ z \end{array} \quad \text{or, in more detail,} \quad \begin{array}{l} x \rightarrow A(x) \rightarrow y \rightarrow\ P \\ \qquad\qquad\qquad\ \uparrow \\ \qquad\qquad\qquad z = (\beta , BW) \end{array}$$

Given the value of y, no other information (e.g., about the values of x or z) is relevant for determining response probability P. In this example, y could be defined as the *internal dose* on which P depends. How well P can be predicted depends only on how well y can

be estimated, and not on the rest of the information in the time series x or the parameter vector z.

Defining internal dose as the minimum information that is sufficient to derive or calculate P provides a concept that can be extended to large risk models. To do so, we introduce a mathematical abstraction of the structure of a model called its *derivation graph.* This conceptual apparatus and closely related ideas and extensions have been used by many modelers using various terms, such as influence diagrams (Schachter, 1990), causal graphs (Pearl, 1998; Galles and Pearl, 1997), Bayesian networks (Pearl,1988, Jensen 1996), and structured modeling graphs (Geoffrion, 1996). A *derivation graph* for a risk model is a directed graph, which we denote by G = (V, E), showing how P is derived from other quantities. In such a graph, V is a set of vertices or "nodes", representing the quantities (parameters or variables) of the model. E is a set of directed edges (“arrows”) between pairs of nodes (i.e., E is an asymmetric binary relation, $E \subset V \times V$). Unless otherwise stated, G is assumed to be acyclic, i.e., it has no directed cycles (Schachter, 1990). The interpretation of the graph is as follows:

(i) *The value of any quantity v in V is determined by the values of the quantities that point into it,* i.e., that have directed edges ending at v, if there are any. Thus, one can associate with each node a formula or algorithm that determines its value from the values of its inputs. From this perspective, the graph structure reflects the order in which values of quantities can be calculated and propagated (Shenoy, 1992). In a non-deterministic model, the conditional probability density function (pdf) for the value of a quantity is determined by the values of the quantities that point into it. If all quantities are scalars (rather than time series, spatial density functions, or other mathematical objects) then the derivation graph is called an *influence diagram* (Schachter, 1990).

(ii) Nodes with only outward pointing edges constitute the model's *input quantities.* Nodes with only inward pointing edges are *derived quantities.* They constitute outputs from the model. *Intermediate nodes,* having both inward and outward arrows, are also outputs, in that their values are determined by the input quantities. The administered dose variable x is an input variable and P is a derived output quantity.

Although not required by this directed acyclic graph ("dag") formalism, it will often be convenient to interpret derivation graphs as showing *causal dependency* relations among quantities, as defined by Iwasaki and Simon, 1994. These relations "are based on the notion of *manipulability* of a variable by *intervention* into the system. When we say that 'variable B depends on variable A', which we will write as

$$A \rightarrow B,$$

we mean that there exists some mechanism that determines the value of B from that of A, whether it be a physical law, an actual device, or some other kind of communication path between the two variables. Thus, everything being equal, a change in A will necessarily result in a change in B." In our dag formulation, a change in A does not *necessarily* change the value of B, since the relation between them may be stochastic or may be such that B has the same value for several different values of A. For example, if A has caused an irreversible change in B, then a further change in A may not change the value of B further. Nonetheless, interpretations of dags in which changes in input variables propagate among the other variables in a model along the arrows linking them are useful for many purposes in risk analysis, as well as in other areas of applied science and engineering. As noted by Iwasaki and Simon (*ibid*, p. 146) "The method for determining causality based on propagation of disturbances is generally consistent with, but somewhat less general than, the more formal method of causal ordering" developed in their paper. This method, drawing on earlier econometric work on structural equation modeling, is based on dependency relations among variables in systems of structural equations (Bentler, 1992). It orders variables in such a way that earlier ones determine the values of later ones, i.e., the values of later ones can be derived from the values of earlier ones. Thus, causal orderings of variables in models can provide an interpretation of, and motivation for, derivation graphs.

To construct an internal dose quantity in a risk model with derivation graph G, define Y as a subset of quantities with the following properties:

(i) *Conditional Independence:* The value of P is *conditionally independent* of the values of all other quantities given the values of the quantities in Y. Thus, learning the value of any other quantities (i.e., in V - Y) would not change P's conditional pdf. For example, $f_P(P \mid Y, v) = f_P(P \mid Y)$, for any v in V - Y, where f_P denotes the

probability density function of P and "|" means "given the information" or "conditioned on".

(ii) *Sufficiency:* The values in Y suffice to determine the conditional pdf of P (or the value of P, in a deterministic model). This property implies, for example, that there are no explanatory variables in V that affect P but that do not affect the variables in Y.

(iii) *Minimality:* No proper (smaller) subset of Y has properties (i) and (ii).

The information in Y is minimally sufficient to determine uniquely, via the risk model, the value or conditional pdf of the response probability P.

Definition: *In a risk model with derivation graph G = (V, E) having response probability P as a derived output, a quantity is a measure of internal dose if and only if it satisfies properties (i) through (iii).*

According to this definition, a set Y satisfying properties (i) through (iii) is a measure of internal dose if and only if it contains a single quantity. The adequacy of a proposed internal dose measure y can be assessed statistically, if y can be measured or calculated, by testing the hypothesis that P is statistically independent of x (and other inputs), given the value of y. In a deterministic model, the conditional pdf of P is a single probability number when the values of the input quantities are uniquely specified. Thus, deterministic risk models are included as special cases.

Internal dose measures may not be unique, even in models where they exist. For example, if y is an internal dose measure, then so is any monotonic transformation of y. The next section introduces units of measurement and measurement scales to eliminate this ambiguity. In addition, consider the following simple model, and assume that there is a one-to-one, deterministic relation between x and y:

$$x \rightarrow y \rightarrow P$$

Then either x or y can be used as a measure of internal dose. However, y is usually determined by the values of several other quantities (e.g., administered dose and metabolic and pharmacokinetic parameters), and its value does not uniquely determine theirs. In

general, if there are multiple measures of internal dose in a risk model, we prefer the one that is closest to P in the derivation graph (excluding P itself).

2. EXISTENCE AND ALGEBRAIC FORM OF INTERNAL DOSE MEASURES

Given a deterministic risk model P = f(x ; z), when does there exist a derived scalar quantity (an "internal dose")

$$y = g(x, z)$$

such that y uniquely determines the value of P? More generally, in a stochastic risk model defined by the conditional probability density function

$$f_P(p \mid x, z),$$

we wish to determine when there exists a scalar quantity y = g(x, z) such that

$$f_P(p \mid x, z) = f_P(p \mid y) = f_P(p \mid x, z, y).$$

That is, P depends on x and z only through the scalar y. If such a y exists, then it is important to determine the algebraic form of the function g(.) by which it is calculated from x and z.

The questions of existence and algebraic form can be addressed constructively in the important special case in which all input quantities are physical quantities measured on ratio scales, i.e., scales that are unique up to choice of a dimensionally consistent system of measurement units. For example, concentration, measured in mass of chemical per unit volume of inhaled air, is a ratio scale variable. The ratio of any two concentrations is independent of the choice of measurement units, e.g., mg/liter, ounces per cubic foot, etc. Suppose that P is determined by a set of primary physical quantities, each measured on a ratio scale. Examples of such primary quantities might include the concentration and duration of an inhalation exposure (if the exposure time series is completely specified by these two quantities), the mass of the exposed animal, its respiratory volume, breathing rate, and so forth. Let y be a secondary quantity, also measured on a ratio scale, derived from the primary quantities by some algebraic combining function. Thus, if the primary quantities are summarized in a vector (x, z), where x contains the exposure information

(e.g., the concentration and duration of exposure) and z contains all other primary quantities (covariates) needed to determine y, then y is assumed to be derived from (x, z) via some formula

$$y = g(x, z).$$

So far, all that has been added to the initial formulation is that y and the components of x and z are all measured on ratio scales, in a dimensionally consistent system of units. If the mapping defined by g is well-behaved (e.g., once-differentiable), this assumption of ratio scale measurements suffices to establish the algebraic form of the combining function, g(.). The product law of dimensional analysis states that *"The measurement of a secondary quantity is expressible as a product of powers of measurements of primary quantities"* (Lin and Segal, 1988, p. 206; Bhaskar and Nigam, 1990, p. 106). Thus, if x has m components and z has n components, then the formula for y is:

$$y = C(x_1)^{a_1}(x_2)^{a_2}\ldots(x_m)^{a_m}(z_1)^{b_1}(z_2)^{b_2}\ldots(z_n)^{b_n}$$

where C is a constant (usually set equal to 1, which fixes the measurement scale of y) and the exponents are constants to be determined. Intuitively plausible formulas such as

$$y = \beta A(x)/BW = \beta(\text{concentration})(\text{duration of exposure})(\text{body weight})^{-1}$$

are special cases of the product law.

The product law is not directly applicable to risk models since the output P, being a probability, is not measured on a ratio scale. However, there is a one-to-one correspondence between probabilities and the arrival rates of (adverse) events in a Poisson process, operated for one period of time, that would produce one or more arrivals with the given probabilities. Namely, if p is a probability number, then the derived quantity

$$r = -\ln(1 - p)$$

is the average arrival rate (measured in expected number of arrivals per unit time, a ratio scale) that would produce at least one arrival in unit time with probability p. This transformation of the probability axis provides for ratio-scale measurement of event likelihoods. Therefore, the product law applies to r, and any risk model that uniquely

derives P [and hence the quantity -ln(1 - P)] from primary ratio-scale quantities (x, z) must have the product form:

$$-\ln(1 - P) = C(x_1)^{a_1}(x_2)^{a_2}\ldots(x_m)^{a_m}(z_1)^{b_1}(z_2)^{b_2}\ldots(z_n)^{b_n}.$$

Inverting yields the final formula for P:

$$P = 1 - \exp[-C(x_1)^{a_1}(x_2)^{a_2}\ldots(x_m)^{a_m}(z_1)^{b_1}(z_2)^{b_2}\ldots(z_n)^{b_n}].$$

This establishes the following result.

Proposition 1: *Under the conditions of the product law of dimensional analysis, if the response probability P is completely and minimally determined by (i.e., derived from) one or more ratio-scale input quantities, then a numerical measure of internal dose exists. It is given by a product of powers of these input quantities. If this internal dose measure is denoted by y, then the form of the dose-response relation is*

$$P = 1 - \exp(-Cy).$$

Corollary: *Suppose that the response probability of an animal exposed to dose x at time 0 (e.g., via injection) depends only on x and on the length of the follow-up period, say, T. Assume that x is measured on a ratio scale (e.g., mg of chemical/kg of animal's body weight). Then the time-dose-response function follows a Weibull survival model, i.e., it is of the form:*

$$P = 1 - \exp(-Cx^aT^b),$$

where a, b, and C are constants.

Generalizations to more complicated dosing scenarios follow immediately. For example, if dosing is sustained for D days, and if P is determined by D as well as x (now measured in mg/kg/day, for example) and T, then the dose-response model would become

$$P = 1 - \exp(-Cx^aT^bD^c),$$

where the constants C, a, b, and c must be determined, e.g., by nonlinear regression applied to sample data for P, x, T, and D.

In summary, dimensional analysis provides useful partial answers to the questions of whether an internal dose measure exists (it does if P is determined by a set of physically measurable, ratio-scale input quantities) and the form that the internal dose quantity must take (a product of powers of the primary input quantities). These results share the usual limitations of dimensional analysis. Specifically, they do not identify the group of primary input quantities needed to form a complete risk model for calculating P. (Indeed, the set of quantities that are treated as primary is typically unique only up to algebraic combinations that preserve group structure; see Doucet and Sloep, 1992). Such causal knowledge must come from biological or other considerations . Bhaskar and Nigam (1990) discuss in detail for simple physical models the interplay between the mathematical constraints implied by dimensional analysis and the qualitative knowledge of system dynamics and of the interdependencies among variables needed to deduce specific formulas describing input-output behaviors.

3. UNCERTAINTY ANALYSIS OF COMPLEX RISK MODELS

That P may depend on a large number of input quantities only through a single intermediate quantity raises the possibility that the value of P has only a limited sensitivity to uncertainties in the values of model inputs. Input uncertainties may be represented in a derivation graph such as Figure 1a.

$$\underline{x} \rightarrow P \leftarrow z \rightarrow \underline{z^*} \leftarrow e$$
$$\downarrow$$
$$\underline{n(x)} \leftarrow \underline{N(x)}$$

Figure 1a: *A Model with Imperfectly Measured Covariates*

In this and subsequent derivation graphs, quantities that can be observed, including derived quantities that are deterministic functions of observed quantities, are underlined. Theoretical constructs or quantities that can not be observed are indicated in plain text. Thus, in this diagram, the dose x and the empirical quantal response at that dose, n(x)/N(x), are each observable. [Notice that this is a nondeterministic model: P and N(x) do not uniquely determine n(x), but they do uniquely determine the binomial pdf of n(x).]

Similarly, the covariate (and/or parameter) vector z can not be observed without error. What can be observed instead is a quantity z*, which depends on both the true (but unknown) value of z and also on an unobserved source of uncertainty, e, interpreted as noise or as measurement error. In simple statistical models, the relation between z* and z is often assumed to be

$$z^* = z + e,$$

where all three quantities are vectors of the same dimension. Another possibility is that z* is a function of z (for example, z* might contain the aggregate quantity of total metabolites produced whereas z could be a higher-dimensional vector that reports the amounts formed of each specific metabolite). In general, z* contains less information than z (in a sense soon to be made precise), either because it is contaminated with exogenous error e or because it is a function of z, or both.

A question of practical interest is whether z* can contain so much uncertainty that risk estimates of P (given x) will be less uncertain if they ignore z* than if they include it in estimating P. This type of question arises in discussions of the practical value to policy makers of incorporating information from PBPK and other complex "biologically-based" risk models. In such a model, there are typically many more uncertain quantities that must be estimated from limited data than in simpler (e.g., purely statistical) risk models. On the other hand, there may be additional measurements that can be taken (e.g., on the approximate values of partition coefficients or cytotoxic potencies) that can help to validate the complex model and raise confidence in its predictions. Which effect dominates? Is uncertainty about the relation between x and P increased or decreased by using a more complex model with partial but imperfect (perhaps highly uncertain) estimates of the additional parameters? These questions are addressed next.

Influence Diagram Models with Observations and Uncertain Inputs

A clear answer can be given if uncertainty and information are formalized as entropies of probability density functions and if derivation graphs are specialized to be influence diagrams. Recall that an influence diagram is a derivation graph in which all quantities are scalars or finite-dimensional vectors. The joint pdf for values of model quantities is represented by storing at each derived node its conditional pdf, specifying the conditional probability density for each of its possible values, given each possible

combination of values for its inputs. Marginal pdfs are stored at input nodes. Any joint pdf may be expressed (or "factored") in terms of marginal and conditional pdfs for the input and output variables, respectively (Schachter, 1990). For practical computational purposes, each variable is typically assumed to have a finite number of possible distinct values (e.g., continuous quantities are rounded to a few significant digits). We will assume discrete variables except where otherwise stated (e.g., Proposition 3).

To illustrate the main ideas before stating general results, consider the schematic representation of uncertainty in a model with a PBPK component, shown in Figure 1b.

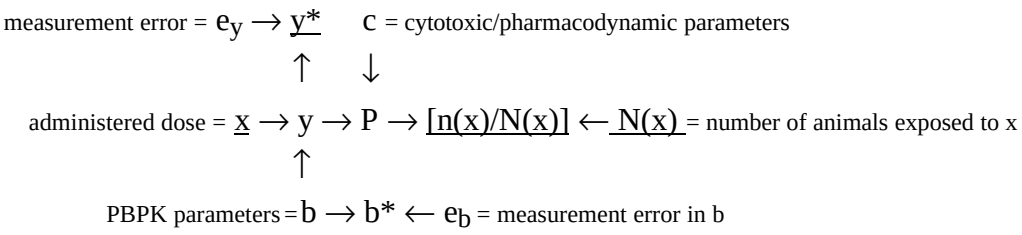

Figure 1b: *A More Complex Derivation Graph with Uncertainties*

In this model, administered dose (x) is assumed to be described by a vector of numerical parameters, such as concentration, hours per day, days per week, and number of weeks of dosing. These quantities and some pharmacokinetic parameters (b) uniquely determine a vector of internal metabolite doses (y) (e.g., described as a vector of AUCs of metabolites in specific physiological compartments). In contrast to Section 2, it is not important here whether y is a scalar or a vector, so long as it contains all the information from x and b that is relevant for predicting P. The pharmacokinetic parameters are not directly observable, but some function of y can be measured with error (y* denotes the measured quantity and e_y the measurement error), while some function of the parameters b can also be measured with error (b* and e_b, respectively). Finally, P is uniquely determined from y and from the values of some pharmacodynamic parameters (c), which may include cell kinetics and cytotoxic potency parameters, mutagenic potency parameters, and so forth. In the diagram, it is assumed that the values of these quantities can not be measured, even with error. At best, only their prior joint frequency distribution is available for purposes of risk assessment. As usual, the unobservable P and the experiment design variable N(x), i.e., the sample size at dose level x, uniquely determine the pdf of n(x), the number of animals with tumors at dose level x, and hence the pdf of the observed quantal response, n(x)/N(x).

In such an influence diagram model, information from observed quantities can be propagated throughout the diagram to obtain posterior (conditioned) pdfs for quantities of interest -- in this case, the value of P corresponding to any specified value of x. Several practical computational algorithms are available to accomplish such "fusion and propagation of evidence"; see, for example, Laurtitzen and Spiegelhalter (1988), Pearl (1988), Schachter (1990), Shenoy (1992). The result is the one that would be obtained by conditioning the full prior joint pdf for all model quantities on all observations using Bayes' rule, and then extracting [by integrating over all other quantities, i.e., "marginalizing out" (Shenoy, 1992)] the pdf for P. However, influence diagram algorithms are often practical for applied work and are available in commercial decision analysis software packages such as HUGIN™ (Jensen, 1996) as well as in a variety of academic and research software. By contrast, straightforward Bayesian computation of the posterior pdf by direct application of Bayes' rule is often computationally unmanageable because it requires multivariate integration or summation over a combinatorially explosive number of possible joint values of model quantities. Henceforth, we assume that Bayesian conditioning of unobserved output quantities (including P) on the values of all observed quantities can be carried out using readily available influence diagram algorithms.

Comparing Derivation Graph Models

Figure 1b seems more complex than Figure 1a. This intuition can be sharpened by defining the following comparative binary relations between derivation graphs.

Definition: *Let G and H be two derivation graphs. Then*

(i) G is an <u>extension</u> of H if H is a subgraph of G and G contains input or output nodes not in H.

(ii) G is an <u>elaboration</u> of H if G has all the same input and output nodes as H, but also contains additional intermediate nodes.

(iii) G is an <u>expansion</u> of H if G is an elaboration or an extension of H, or both. Conversely, if G is an expansion of H, then H is a <u>contraction</u> of G.

(iv) G is a <u>reduction</u> of H if G is a contraction of H and G and H have the same observable quantities.

(v) G is more informed than H if the observable quantities of H are a proper subset of the observable quantities of G.

With this terminology, refinements of the notion of comparative complexity between models can be made. The following relation is useful in discussing whether greater model complexity increases or reduces uncertainty about its predictions.

(vi) G is at least as structurally complex as H if G is a more informed expansion of H or if G is an expansion of H and G and H have the same observed quantities.

To illustrate these definitions, note that

(i) $\underline{x} \rightarrow y \rightarrow P$ is a more-informed extension of $\underline{x} \rightarrow y \rightarrow P$

$$\downarrow$$

$$\underline{y^*} \leftarrow e_y$$

(ii) $\underline{x} \rightarrow P$ is a reduction of $\underline{x} \rightarrow y \rightarrow P$; neither is more informed.

If G and H are two derivation graph models for the risk variable P [or for its observable output, n(x)/N(x)] and if G is at least as structurally complex as H, then which model is better able to predict P? Proposition 3 below answers this question.

Reducing Derivation Graph Models

This section introduces a concept of *model reduction,* motivated in part by the use of reduced-form models in econometrics to replace causally interpretable structural equations and to improve estimation (Judge *et al.*, 1985). It is useful for analyzing the information about risk (P) provided by a risk model and a set of observations on some of its quantities.

An unreduced model typically contains quantities that are not essential for determining (or constraining) the values of its observable quantities. Reducing a derivation graph model requires one to compose (via function composition) the input-output mappings of its nodes, so that the effects on the model's observable quantities that would be transmitted by the deleted nodes are still represented in the reduced model. For example, if the diagram

$\underline{x} \rightarrow y \rightarrow \underline{P}$

corresponds to the deterministic algebraic model

$y = g(x), P = f(y),$

then the reduced diagram

$\underline{x} \rightarrow \underline{P}$

would correspond to the *reduced model*

$P = f(g(x)) = (f \circ g)(x),$

where $\circ$ denotes function composition. (In practice, the response probability P is not observable, so this is only an illustration.) If the mappings from x to y and from y to P were nondeterministic, specified by conditional pdfs, then they could still be composed to yield the pdf of P conditioned on x. For example, suppose that Pr(y | x) and Pr(P | y) are defined by the following matrices (probabilistic mappings):

Pr(y | x):	x = 0	x = 1	Pr(P | y):	y = 0	y = 1
y = 0	1	0.5	P = 0	1	0.2
y = 1	0	0.5	P = 1	0	0.8

The composition of these mappings is the matrix product Pr(P | x) = Pr(P | y)Pr(y | x):

Pr(P | x):	x = 0	x = 1
P = 0	1	0.6
P = 1	0	0.4

Given a probability distribution column vector for the input x, say p (with non-negative components summing to 1), the column vector Pr(P | x)p gives the corresponding probability distribution for P. This process of replacing unobserved intermediate quantities (such as y) with direct mappings from observable inputs to observable outputs will be called *model reduction.* A model is said to be *fully reduced* if no further reductions are possible, i.e., if it contains only observable quantities for its intermediate and output nodes.

Reducing a model usually yields a smaller set of *reduced parameters* (consisting of algebraic combinations of the parameters in the original model) that still fully determine the input-output mappings and the other constraints among observable quantities implied by the original model. For example, in the model [P = f(y), y = g(x)], if the original model equations are

$$g(x) = a + bx \quad \text{and} \quad f(y) = c + dy,$$

then the four parameters a through d in the original model (referred to as the "structural" model parameters in econometrics) may be combined to form the reduced model (Judge *et al.*, 1985):

$$P = a' + b'x$$

having the two reduced parameters

$$b' = bd \text{ and } a' = (c + ad).$$

Notice that only the reduced parameters can be estimated from the observable quantities x and P. The four original parameters are not uniquely identifiable from the observed quantities, no matter how many observations are available.

Entropy Inequalities

Model reduction prunes away quantities that do not affect observable outcomes, even if they have clear, perhaps conceptually important, theoretical or biological interpretations. Uncertainties about those quantities are superfluous, in that resolving them will not improve ability to predict risk (or other observable output quantities) from the input quantities. In general, uncertainty about reduced parameters is less than uncertainty about their constituent (structural) parameters.

To make this claim precise in the context of influence diagram models, a quantitative measure of the "uncertainty" about a discrete random variable is needed. Recall that the *entropy*, measured in bits, of a discrete random variable X is defined as

$$ent(X) = -\Sigma_x p(x) \log_2 p(x),$$

where x ranges over all possible values of x and p(x) = Pr(X = x) is the probability density assigned to value x (Cover and Thomas, 1991). If $-\log_2 p(x) = \log_2[1/p(x)]$ is interpreted as the amount of information gained when it is learned that X has value x, then ent(X) may be interpreted as the expected amount of information that will be gained by learning the value of X. It also gives a lower bound on the average number of binary questions with equally likely answers (each question asking whether the value of X belongs to a specified subset of values) required to determine the value of X when an optimal questioning strategy is used (Cover and Thomas, 1991). Notice that this number depends only on the number of distinct values of X and their probabilities, but not on their magnitudes or on whether they are closely spaced. For many purposes, ent(X) is a uniquely satisfactory measure of the uncertainty about a discrete random variable (Buckley, 1985). It is maximized when all of the possible values of X are equally likely (i.e., no information is available to help determine the correct value) and is equal to zero when the value of the random variable X is known (probability = 1). The next two results use entropy as a quantitative measure of uncertainty about model quantities.

Proposition 2: *Let b be a vector of quantities (parameters or variables) in an influence diagram model. Let r(b) be the vector of reduced quantities (algebraic combinations or other functions of b) obtained by reducing the model. Then the uncertainty about r(b) is less than or equal to the uncertainty about b, i.e.,*

$$ent[r(b)] \leq ent(b).$$

Proof: This inequality follows from the following general result in information theory. If X is any discrete random variable, then the entropy of any function of X is less than or equal to the entropy of X (Cover and Thomas, 1991, p. 43).

Intuitively, the value of r(b) may fail to distinguish between two or more values of b (i.e., r(.) may be many-to-one), so that less information is needed to uniquely determine the value of r(b) than to uniquely determine the value of b.

Corollary: *If response probability P depends on a vector of input quantities, b, only through an internal dose quantity that is a scalar function of them, say, y = g(b), then*

ent(P) £ ent(y) £ ent(b).

This corollary gives specific meaning to the claim that uncertainty about internal dose is less than or equal to uncertainty about the input quantities (e.g., administered dose and physiological parameters) used to derive it.

An objection to entropy as a measure of uncertainty about a random variable is that its value does not depend on the range of possible values, but only on the number of distinct values and their probabilities. Entropy is often interpreted as a measure of the "surprise" that one might feel on learning the specific value of a random variable, but it does not try to express whether the surprise reflects a value near or far from the expected value. Yet, intuitively, a random variable that is uniformly distributed between 0 and 100 is "more uncertain" than one that is uniformly distributed between 0 and 0.01 because the range of its possible values is larger. (Note, however, that this intuition implies that "uncertainty" is not invariant under changes of scale; see e.g., Goodman, 1987.) By dropping the restriction to discrete random variables, this element of "uncertainty" can also be incorporated into a quantitative measure. Namely, for a continuous random variable, the *differential entropy* is defined as

$$ent(X) = -\int p(x)\log_2 p(x)dx$$

where p(x) is the probability density of x and the integral is taken over the entire set of possible values of X. For a continuous random variable X,

$$ent(aX) = ent(X) + \log_2|a|,$$

where |a| is the absolute value of a (or of the determinant of a, if X is a vector and a is a matrix). Thus, scaling up the range of X increases differential entropy ($|a| > 1$), while scaling the range of X down reduces differential entropy ($|a| < 1$). Interpretively, ent(X) is related to the logarithm of the volume of the support set (the set of values with positive probability measure) for X (Cover and Thomas, 1991, p. 226). For a scalar random variable uniformly distributed over an interval of length L, $2^{ent(X)} = L$.

The next result applies to models with probabilistic mappings, including deterministic models as a special case. It also holds for continuous as well as for discrete random variables, with ent(X) being interpreted as differential entropy if X is continuous.

Proposition 3 (Entropy Inequality): *Let G and H be two influence diagram models for P, and suppose that G is at least as structurally complex as H. Let ent(P | G) denote the entropy of the pdf of P after conditioning on all observable quantities in G, and similarly for ent(P | H). Then*

$$ent(P \mid G) \leq ent(P \mid H),$$

with equality if and only if P is statistically independent of the observations in G, given the observations in H.

Proof: For any two random variables X and Y, ent(X | Y) £ ent(X), with equality if and only if X and Y are independent (i.e., "conditioning reduces entropy") (Cover and Thomas, 1991, Theorem 16.1.5, p. 483). Let X = (P | H), i.e., it is the random variable P with its pdf conditioned on all observations in H. Let Y denote any observations in G that are not in H. These interpretations give Proposition 3. (If Y is empty, then G and H contain the same information and equality holds in the proposition.)

Proposition 3 implies that, if one influence diagram model for P is more complex than another, in that it contains more observable quantities, then the more complex model will produce lower-entropy (less uncertain) predictions of P after conditioning on all observations. On the other hand, if one model for P contains the same observations as another but includes more derived, non-observable quantities, these do not increase uncertainty about P. They may be eliminated via model reduction without affecting calculations of the pdf of P. Thus, a structurally simpler model, in which some observations are ignored in order to reduce model complexity, does *not* lead to better (more certain) risk estimates. The extra complexity is worthwhile as long as it carries with it an opportunity to incorporate more relevant observations into estimation of the dose-response relation.

That an influence diagram model with more observed quantities should always lead to conclusions about P that are at least as certain as those from a less informed model is intuitively appealing. The practical importance of this result is that it implies that *observed*

quantities that are relevant for estimating P should be included in the estimation process, even if they contain very large measurement errors. Doing so can never make the estimate of P worse (i.e., make its posterior pdf, after conditioning on observed quantities, have higher entropy) than it would be if some observations were ignored. This is in contrast to the recommendations from econometric procedures that seek to minimize mean square prediction error. The squared error loss criterion conventionally adopted in econometrics often prescribes that a subset of regressors should be used in preference to the complete set of regressors (e.g., Judge *et al.*, 1985, p. 866). Including independent variables that are measured with sufficiently great error variance can increase the error sum of squares for the estimated regression coefficients. However, to reduce uncertainty about the dependent variable, P, as measured by the entropy of its posterior distribution, all independent variables should be included, unless P is conditionally independent of some of them, given the rest. In this case, the variables that are conditionally irrelevant for predicting P may be dropped without affecting the posterior distribution of P (Galles and Perl, 1997).

As an example, in Figure 1b, suppose that P is estimated either by its full posterior pdf after conditioning on observed information, or by some functional of this posterior pdf (e.g., its mean, median, or mode). Then conditioning on the observed value of b*, which is an error-prone estimate of the true but unknown parameter vector b, will always reduce the entropy of P, no matter how broad is the distribution of the error term e_b, so long as b and b* are not statistically independent.

4. COMPATIBILITY WITH DECISION-ANALYTIC MODELS FOR THE VALUE OF INFORMATION

Influence diagrams for risk management applications typically include decision variables, i.e., quantities whose values are set by a risk manager or policy maker. Uncertainty about model quantities may then be viewed as important only to the extent that it jeopardizes good decision-making. To formalize and quantify this concept of *decision-relevant uncertainty*, let *A* denote a set of possible alternative actions or decisions (e.g., occupational exposure standards for exposure to a chemical) and let *B* denote the set of possible parameter vectors for an influence diagram model of risk. Introduce a loss function (a real-valued bivariate function $L: A \times B \rightarrow R$) defined as follows:

L(a, b) = loss from decision or action *a* when the parameters of the risk model are *b, " a Î A, " b Î B.*

The number *L(a, b)* indicates the undesirability of the outcome in which policy decision *a* is implemented when the true dose-response relation corresponds to parameter vector *b*. Under most normative theories of decision making, loss is measured on a difference scale (unique up to choice of origin and scale) that may be normalized to run from 0 for the best outcome to 1 for the least desirable outcome. Equivalently, the *(a, b)* pairs may be assigned utilities scaled to run from 0 for the least-preferred pair to 1 for the most-preferred pair. Since decisions that maximize expected utility minimize expected loss, either formulation may be used (DeGroot 1970, p. 123).

Given a set of observations, suppose that a decision from *A* is made that minimizes expected loss with respect to the posterior pdf of *b*. The procedure is that, first, the impact of the observations on the probable value of *b* is computed, perhaps using commercial influence diagram software. The result is a posterior pdf for *b*. This pdf induces a unique pdf for *L(a, b)* for each risk management decision or act, *a Î A.* A *minimum expected loss* or *maximum expected utility* decision rule chooses *a* so that the induced pdf of *L* has the smallest mean value among all pdfs that can be obtained by selecting acts in *A* (DeGroot 1970). Such an *a* is also called a *Bayes action* (Shao, 1999, p. 195). The *value of information* can be measured by the reduction in the expected loss, i.e., in the expected value of *L*, when a Bayes action is taken after using the information to compute the pdfs of *L* for different acts in A (Laffont, 1990, Chapter 4). The question then arises of whether risk models with greater structural complexity lead to better decisions, that is, to Bayes actions with smaller expected losses.

A principle results of modern decision theory gives an affirmative answer. Under most normative models of how decisions "ought" to made, including models that justify Bayes action (maximum expected utility) decision-making, *the expected value of information is nonnegative*. Conditioning the pdf of *b* on more information before choosing an action leads to expected losses that are never greater than those based on less information.

The intellectual framework of modern normative decision theory, within which this result is formalized and proved, is large and sophisticated, defying brief summary. See Quiggin (1993) for an excellent review of both “standard” (subjective expected utility, or

"SEU") theories and some alternatives, including controversial ones in which the value of information need not be positive. Without attempting to recapitulate the SEU framework, we note that most widely accepted theories of "rational" behavior imply that a rational decision maker will prefer to condition on all available relevant information, no matter how imperfect, because doing so is expected to reduce decision loss. This conclusion – that information has non-negative value – holds regardless of the decision maker's specific loss function or utility function. Like the entropy inequalities, this information value inequality implies that even poor observations, in which observed values only weakly associated with the true values, will not lead to worse risk management decisions.

Applied to the influence diagram setting, this result may be stated as follows.

Proposition 4 (Expected-Loss Inequality): *Let G and H be two influence diagram models with the same uncertain quantities b and decision variables a. Suppose that G is at least as structurally complex as H. Let* $E_{b|G}[L(a, b)]$ *denote the expected loss associated with the pdf of b after conditioning on all observable quantities in G, assuming that a Bayes action is selected, and define* $E_{b|H}[L(a, b)]$ *similarly. Then, under the usual assumptions of subjective expected utility (SEU) theory (Quiggin, 1993),*

$$E_{b|G}[L(a, b)] \pounds E_{b|H}[L(a, b)].$$

Thus, the expected loss from an optimal decision made using model G is less than or equal to the expected loss from an optimal decision made using model H.

Proof: Laffont (1990), Section 4.2, Theorem 1, gives the following result. Information structure 1 is finer than information structure 2 if and only if, for *any* prior probability distribution and for *any* utility function, the expected utility of a Bayes action conditioned on information structure 1 is at least as great as the expected utility of a Bayes action conditioned on information structure 2. An *information structure* is defined as a function mapping "states of nature" (which we interpret as specifying the values of all quantities in a model) into a set of observations. One information structure is *finer than* another if the set of states that could have generated an observation in the first is a subset of the set of states that could have generated the corresponding observation in the second, for all possible pairs of observations from the two information structures. Letting the first information structure correspond to the set of observations in G and the second correspond to the set of observations in H establishes our Proposition 4.

For expected utility decision-makers, increased structural complexity in a risk model not only leads to more certain inferences about risk (Proposition 3), but also to

decisions with lower (or equal) expected losses. In this sense, decisions based on G are truly better-informed than decisions based on H.

As an important application, let *A* be the set of possible unit risk estimates for a carcinogen, so that a "decision" now refers to making (or choosing) a unit risk estimate. Interpret *b* as the true unit risk, defined as the slope of the dose-response curve as dose approaches zero. Then *L(a, b)* is the loss associated with estimating the unit risk estimate as *a* when its true value is *b*. If *a* is chosen to minimize expected loss, then conditioning the estimate on more observations will reduce (or, if the information is irrelevant to *b*, leave unchanged) the expected loss. This is true for *any* loss function. It is true even if the observations incorporated into the estimate (via conditioning) are very noisy or very imperfect surrogates for the biological quantities that determine *b*. Therefore, a complex model that incorporates more relevant observations and measurements will lead to better (or to unchanged) risk management decisions, as evaluated by the expected loss criterion, than a simpler model that ignores the additional information.

Example

Figure 2 shows an influence diagram with a single decision node, represented by x_2. The observed quantities are as follows:

- x_1 = vector of dose levels to which a cohort of workers was exposed in the past
- N_1 = a vector giving the number of workers in each dose group
- n_1 = a vector giving the number of workers in each dose group that contracted a disease
- N_2 = size of a currently exposed public population
- x_2 = enforced standard specifying the maximum allowed dose to which the public population may be exposed (a decision variable).

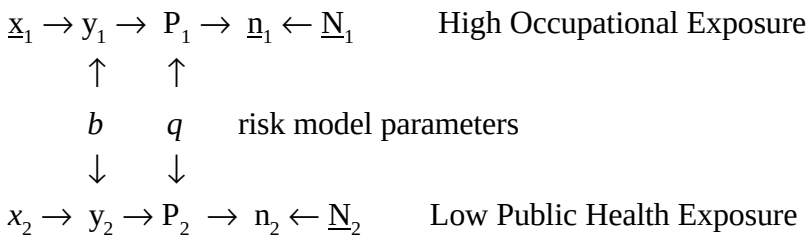

Figure 2: *A Risk Management Example Requiring Inference, Extrapolation, and Decision-Making in a Model with Uncertain Input Values*

The remaining quantities are not observed. They include y_1 and y_2, the internal doses created by administered doses x_1 and x_2, respectively; P_1 and P_2, the response probabilities due to these doses; *b* and *q*, the parameter vectors that determine the internal dose from the administered dose and probability of response from the internal dose, respectively; and n_2, the number of cases of the disease that will occur in the exposed public population.

If x_2 had already been fixed and n_2 observed, then the inference problem would be to combine the evidence from the two studies, 1 and 2, to determine the probable values of the unobserved quantities. In Figure 2, however, such pooling or meta-analysis is unnecessary, since only the data from the x_1 study is available. Instead, the inference problem is to use the evidence contained in x_1, n_1, and N_1 to estimate *b* and *q* in order to determine the dose-response relation between x_2 and P_2. The decision problem is then to set a level for x_2 that will minimize expected loss, given all relevant information about P_2.

For simplicity and concreteness, suppose that *b* is a scalar and that the node formulas are as follows:

$y_i = bx_i$ for i = 1, 2
$P_i = 1 - \exp(-qy_i)$ for i = 1, 2
$\text{Loss} = (v/x_2) + wn_2$

where v and w are value weights reflecting a preference or value trade-off between the costs of control and the loss due to cases of the disease. Let the empirical evidence be the following data set:

data = {$n_1 = 0$ for $x_1 = 0$ ppm-years and $N_1 = 200$ workers;
$n_1 = 10$ for $x_1 = 100$ ppm-years and $N_1 = 1000$ workers}

Then the expected loss associated with decision x_2, given the evidence contained in this data set (i.e., after conditioning on the observations) is

$$E(\text{Loss} \mid \text{data}) = (v/x_2) + wE(n_2 \mid \text{data}) = (v/x_2) + wN_2E(P_2 \mid \text{data}),$$

where

$$E(P_2 \mid \text{data}) = 1 - E[\exp(-bqx_2) \mid \text{data}]$$

Choosing x_2 to minimize the expected loss is the same as choosing it to minimize the simpler expression

$$(1/x_2) + kE(P_2 \mid \text{data})$$

where $k > 0$ is defined as $k = (wN_2/v)$. All relevant value trade-offs in this example can be reduced to selection of a value for k.

The term $E(P_2 \mid \text{data})$ may be quantified via Bayes' rule if a prior joint pdf for the reduced parameter *bq* is specified. (Note that it is impossible to estimate *b* and *q* separately from the data. Therefore, assigning each a pdf to each independently and then sampling their values would introduce spurious uncertainty into the model. Using the reduced parameter *bq* avoids this difficulty.) A perhaps more natural, but logically equivalent, alternative is to define a prior pdf for P_1 and update it by conditioning on the data. The posterior pdf for P_1 thus obtained induces a unique corresponding posterior pdf for *bq*, since the two quantities are related by the transformation

$$bq = -\ln(1 - P_1)/x_1.$$

For example, suppose that the prior pdf for P_1 is the maximum-entropy prior, i.e., the uniform density between 0 and 1. Then, after observing the data, the posterior pdf for P_1 has a beta-distribution with parameters $(n_1 + 1)$ and $(N_1 - n_1 + 1)$. Its mean value is:

$$E(P_1 \mid \text{data}) = (n_1 + 1) / (N_1 + 2) = 11/1002$$

(DeGroot, 1970, p. 160). Note that the observation of zero responses when dose = 0 has no effect on inferences in this model, since it is implied by the node formulas. If the correct node formula for determining P were uncertain, however, then the zero-response observation could help to select among alternative formulas. The influence diagram framework is useful for addressing such *model selection* issues as well as data uncertainty issues.

The posterior pdf for P_1 implies a posterior pdf for P_2, for any choice of the decision variable x_2. Namely, the formula

$$P_2 = 1 - \exp(-bqx_2) = 1 - \exp[\ln(1 - P_1)(x_2/x_1)],$$

transforms the beta pdf for P_1 into a pdf for P_2 by a change of variables.

To minimize the scaled expected loss function

$(1/x_2) + kE(P_2 \mid \text{data})$

by appropriate choice of x_2, it is necessary to be able to compute the expected value of P_2 for each choice of x_2. This can be done conveniently using influence diagram tools such as the *Analytica*™ software, originally developed at Carnegie-Melon University and now sold by Lumina Decision Systems, Inc. of Palo Alto. Using this software, and assuming that k = 1000, the example can quickly be solved for the optimal value of x_2. The Bayesian action is found to be x_2 = 3 ppm-years, with expected loss 0.67. Making the standard too relaxed or too stringent by a factor of 10 (i.e., setting x_2 = 30 ppm-years or x_2 = 0.3 ppm-years, respectively) increases the expected loss 5-fold in each case, to about 3.36, and shifts the cumulative loss distribution function rightward.

Discussion

The Bayesian decision-analytic approach requires a prior pdf for uncertain quantities as an input. Its defenders have developed several responses to the criticism that the need to specify a prior pdf injects undesirable subjectivity into the analysis. One defense is that use of subjective prior information is not necessarily undesirable. A second, pragmatic defense is that, when there are adequate data, the posterior pdf is not sensitive to the specific prior selected. Finally, there have been numerous attempts to cope with "vague priors", e.g., by using sets of prior pdfs, by relaxing the requirement that prior beliefs be represented by additive probability measures (as in super-additive Dempster-Shafer belief functions), by selecting maximum-entropy priors as defaults, and so forth (Dubois *et al.*, 1992). Powerful "coherence" arguments support the use of conventional Bayesian prior pdfs – for example, on the grounds that beliefs can be updated in a dynamically consistent fashion only if they are capable of being represented by a Bayesian prior pdf (Epstein and Breton, 1993). In any case, use of Propositions 3 and 4 does not require resolution of the debate over selection and use of priors. If a joint pdf for model inputs is available *in principle*, whether or not it is specified subjectively, the entropy and expected-loss inequalities will hold.

5. CONCLUSIONS

This paper has the following two main practical implications for risk analysts building and using complex risk models:

1. *Always use model reduction* to eliminate spurious uncertainties, prior to model parameter estimation and quantitative uncertainty analysis. For example, Monte-Carlo analyses that randomly sample values for all the original input quantities in a model, instead of only for reduced input quantities, may produce unnecessarily large uncertainties in the outputs.

2. *Always include model complexities that allow additional relevant (although possibly very noisy or imperfect) observations to be incorporated* into the estimation of input-output relations, including dose-response relations. Doing so reduces or leaves unchanged both the uncertainty in model predictions and the expected loss from decisions based on the model.

Concern that adding details and complexity to risk models will reduce their value by impairing their ability to obtain accurate risk estimates may sometimes be justified in a traditional multivariate regression setting, based on sum of squared error loss functions. But this concern is not justified in the Bayesian risk estimation framework we have discussed. Provided that model reduction is used to eliminate "spurious" complexities, which we have defined as those not grounded in observations, greater model complexity enables more certain risk estimates and better decisions.

REFERENCES

Bentler, P.M., and J.A. Stein, 1992. "Structural equation models in medical research." *Stat Methods Med Res*;**1**(2):159-181

Bhaskar, R., and A. Nigam, "Qualitative physics using dimensional analysis," *Artificial Intelligence*, 45, 73-112, 1990.

Buckley, J.J., "Entropy principles in decision-making under risk," *Risk Analysis*, **5**, 303-313, 1985.

Chatfield, C., *Statistics for Technology, 3rd Ed.* Chapman and Hall, New York, 1983.

Cover, T.M., and J.A. Thomas, *Elements of Information Theory*. Wiley, New York, 1991.

DeGroot, M.H., 1970. *Optimal Statistical Decisions.* McGraw-Hill.

DeGroot, M.H., *Probability and Statistics.* Addison-Wesley, Reading, MA, 1975.

Doucet, P., and P.B. Sloep, *Mathematical Modeling in the Life Sciences.* Ellis-Horwood, New York, 1992.

Dubois, D., M.P. Wellman, B. D'Ambrosio, P. Smets (eds), *Uncertainty in Artificial Intelligence: Proceedings of the Eighth Conference.* Morgan Kaufmann, San Mateo, CA, 1992.

Epstein, L.G., and M. Le Breton, "Dynamically consistent beliefs must be Bayesian", *Journal of Economic Theory*, **61**, 11-22, 1993.

Galles, D., and J. Pearl, 1997. "Axioms of causal relevance." *Artificial Intelligence*, **97** (1-2) pp. 9-43

Geoffrion, A.M., 1994. "Structured modeling: Survey and future directions." *ORSA CSTS Newsletter*, **15**:1. An update of this paper and references to other work are at http://personal.anderson.ucla.edu/~art.geoffrion/home/docs/csts/index.htm

Goodman, J., 1987. "A comment on the maximum entropy principle," *Risk Analysis*, **7**, 2, 269-272.

Iwasaki, Y., and H. A. Simon, 1994. "Causality and model abstraction," *Artificial Intelligence* **67**(1), 143-194.

Jensen, F.V., 1996. *An Introduction to Bayesian Networks.* Springer-Verlag, New York.

Judge, G.C., W.E. Griffiths, R.C. Hill, H. Lutkepohl, and T-C Lee, 1985. *The Theory and Practice of Econometrics: Second Edition.* Wiley, New York.

Laffont, J.J., 1990. *The Economics of Uncertainty and Information.* M.I.T. Press.

Lauritzen, S.L., and D.J. Spiegelhalter, "Local computations with probabilities on graphical structures and their application to expert systems," *Journal of the Royal Statistical Society, Series B,* **50**, 2, 157-224, 1988.

Lin, C.C., and L.A. Segel, *Mathematics Applied to Deterministic Problems in the Natural Sciences.* Society for Industrial and Applied Mathematics, Philadelphia, PA, 1988.

Pearl, J., *Probabilistic Inference in Intelligent Systems: Networks of Plausible Inference.* Morgan Kaufmann, San Mateo, CA, 1988.

Quiggin, J., 1993. *Generalized Expected Utility Theory: The Rank-Dependent Model.* Kluwer Academic Publishers, Boston.

Schachter, R., "An ordered examination of influence diagrams," *Networks,* **20**, 535-563, 1990.

Shao, 1999. *Mathematical Statistics.* Springer-Verlag, New York.

Shenoy, P.P., "Valuation-based systems for Bayesian decision analysis," *Operations Research,* **40**, 3, 463-483, 1992.