17 Health Risk Analysis for Risk Management Decision-Making

Louis Anthony (Tony) Cox, Jr.

ABSTRACT

Health risk assessment offers a framework for applying scientific knowledge and data to improve "rational" (consequence-driven) risk management decision-making when the consequences of alternative decisions are uncertain. It does so by clarifying both: (a) The *probable consequences* of alternative decisions (usually represented by conditional probabilities of different consequences occurring, given specified current information and probabilistic risk models); and (b) How *current uncertainties about probable consequences might change* as more information is gathered. This chapter summarizes methods, principles, and high-level procedures for using scientific data (e.g., biological and epidemiological knowledge) to assess and compare the probable human health consequences of different exposures to hazards (i.e., sources of risk); to predict likely changes in exposures and risks caused by alternative risk management interventions; and to evaluate and choose among interventions based on their probable health consequences. The usual goal of these methods is to identify and select actions or interventions that will cause relatively desirable probability distributions of human health consequences in affected populations. We discuss the steps of hazard identification (including causal analysis of data), exposure assessment, causal dose-response modeling, and risk and uncertainty characterization for improving health risk management decision-making.

CONTENTS

Public health risk analysis deals with decisions about which of a set of available risk management *interventions* (usually including the *status quo* or "do-nothing" option) should be implemented. For example, should cell phone use in cars be banned? Under what conditions, if any, should cattle be imported from countries with low levels of diseases such as BSE? Should antibiotics used in human medicine be prohibited from uses in food animals, even if doing so can cause more sick animals (and hence perhaps more sick people), in order to preserve the effectiveness of the antibiotics in treating human patients? To what extent should industrial emissions of specific compounds be restricted?

Health risk analysis provides a set of methods, principles, and high-level procedures for using scientific data (e.g., biological and epidemiological knowledge) to assess and compare the probable human health consequences of different exposures to hazards (i.e., sources of risk); to

assess the likely changes in exposures and risks arising from alternative risk management interventions; and to evaluate and choose among alternative risk management interventions based on their probable health consequences. The goal is usually to identify and select actions or interventions that will cause relatively desirable (e.g., stochastically undominated) probability distributions of human health consequences in the affected population. Health risk analysis is often divided into the overlapping stages of *risk assessment risk management,* and *risk communication,* organized as an iterative process. Table 1 summarizes several traditionally defined steps in this process.

Hazard identification deals with how to establish cause-and-effect relations from data. *Exposure assessment* quantifies the changes in exposures caused by alternative interventions, while dose-response modeling (or exposure-response modeling) quantifies the causal relation between changes in exposures and probable resulting changes in adverse consequences. Finally, *risk characterization* integrates the preceding components to predict the probable changes in health that will be caused by a risk management action that changes exposures.

Health risk assessment uses available facts, data, and models to estimate the health risks to individuals, to an entire population, and to selected subpopulations (e.g., infants, the elderly, immunocompromised patients, and so forth) caused by hazardous exposures and by the decisions and activities that create them. Health risks of sporadic illnesses due to exposures to chemicals, radiation, bacteria, or other hazards are measured quantitatively by the changes in the *frequencies and severities* of adverse health effects caused by the exposures.

Quantitative Definition of Health Risk

For sporadic illnesses (as opposed to epidemics), individual and population health risks can be defined as follows:

- The *individual risk* of sporadic illnesses (or accidents, injuries, or other adverse outcomes) caused by an exposure can be represented by the *frequency and severity of additional adverse health effects per capita-year caused by that exposure*. It can often be tabulated or plotted as the expected number of cases per capita-year in each severity category—e.g., mild, moderate, severe, or fatal, as defined in Buzby, *et al.* (1996) based on illness-days and mortality. To avoid having to carefully define, describe, and compare the severities of different illnesses, one can simply use days of illness per year for each category of illness (e.g., mild, moderate, or severe) to summarize morbidity impacts, perhaps broken down by different age groups or other population sub-groups.

Table 17.1. *Traditional Steps in Health Risk Analysis*

Step	**Purpose and Description**	**Relevant information and techniques**
Hazard identification	Identify potential sources of harm or loss. These sources are called *hazards.* Hazard identification identifies possible adverse health effects of activities or exposures and possible causes of observed adverse effects.	• Human data: Epidemiology, clinical and public health statistics; surveillance data • Animal tests and bioassays • *In vitro* tests • Structure-activity patterns, molecular modeling, pattern recognition and statistical classification techniques
Exposure assessment	Quantify the number of people receiving various levels or intensities of exposure to a hazard over time. Relevant exposure metrics may depend on dose-response relations.	• Environmental fate and transport models, possibly summed over multiple media (paths) and sources. • Studies of human activity patterns. • Biological monitoring of exposed individuals and receptors.
Quantitative exposure-response and dose-response modeling	Quantify the magnitude of risk created by exposure of a target to a hazard. Characterize the probable frequency and severity of adverse health outcomes or losses caused by exposure to the hazard.	A quantitative risk assessment (QRA) runs multiple exposure scenarios through *dose-response models* to predict likely health impacts. Statistical, simulation, or biomathematical models of biological processes are used to quantify dose-response relations.

Risk characterization and uncertainty analysis	Combine estimated probabilities and severities of health harm (adverse consequences), together with indications of uncertainty or confidence, to create an overall summary and presentation of risk.	Monte Carlo simulation calculates risks by sampling many scenarios. Risk profiles, probability distributions, and trade-off and sensitivity analyses display risk, uncertainty, and variability.
Risk communication	Deals with how to present risk information to stakeholders. Considers how different types of recipients perceive risks and internalize/act on messages about them, in deciding what messages to send via what media.	Psychological theories and models and behavioral/experimental findings on risk perception and effective risk communication
Risk management decision-making	Decide what actions to take to control risks and hazards – i.e., accept, ban, abate, monitor, further research, reduce, transfer, share, mitigate, or compensate.	Risk-cost-benefit analysis, formal decision analysis for groups and individuals, risk quantification and comparison

Alternatively, and often more conveniently, the loss due to increased mortality and morbidity can be expressed in terms of quality-adjusted life-years (QALYs), which can serve as a single summary measure of severity if the required preference-independence conditions justifying QALYs are accepted (Hazen, 2003; Miyamoto, 1999). Individual risk is then given by the joint probability distribution of the number of cases per capita per year and the associated severities (i.e., QALYs lost per case).

- *Population risks* are described by the sum (or, in more detail, by the frequency distribution) of individual risks over all person-years in the population. They can be expressed as *numbers of additional adverse health effects per year* (of each type or severity category) occurring in the population. Population risks can also be further characterized by identifying subpopulations with especially high individual risks.

Technical Note: Use of Expected Values. Use of the expected number of events per year to quantify risk is justified for sporadic illnesses that occur independently, or with only weak statistical dependence, in large populations, when the Poisson approximation (Janson, 1994) or the compound Poisson approximation (Barbour and Mansson, 2000) holds. The expected number of cases per year then determines the full probability distribution of the number of illnesses per year, to a close approximation (made precise in the above references). Moreover, the Poisson probability distribution is stochastically increasing in its mean; thus, larger numbers of expected cases correspond to less preferred distributions for *all* decision-makers who prefer fewer cases per year to more. The formulae *Individual risk = expected number of additional illnesses per year × expected QALYs lost per illness* and *Population risk = sum of individual risks* are useful for sporadic illnesses, although they must be generalized for other types of risks, e.g., to allow for risk aversion (Cox, 2001).

The main goals of risk assessment are to produce information to improve risk management decisions by *identifying and quantifying valid cause-effect relations between alternative risk management decisions and their probable total human health consequences*, and by identifying decisions that make preferred outcomes more likely. Health risk assessments typically use explicit – and, if possible, validated – analytic models (e.g., statistical, biomathematical, or simulation models) of causal relations between actions and their probable health effects. In general, quantitative risk assessment applies specialized models and methods to quantify likely exposures and the frequencies and severities of their resulting health consequences.

Example: Statistical and Causal Risk Relations May Have Opposite Signs

As illustrated by the following (perhaps counter-intuitive) example, *there is no necessary relation between statistical exposure-risk associations and the change in risk that would be caused by changing exposure.* As a simple counterexample, consider a hypothetical population in which 100% of men and 0% of women are exposed (i.e., Exposure = 1 for men, Exposure = 0 for women); and in which Risk = 0 for women, Risk = 100% for unexposed men, and Risk = 10% for exposed men. In this example, exposure reduces risk, but the statistical association between them is positive. The *statistical* relation between exposure and risk in this population is:

Risk = 0.1 × Exposure,

That is, when Exposure = 1, Risk = 10% (for exposed men), and when Exposure = 0, Risk = 0 (for unexposed women.) Yet, the *causal* effect of reducing Exposure is to *increase* risk in the population, by shifting men from the lower-risk exposed group to the higher-risk unexposed group. The causal relation between Exposure and Risk in this population is thus:

Risk = 1 – 0.9 × Exposure for men; Risk = 0 for women.

In general, fitting a simple reduced-form statistical model to data does *not* allow one to correctly predict the effects of changing the independent variables on resulting changes in the dependent variable (Shipley, 2000; Freedman, 2004). [This example is motivated by empirical relations found in a real data set collected by CDC (Friedman et al., 2000) for the foodborne bacterial pathogen Campylobacter. Men *do* appear to have greater susceptibility to campylobacteriosis than women; they *do* appear to have greater exposure to risk factors such as eating undercooked meat in restaurants and swimming in untreated water; and exposure to chicken (e.g., buying and handling raw chicken, preparing and eating chicken at home, etc.) *does* appear to reduce risk of campylobacteriosis, for both sexes. The above counterexample exaggerates these empirical

patterns to extremes to provide a simple illustration of the disconnect between statistical and causal relations.]

A Bayesian Network Framework for Health Risk Assessment

To support effective risk management decisions, human health risk assessments must characterize known or suspected potential causal relations between risk management actions (including the *status quo* or "do-nothing" option), on the one hand, and probable resulting human health consequences on the other. Actions typically affect exposures to sources of risk (i.e., hazards), while consequences typically include changes in the frequency or severity of resulting illnesses or deaths in affected populations. *Hazard identification* identifies causal relations (possibly including causal paths) leading from risk management actions to their human health consequences. Hazard identification often precedes any plan to develop a risk management strategy, as effective risk management is often impossible if causal relations are not understood.

Figure 17.1 outlines a causal graph (Shipley, 2000; Greenland and Brumback, 2002) for assessing risks to humans from changes in exposures to hazards. In this template, risk management *actions* can change *exposures* of individuals to potentially harmful agents (the hazards). Changes in exposures, in turn, change expected *illness rates* and hence adverse *health consequences* (e.g., illness-days or early deaths per capita-year) in susceptible members of the exposed population. If desired, different human health consequences can be aggregated into a single summary measure, such as quality-adjusted life-years (QALYs), if the required preference conditions hold (Hazen, 2003), but this is optional. The effects of such changes on the number of QALYs lost per year in the population can be mediated by individual behaviors or attributes (e.g., immune status, age, gender, diet, behaviors, and other covariates that affect susceptibility

to infections). These covariates may also influence each other (indicated by the brackets [] around them in Figure 17.1. For example, an AIDS patient may have food consumption and preparation behaviors and medical treatments that differ from those of a non-AIDS patient. Risk assessment helps to identify risk management options (acts) that decrease adverse health consequences, taking into account the distribution of covariates in the population.

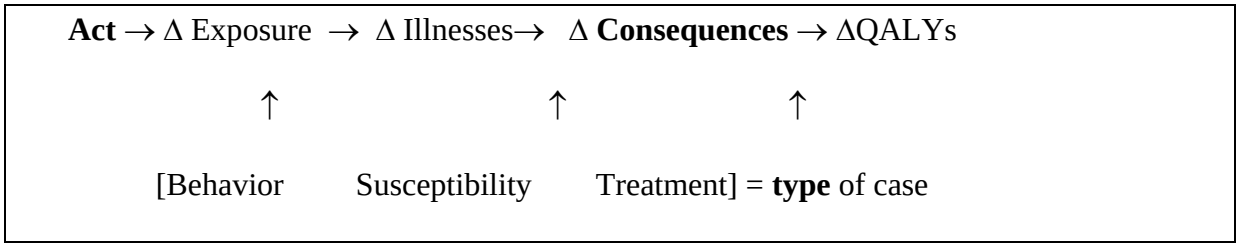

Figure 17.1 A causal graph for health risk analysis

Technical note: Influence diagram interpretation. Figure 17.1 can be interpreted as a Bayesian belief network or causal graph model (Greenland and Brumback, 2002; Chang and Tian, 2002). In this framework, each variable to which arrows point is interpreted as a random variable with a conditional probability distribution that is completely determined by the values of the variables that point into it. Because this diagram has a decision node ("act") and a value node ("ΔQALYs"), it is an example of an influence diagram (Owens *et al.*, 1997). Important details are represented only implicitly, by conditional probability distributions. Algorithms to identify possible causal graph structures from data (and hence to test whether hypothesized causal theories are consistent with data) have been developed (e.g., Tsamardinos *et al.*, 2003), but are not yet routinely applied in risk assessment. Such causal graph models are useful because effective algorithms exist to (a) Quantify the conditional probability distributions of any subset of their variables, given observed values of the rest; and (b) Solve for acts that give maximum expected utility (once a utility function has been defined for outcomes such as ΔQALYs) (see Crowley, 2004, http://www.cs.ubc.ca/~crowley/academia/papers/aiproj.pdf).

Each choice of a risk management **act** in Figure 17.1 generates a corresponding random number of incremental illness cases ("responses") caused or prevented each year in each severity class of consequences (e.g., mild, moderate, severe, fatal) in the population (and in each subpopulation, if there are several). The expected health consequences of this change can be calculated from the following three submodels, which are common to most risk assessments:

- An *exposure model* (the "act → Δexposure" link in Figure 1) that quantifies the amounts of exposure received per unit time by exposed individuals.
- A *dose-response* or *exposure-response* model (the "Δexposure → Δillnesses" link in Figure 1) that quantifies the probability of illness, or the expected incremental number of cases at each given severity level, per unit of exposure. In general, this relation may depend on the individual's "type (i.e., on the combination of covariate values that influence risk for that individual), as well as on the dose (units of exposure) received.
- A *health consequence model* (the "Δillnesses → Δconsequence" link in Figure 1) quantifying the conditional probabilities of different health outcomes (e.g., survival *vs.* fatality, or number of QALYs lost) from each case. These outcome probabilities may depend on factors such as physician prescription behavior or hospital infection-control standards.

These three submodels determine the expected illnesses and QALYs lost per year in each severity class for each act. Multiple exposure pathways and at-risk populations (perhaps including groups receiving different medical treatments) can be included to quantify the *total* human health impact of different acts. Summing health impacts over all distinct combinations of hazards, exposure routes, and target populations (each corresponding to an instance of Figure 1) gives the total probable change in human health consequences for the act.

Technical Note: Monte Carlo simulation. If there are too many combinations of hazards, exposure routes, and target populations for explicit summation over all of them to be practical, then Monte Carlo simulation can be used to obtain accurate numerical approximations of the average risk (and the distribution of health effects). For example, suppose that risk is given by $f(x_1, x_2 \ldots x_n) = f(\mathbf{x})$, and that one can sample from the joint probabilitity density function (PDF) of the x_i, $Pr(x_1, x_2 \ldots x_n) = Pr(x_1)\, Pr(x_2 \mid x_1) \ldots Pr(x_n \mid x_1, \ldots, x_{n-1})$. Then Markov Chain Monte Carlo (MCMC) simulation techniques such as Gibbs sampling (Andrieu *et al.*, 2003; Lange, 2003 can be used to generate random samples from the joint PDF of **x**. Taking a simple arithmetic average of the values of f(**x**) obtained for a sufficiently large random sample of **x** values will give an accurate estimate of the true average risk $E_{Pr(x)}[f(\mathbf{x})]$ implied by f(**x**) and Pr(**x**). Commercial risk analysis software tools such as Analytica™, @RISK™, and Crystal Ball™ include Monte Carlo simulation routines that can generate estimated means, confidence bands, and entire estimated probability distributions for f(**x**). Vose (1998) provides a basic introduction to Monte Carlo simulation in spreadsheet models for microbial risk assessment and Cassin *et al.* (1998) discusses how to use Monte Carlo simulation for tasks such as priority-setting and risk management.

The conceptual framework in Figure 1 can be implemented with greater or lesser degrees of sophistication. Perhaps the simplest approach is to generate point estimates for each risk management act and exposure pathway for each of the following:

- *Exposure factor* = units of exposure received per capita per year;
- *Dose-response* factor = expected cases of illness per unit of exposure;

- *Health consequence factor* = expected QALYs lost (or illness-days created, etc.) per case of illness. (Alternatively, a vector of expected numbers of different health outcomes can be estimated; e.g., mild, moderate, severe, and fatal outcomes per case.)

In this approach, each sub-model (corresponding to a horizontal arrow in Figure 1) is represented by a single number. One can then multiply these numbers together, and multiply by the number of people affected, for each causal path and each risk management action. (Causal paths may include not only different exposure paths, but may encompass all three links.) Summing the results over all causal paths provides an estimate of the total human health impact per year for each action. A more refined calculation can be made by considering how these factors might change over time, and then summing over time periods (perhaps with discounting).

At the other end of the spectrum, Figure 1 can be applied to risk estimation using conditional probability algorithms developed for Bayesian networks and causal graphs (Chang and Tian, 2002). In this case, *hazard identification* can be thought of as identifying instances of Figure 17.1 that are consistent with available data. Statistical methods are available to test whether specified causal graph models are indeed consistent with data (Greenland and Brumback, 2002; Shipley, 2000), and practical algorithms have been developed to identify potential causal graph models from multivariate data (Aliferis *et al.*, 2003; Tsamardinos *et al.*, 2003). The remaining steps in the risk assessment process can then be interpreted as quantifying and applying the resulting Bayesian network. In this framework, the simple approach of multiplying exposure, dose-response, and consequence factors generalizes to allowing arbitrary probability distributions for inputs and conditional probability relations to be combined via Monte Carlo simulation (Andrieu *et al.*, 2003) to derive the joint probability distributions of the outputs.

Bayesian network methods – combined with objective statistical tests for potential causality, such as conditional independence tests (Shipley, 2000; Greenland and Brumback, 2002) – appear promising for providing more effective, data-driven risk assessments, while also allowing for the use of expert judgment when necessary. See for example Parsons *et al.*, 2005 for preliminary work on Bayesian networks and related methods for risk assessment.

Hazard Identification

Risk assessment begins with *hazard identification* the process of specifying the scope of the assessment and summarizing the available empirical evidence that exposure to a specific "hazard" causes specified adverse health effects in exposed individuals or populations. Thus, hazard identification can serve to:

1. *Rapidly screen potential hazards* by identifying whether available data support the hypothesis that the hazard might cause specific adverse health effects (possibly using formal statistical methods of causal analysis; e.g., Shipley, 2000.
2. *Identify causal relations between specific hazards and specific adverse human health effects.*
3. *Identify risk factors, behaviors, and exposure conditions that increase risks to specific exposed populations* (e.g., the old, the young, the immuno-compromised, etc.)
4. *Summarize empirical evidence both for and against the hypothesis that exposures to specific hazards cause specific adverse human health effects* (Patton, 1993.)

In reality, of course, *joint causation* is common; i.e., observed adverse consequences are often due to a combination of a hazardous agent, activities resulting in exposures to that agent, failure to undertake protective actions, and possibly other confounding factors, such as decreased immunity in a sub-population. In general, *any* event or condition that hastens the occurrence of an adverse effect or increases its likelihood can be viewed as a contributing "cause" of the effect.

For more on the philosophical definition and ambiguities of "causation", see Williamson (2005). Thus, "the cause" of an adverse health effect is often not uniquely defined. Nonetheless, for purposes of risk management, it often suffices to predict the effects of alternative risk management interventions on the rates of adverse events of different severities. Hazard identification helps to identify such interventions.

Table 17.2 below outlines steps for forming and testing causal hypotheses about exposure-response relations using epidemiological data. As more of these steps are completed, the empirical support increases for a causal relation between exposure and risk. Most statistical methods in epidemiological risk analysis focus on steps 1-3; i.e., identifying non-random associations, and then eliminating potential biases and confounders as likely explanations. These steps can often be carried out using observational data, even without experimental controls, by using the *refutationist approach* (Maclure, 1990, 1991). This systematically enumerates possible competing explanations for the observed data, and then eliminates each of those potential non-causal explanations (if possible) using statistical tests on the available data.

Table 17.2. *Steps to Establish a Causal Exposure-Risk Relation*

1. *Identify a statistically significant exposure-response association*; e.g., using case-control, prospective cohort, or other cross-sectional or longitudinal epidemiological data.
2. *Eliminate confounding* as a possible explanation of the association, by accounting for factors such as lifestyle, age, or exposure to other hazards, e.g., using conditional independence tests (Grimes and Shulz, 2002; Feldman, 1998; Greenland and Morgenstern, 2001).
3. *Eliminate biases in sampling, information collection, and modeling choices* as possible explanations for the association (Choi and Noseworthy, 1992; Deeks *et al.*, 2003).
4. *Test and confirm hypothesized causal and conditional independence relations*, for example, by showing that the response is *not* conditionally independent of the hypothesized exposure that causes it, given other variables (Shipley, 2000; Friedman, 1996; Frey *et al.*, 2003).

5. *Confirm efficacy of interventions*, e.g., by experimental manipulations and/or intervention and change point analyses of time-series data (e.g.., Swanson *et al.*, 2001; Green, 1995).
6. *Identify and elucidate causal mechanism(s)*, identified from experimental data and/or from generally accepted principles.

Many epidemiologists have recognized that, to draw valid causal inferences, it is necessary to refute competing (non-causal) hypothesized explanations for observed exposure-response associations (Maclure, 1990, 1991). Table 3 summarizes common competing explanations (mainly, confounding and/or sampling, information, or modeling biases), and some suggested statistical methods to refute them (Cox, 2001, Chapter 3).

Table 17.3. *Potential Non-Causal Explanations for Exposure-Response Associations*

Potential Non-Causal Explanations	**Methods to Refute Potential Explanations**
Modeling Biases	
Variable selection bias (includes selection of covariates in model)	Bootstrap, Bayesian model averaging (BMA), and cross-validation for variable selection (Wang et al., 2004).
Omitted explanatory variables (including omitted confounders)	Include potential confounders in an explicit causal graph model; test for unobserved latent variables
Variable coding bias (coding may affect apparent risk)	Don't discretize continuous variables. Use automated variable-coding methods (e.g., classification trees).
Aggregation bias/Simpson's paradox	Test hypothesized relations at multiple levels of aggregation, down to individual-level data.
Multiple testing/comparisons bias	Adjust p-values (Romano and Wolf, 2005).
Choice of exposure and dose metrics; choice of response effect definitions	Use multiple exposure indicators (e.g., concentration and time). (Don't combine.) Use survival functions and/or transition rates among observed health states.
Model form selection bias; uncertainty about correct model	Use flexible non-parametric models (e.g., smoothers, wavelets) and BMA for multiple models. Report model diagnostics and sensitivities of results to model forms (Greenland, 1989).
Missing data	Use data augmentation, expectation maximization (EM) algorithm, MCMC algorithms (Schafer, 1997).
Measurement and misclassification errors in explanatory variables	Use Bayesian measurement error models, data augmentation, missing-data techniques (Schafer, 1997; Ibrahim *et al.*, 2005).
Unmodeled heterogeneity in individual response parameters	Latent variable and finite mixture distribution models, frailty models of inter-individual variability.

Biases in interpreting and reporting results	Report results (e.g., posterior PDFs) *conditioned* on data, models, and statistical methods. Show sensitivities.
Sample Selection Biases	
Sample selection (sample does not represent population)	Randomly sample *all* cohort members if possible.
Data set selection bias (i.e., selection of studies may affect results)	Meta-analysis of sensitivity of conclusions to studies. Use causal graph models to integrate diverse data sets.
Health status confounding, hospital admission/referral bias	If possible, use prospective cohort design and population-based cases and controls (Choi and Noseworthy, 1992).
Selective attrition/survival (e.g., if exposure affects attrition rates) Differential follow-up loss	Use a well-specified cohort. "Include non-surviving subjects in the study through proxy interviews" (Choi and Noseworthy, 1992). Compare counter-factual survival curves.
Detection/surveillance bias	Match cases to controls (or exposed to unexposed subjects) based on cause of admission.
Membership bias (e.g., lifestyle bias, socioeconomic history)	• In cohort studies, use multiple comparison cohorts. • Hard to control in case-control studies.
Self-selection bias; Response/volunteer bias	Achieve response rate of at least 80% by repeated efforts. Compare respondents with sample of non-respondents.
Information Collection Biases	
Intra-interviewer bias	Blind interviewers to study hypotheses, subject classifications.
Inter-interviewer bias	Use same interviewer for study and comparison groups.
Questionnaire bias	Mask study goals with dummy questions; avoid leading questions/response options.
Diagnostic suspicion bias Exposure suspicion bias	Hard to prevent in case-control studies. In cohort studies, make diagnosis and exposure assessments blind to each other.

As stated by Savitz *et al.* (1990), "Biases that challenge a causal interpretation can always be hypothesized… It is essential to go beyond enumerating scenarios of bias by clearly distinguishing the improbable from the probable and the important from the unimportant." Fortunately, well-developed statistical methods and algorithms are now available to: (a) identify significant statistical associations from data showing spatial and temporal associations between exposures and health effects (e.g., Mather *et al.*, 2004); and (b) screen them for potential causality based on the above criteria.

Technical Note: Statistical tests for assessing potential causality. Over the past forty years, intuitive criteria for causality used in in epidemiology (such as the Bradford Hill criteria,

emphasizing strength, consistency, biological gradient, coherence, etc. of an association) have been made more rigorous, general, and quantitative by advances in applied decision sciences. For example, an approach based on information theory proposes that, roughly speaking, a data set provides evidence that exposure variable X is a *potential cause* of response variable Y if and only if X is: (a) INFORMATIVE about Y, i.e., the mutual information between X and Y, denoted by I(X; Y) and measured in bits (Cheng *et al.*, 2001), is positive in the data set. (This allows for non-linear and even non-monotonic relations.) (b) UNCONFOUNDED: X provides information about Y that cannot be removed by conditioning on other variables, i.e., I(X; Y | Z) > 0 for all subsets of variables Z disjoint from X and Y. (c) PREDICTIVE: Past values of X are informative about future values of Y, even after conditioning on past values of Y. (This generalizes the concept of Granger causality for time series, e.g., Guatama and Van Hulle, 2003.) (d) CAUSALLY ORDERED: Y is conditionally independent of the parents of X, given X, i.e., I(P ; Y | X) = 0, for any parent or ancestor P of X.. These principles yield practical algorithms (e.g., BayesiaLab™, Tsamardinos *et al.*, 2003) for detecting potential causation in cohort, case-control, and time series data sets, even if the functional relations involved are nonmonotonic. (Causation may be present even if these conditions are not satisfied, but then the data do not provide evidence of it.) Formal tests for statistically significant associations between the timing of one event (e.g., introduction or cessation of exposures) and subsequent changes in a series of measurements (e.g., human illness rates in a surveillance program) can be based on *intervention analysis* and *change point analyses* (Green, 1995) for time series. These methods for testing for potential causality are entering common biostatistical and risk analysis practice only slowly, but appear to be very promising (Shipley, 2000).

Exposure Assessment

For environmental risk assessment, US EPA experts have stated that "Questions raised in the exposure analysis concern the likely sources of the pollutant… its concentration at the source, its pathways (air, water, food) from the source to target populations, and actual levels impacting target organisms" (Patton, 1993). Similarly, for microbial hazards, the US FDA has defined exposure assessment as "A component of a risk assessment that characterizes the source and magnitude of human exposure to the pathogen". The magnitude of human exposure, also called the dose, is defined as "The amount or number of a pathogen that is ingested or interacts with an organism (host)" (http://www.foodsafety.gov/~dms/lmriskgl.html).

Exposure assessment seeks both to identify exposed subpopulations at risk from exposures to hazards and also to identify conditions leading to high-risk exposures. It describes the extent of exposures (frequency and magnitude of individual exposures in the population in relation to susceptibility and covariates) and uses models to predict how risk management decision options will probably affect them. A successful exposure assessment should describe the frequency distribution of exposures received by members of exposed populations and subpopulations and should show how these distributions change for different risk management decisions. The descriptions should contain enough detail to discriminate among different exposure distributions that would cause significantly different health outcomes. This information is used, together with dose-response information, to characterize risks

The shape of the frequency distribution of exposures relative to the dose-response relation (e.g., how frequent are exposures that are likely to cause illness?) drives quantitative risk. It is common for exposures to be very uncertain, especially if they depend on unmeasured and/or highly variable processes. The exposure assessment influence diagram may then look like this:

Act → exposures → illnesses ← individual covariates

↓

measured exposure surrogates

For example, available data may consist of surrogate measurements (e.g., contaminant levels in exposure pathways) rather than direct measurements at the point of exposure. True exposures then play the role of *latent variables* in causal modeling, i.e., they affect observed outcomes but are not observed themselves. Appropriate statistical techniques for causal diagrams with latent variables (e.g., Shipley, 2000 for linear models; Pearl, 2002 and Hartemink *et al.*, 2001 for more general Bayesian Network models) can be applied to the above diagram with surrogate measurements of exposure for data. Software such as WinBUGS helps to automate the required computations for inference with missing data and unobserved or surrogate variables.

Exposure models describe the transport and distribution of hazardous materials through different media and pathways (e.g., air, foods, drinking water) leading from their source(s) to members of the exposed population. In addition, exposure models may consider the distribution over time of human populations among locations and activities result in exposures. Simulation models of transport and behavioral processes, often developed using discrete-event simulation software, can be used to estimate frequency distributions of population exposures from assumptions about or sub-models of the more detailed micro-processes involved.

Example: Simulation of Exposures to Pathogens in Chicken Meat

The World Health Organization (WHO) has described a process simulation model of human exposures to the foodborne pathogen Salmonella as follows:

The exposure assessment of *Salmonella* in broiler chickens mimics the movement of *Salmonella*-contaminated chickens through the food chain, commencing at the point of completion of the slaughter process. For each iteration of the model, a chicken carcass was randomly allocated an infection status and those carcasses identified as contaminated were randomly assigned a number of *Salmonella* organisms. From this point until consumption, changes in the size of the *Salmonella* population on each contaminated chicken were modeled using equations for growth and death. The growth of *Salmonella* was predicted using random inputs for storage time at retail stores, transport time, storage time in homes, and the temperatures the carcass was exposed to during each of these periods. Death of *Salmonella* during cooking was predicted using random inputs describing the probability that a carcass was not adequately cooked, the proportion of *Salmonella* organisms attached to areas of the carcass that were protected from heat, the temperature of exposure of protected bacteria, and the time for which such exposure occurs. The number of *Salmonella* consumed were then derived using a random input defining the weight of chicken meat consumed, and the numbers of *Salmonella* cells in meat as defined from the various growth and death processes. Finally, in the risk characterization, the probability of illness was derived by combining the number of organisms ingested (from the exposure assessment) with information on the dose-response relationship (hazard characterization)."

(www.who.int/foodsafety/publications/micro/Salmonella/en/).

The results of the Monte Carlo simulation exposure modeling are presented as: (a) an estimated 2% prevalence of contaminated chicken servings; and (b) the following conditional frequency distribution for the dose (CFUs)-per-serving from contaminated servings:

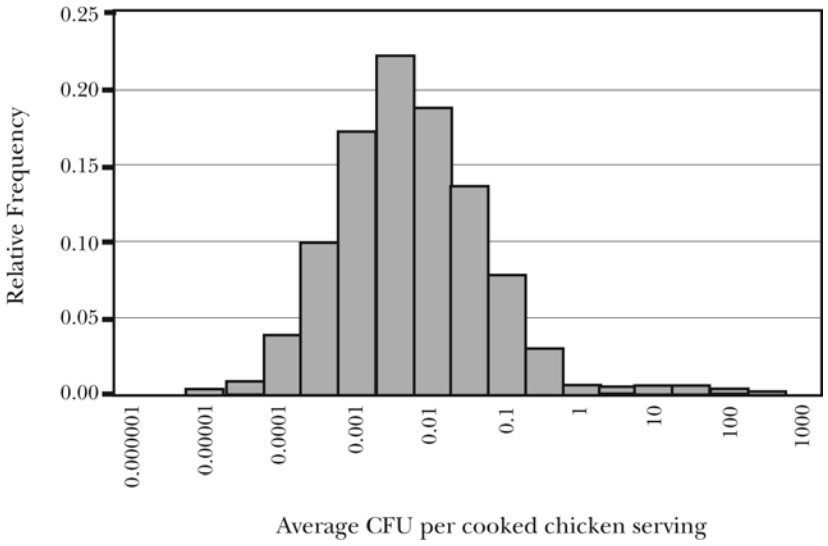

Source: http://www.who.int/foodsafety/publications/micro/*Salmonella*/en/

Figure 17.2. Average CFU per cooked chicken serving

This frequency distribution shows how large an exposure a person is likely to receive from a serving of contaminated, undercooked broiler chicken. This is the main output of the exposure assessment and the main input to the dose-response model for calculating illness risk per serving.

Example: Mixture Distributions and Unknown Exposure-Response Models

Unknown or uncertain exposure-response relations in a population can often be estimated by decomposing the risk as follows:

$\Pr(\text{Illness} \mid \text{exposure} = x) = \Sigma_r \Pr(\text{Illness} \mid \text{exposure} = x \ \& \ \text{response type} = r) \times \Pr(\text{response type} = r)$.

Here, “response type” is an unobserved (latent) variable summarizing all of the missing information needed to predict the probability of illness from a known level of exposure. (For

example, if each individual has an unknown threshold number of bacteria that must be ingested in one meal to cause illness, then r would be that threshold number. If there is a continuum of response "types", the above sum is replaced by an integral.) A useful development in mathematical statistics is the recognition that the uncertain quantities Pr(response type = r) can be interpreted as *statistical coefficients* to be estimated directly from data on the aggregate number of responses observed in populations for different exposure conditions, while the conditional response probabilities that are paired with these coefficients, Pr(illness | exposure = x, type = r) can be estimated simultaneously from the same data (provided that technical identifiability conditions are met. These are automatically satisfied by many families of statistical distributions.) The required statistical methodology is that of *finite mixture distribution models* if the number of types is finite; or continuous mixture models if types are continuous. Well-developed computational Bayesian algorithms can be applied to estimate the number of components in the mixture (i.e., the number of statistically significantly different "types") and the corresponding coefficients and conditional response probabilities (see e.g., Richardson and Green, 1997; Stephens, 2000; Miloslavsky and van der Laan, 2003.) In this construction, the exposure variable x can be any measured quantity that can be paired with corresponding illness rates. All unobserved details are absorbed into the latent "type" variable, r. Missing values and errors in measured values of x can also be handled within the computational Bayesian framework (e.g., using the data augmentation algorithm, Schafer, 1997) to allow the conditional distributions of outputs given observed data to be quantified, even when other data are missing. There is thus great flexibility within simulation approaches to use all available data (via conditioning), but without requiring use of unavailable data.

Dose-Response Modeling

Dose-response models quantify the conditional probability of illness caused by each level of exposure; thus, the term *exposure-response model* is also appropriate. Figure 17.2 shows an example of a dose-response model developed for *Listeria monocytogenes* in ready-to-eat foods. A specific parametric dose-response model was assumed (an exponential model) and fit to epidemiological data for immunocompromised ("High risk") and non-immunocompromised ("Normal") subpopulations. The dark solid curve in Figure 2 is the estimated dose-response model for the "Normal risk" subpopulation. The dashed line above and to the left of it is the dose-response model for the "High risk" subpopulation. The lighter gray curves indicate estimated statistical confidence bands around these best-estimate curves – an upper confidence band for each (corresponding to the upper end of the 95% confidence interval estimated for the parameter of the exponential dose-response model), and a lower 95% confidence band for the right-most (Normal) dose-response model.

As in Figure 17.3, it is often necessary to fit separate dose-response models to "normal" and "susceptible" subpopulations within the general population to account for inter-individual variability in dose-response relations. While more than two gradations of susceptibility can be modeled using finite mixture distributions, distinguishing between only two levels or response "types" in the population, i.e., susceptible and normal, often suffices to explain most of the variability in the data. If different degrees or severities of illness are distinguished, ranging from mild through severe to fatal, then a *health consequence model* describing the conditional probabilities of different levels or severities of health outcomes, given that illness occurs, is needed to augment the conditional probability of illness as a function of exposure. In general,

risk characterization requires describing the *severities* as well as the *frequencies* of adverse health outcomes caused by exposures.

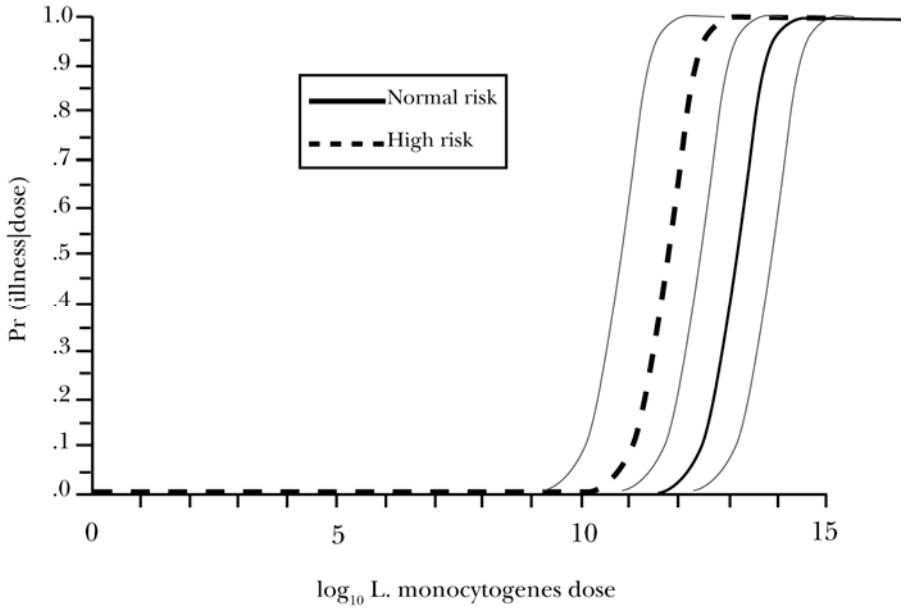

Figure 17.3 Example Dose-Response Function for Listeria monoctogenes

In practice, biologically motivated parametric dose-response models are the most common, and usually the best justified, models in widespread use. They are typically fit to data by a combination of maximumum likelihood estimation (MLE) for point estimates and computationally intensive resampling techniques (e.g., bootstrapping algorithms) for confidence intervals, simultaneous confidence bands around the dose-response curve, and joint confidence regions for model parameters (e.g., Haas *et al.*, 1999, Chapter 7, c.f. p. 293).

Source: FAO/WHO, 2001. http://www.who.int/foodsafety/publications/micro/en/may2001.pdf

Example: Best-Fitting Parametric Models May Not Fit Adequately

Figure 17.4 for *Salmonella* feeding trial data show that even the best-fitting model in a certain class of parametric models (here, the approximate Beta-Poisson dose-response family, widely

used in microbial risk assessment) may not adequately describe the observed data. The parametric family of models is then said to be *misspecified* for the data, i.e., it is not appropriate for describing the empirical relation. In this example, the approximate Beta-Poisson model family is inappropriate for the data because even the best-fitting curve in the family dramatically under-predicts low-dose risks.

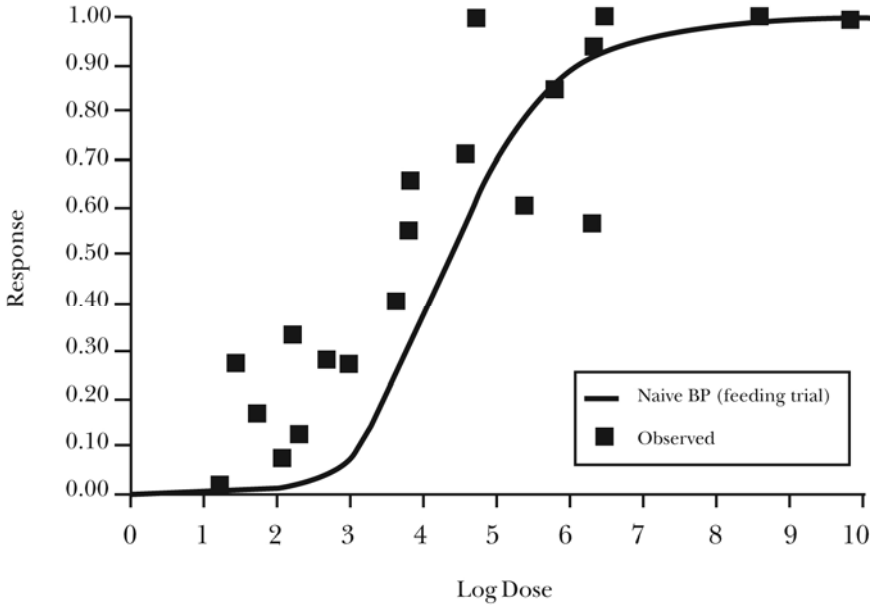

Source: WHO/FAO, 2002. (Naïve BP = approximate Binomial Poisson)

Figure 17.4 The Best-Fitting Beta-Poisson Model Under-Predicts Low-Dose Risks

If the correct dose-response model is unknown and several models all provide adequate fits to the available data, multiple plausible models may be used to carry out the rest of the assessment. In this case, the analysis can be organized and presented as a *model uncertainty decision tree* in which different modeling choices correspond to different branches in the tree. The results of the risk analysis at the end of each branch are contingent on the assumptions and modeling choices that lead to it. Different branches may be weighted by the relative strength of the evidence supporting them (Kang *et al.*, 2000). Bayesian Model Averaging (BMA) provides a

more formal version of this approach (Viallefont *et al.*, 2001; Keiding and Budtz-Jorgensen, 2004). Model uncertainty decision trees can also be used to present and analyze uncertainties due to choices of dose metrics, response definitions, and other modeling decisions, as well as choices of particular dose-response models.

Uncertainty about illness probabilities caused by a given dose is often dominated by uncertainty about the most appropriate dose-response model. A decision tree presentation of alternative modeling choices and the resulting predicted risks – or even a simple plot of different plausible dose-response curves – can express much of the relevant uncertainty with a minimal amount of statistical sophistication. Other important computational methods and algorithms for uncertainty analysis include:

- *Monte Carlo uncertainty analysis* using commercial software products such as Analytica™, @RISK™, Crystal Ball™ (Vose, 2000). For more on uncertainty and sensitivity analysis software, see the descriptions at product web sites.
- *Bayesian uncertainty analysis* for model parameters and predictions (e.g., based on the WINBUGS software for inference with missing data.)
- *Bootstrapping and other resampling techniques* for estimating joint confidence regions for model parameters and predictions.
- *Model cross-validation* techniques for estimating the accuracy and prediction error characteristics of model predictions from performance on multiple subsets of data.

These methods are discussed in general computational statistics texts and, for dose-response modeling, in risk analysis texts such as Haas *et al.*, 1999, Vose, 2000, and Cox, 2001.

Risk And Uncertainty Characterization For Risk Management

Risk characterization is the ultimate output of a risk assessment. It integrates hazard identification, exposure assessment, and dose-response information to determine the probable frequency and severity of adverse health effects in a population caused by exposures to a hazard. Characterizing the change in risk for different risk management interventions helps decision-makers choose among them. Risk characterization also includes characterization of current *uncertainty about risk*. This allows the value of gathering additional information to be assessed as part of risk management deliberation and decision-making, based on the potential value of such information (VoI) to enable risk managers to make choices that are more likely to result in desired consequences (Yokota and Thompson, 2004).

Given the results from:

(a) Exposure assessment (i.e., the conditional probability distribution of exposures, for each act);

(b) Exposure-response/dose-response modeling (i.e., the conditional probability of illness for each exposure pattern); and

(c) Consequence modeling (e.g., the conditional probability distribution of adverse consequences given illness),

the risk characterization step calculates, for each act being assessed, the resulting probability distributions for adverse consequences. (This can be done by literally summing or integrating expressions such as Pr(consequence = c | illness) × Pr(illness | exposure = x) × Pr(exposure = x | Act) over all exposure levels x, to obtain the probability of each consequence, c.)

Example: Risk Characterization Outputs

Figure 17.5 shows one of the risk characterization outputs from a risk assessment of *Listeria monocytogenes* (FAO/WHO, 2001). The solid curve shows the median estimate of the mortalities per year caused among the elderly subpopulation by *L. moncytogenes* in deli meats, for different maximum allowed storage times. The dotted curves represent the 5th and 95th percentiles of the uncertainty distribution, as assessed by Monte Carlo uncertainty analysis.

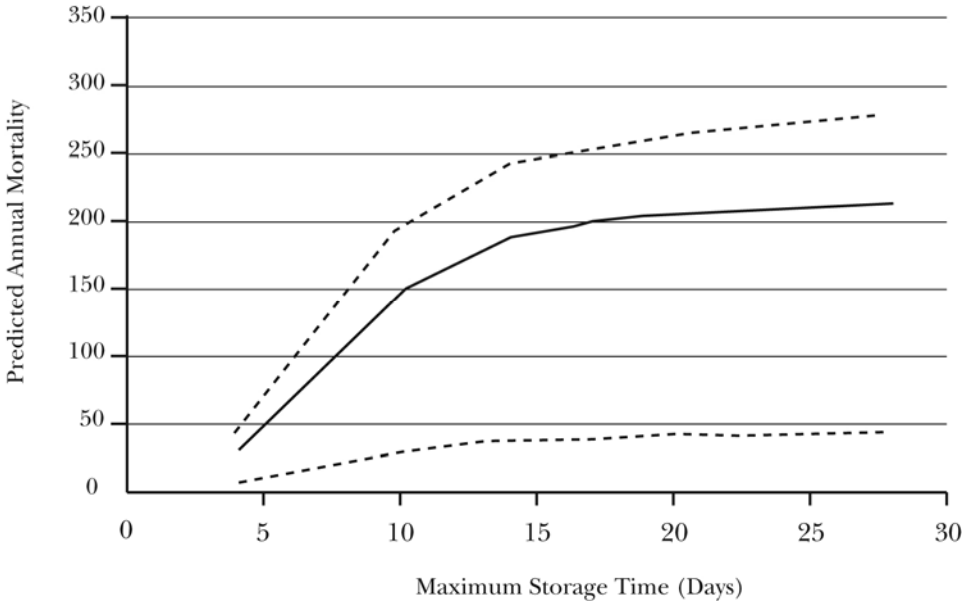

Source: FAO/WHO, 2001, http://www.cfsan.fda.gov/~dms/lmr2-6.html

Figure 17.5 Predicted annual mortality in the elderly subpopulation attributable to deli meats as a function of maximum storage time

This display shows how predicted risks in this subpopulation vary with the effects of different potential interventions that would limit the maximum storage times allowed for deli meats. Similar curves can be shown for the effects of such interventions for other foods or groups of foods (e.g., dairy products, produce, sea food products, etc.) and for other subpopulations and the U.S. population as a whole.

Risk management is often viewed as a decision process that takes scientific information obtained from risk assessment as an input, along with value judgments and policy goals and constraints, and that recommends choices of risk management actions as its output. Alternative risk management approaches include risk acceptance, prevention or avoidance (e.g., by reducing exposures), mitigation of consequences (e.g., by appropriate clinical screening, diagnosis, and prescription procedures), transfer (e.g., health insurance), and compensation.

A successful risk analysis shows the estimated changes in the frequencies and magnitudes of adverse human heath consequences resulting from different risk management decision options. (Of course, if hazard identification and risk management reveal that the risk from the *status quo* is so small that no risk management action is needed, analysis may stop there. A full risk analysis is usually carried out only when a risk management intervention is being contemplated.) Risk analysis uses probability distributions, confidence intervals, and other displays to show uncertainties about the human health consequences of different decisions. It identifies a subset of one or more decision options leading to preferred (e.g., stochastically undominated) probability distributions of health risks and other outcomes.

The outputs of a health risk analysis should allow a risk manager to answer the following questions for each risk management decision alternative being evaluated or compared:

- *What probable change in human health risk would result from each risk management intervention?* If the risk management decision option or action being assessed is implemented, how will the probable adverse human health effects (e.g., expected numbers of mild, moderate, severe, and fatal illnesses per year; expected numbers of illness-days and, if desired, quality-adjusted life-years (QALYs) lost per year) change, both in the whole population and in subpopulations with distinct risks?

- *How certain is the change in human health risk that would be caused by each risk management action?* Instead of a single value, i.e., a point estimate of risk, uncertain risks are characterized by intervals or probability distributions indicating how closely the change in human health risk caused by a proposed risk management intervention can be predicted. There are several technical options for expressing uncertainty around point estimates (e.g., plausible upper and lower bounds, confidence limits, coefficients of variation, tolerance intervals, prediction intervals, Bayesian posterior probability intervals and distributions, etc.) The essential information to provide about uncertainty in any risk assessment is how large or how small the true risks might be, consistent with the data and with the specified assumptions of the risk assessment. Point estimates that are "best" with respect to various technical statistical criteria will typically fall between these extremes.

  Technical note: Statistical point estimates and interval estimates. Many criteria have been used to define and identify "best" point estimates in risk models, e.g., maximum likelihood estimates (MLE), maximum *a posteriori* (MAP) Bayesian estimates, maximum entropy "maximum entropy",minimum description length, least squares, minimum absolute deviation, and minimum expected loss (for various loss functions) (see Cox, 2001 for a survey for risk analysts). While these criteria have led to useful theory and algorithms for estimating the parameters of risk models, *none* of them is satisfactory as the sole output from a risk assessment. *It is essential to provide intervals or probability distributions around any point estimate of risk* to inform the users of a risk assessment about the full range of risks that might be prevented (or caused) by a risk management intervention. This principle applies to qualitative and fuzzy risk ratings as well. If a point estimate of a risk is "High", then some indication must be given of how certain this value is and of how compatible the frequency

and severity components of the risk are with other qualitative labels, such as “Low.” A risk assessment that produces a single overall value for risk with no indication of uncertainty should be avoided.

- *What are the key drivers of risks and uncertainties for each option?* The analysis should make clear to the user the main reasons *why* the estimated risk from each decision option is as high or low as it is. Are the results driven mainly by predicted exposure levels, by the responses of sensitive subpopulations, by genetic or epidemiological data that establish tight constraints on the plausible values, or by other factors? Sensitivity analyses that plot how estimated risks would change as input assumptions and estimates vary within plausible ranges (e.g., within a few standard deviations of their median values) can help to identify the combinations of input values that drive the main conclusions and the extent to which these could be changed without changing the comparison of different risk management interventions.
- *Which risk management interventions are undominated?* One risk management intervention *dominates* another if it produces smaller probabilities of exceeding any specified level of adverse consequences per year. For example, if two different interventions lead to different expected numbers of sporadic illness cases per year (with the actual number being a Poisson random variable), and if the probable health consequences per case (e.g., the distribution of the number of days of illness of given severity) is the same for each intervention, then the one giving the smaller expected number of illnesses per year dominates the other. Scientific risk assessment can, at most, identify undominated risk management alternatives for risk managers to further assess and choose among.

CONCLUSIONS

This chapter has briefly described how risk analysis can promote improved risk management decision-making. A successful risk analysis estimates the causal relations between decisions and probable resulting exposures, and between exposures and their probable total human health consequences. To guide rational decision-making, a risk analysis should yield evaluations and comparisons of proposed risk management *actions and interventions*, not simply descriptions of the current situation. It should show the estimated changes in frequencies and magnitudes (and uncertainties) of human heath consequences resulting from different proposed risk management decisions. It is important to identify an adequate range of risk management options to assure that dominant alternatives are not overlooked. For each option, total health consequences are found by summing the impacts of proposed actions on human exposures over all relevant pathways that contribute significantly to the outcome. Applying an exposure-response model to the changed exposures for different decisions then yields the estimated risks associated with them.

A well-conducted risk analysis enables its recipients to participate more effectively in risk management deliberations and to communicate questions and concerns more clearly and concisely than would otherwise be possible. It does so by providing them with the relevant information needed to determine the probable consequences of proposed actions and by showing how sensitive these predicted consequences are to specific remaining uncertainties.

ACKNOWLEDGMENTS:

I thank Professor Vicki Bier of the University of the University of Wisconsin at Madison for comments and suggestions on an early draft of this chapter that substantially improved the final exposition. Parts of this chapter, including several examples, are condensed and adopted from

my book Q*uantitative Health Risk Analysis Methods: Modeling the Human Health Impacts of Antibiotics Used in Food Animals,* published by Springer in 2005 (http://www.springer.com/sgw/cda/frontpage/0,,5-40521-22-50492725-0,00.html).

DEDICATION

This chapter is dedicated to the memory of Professor Alvin W. Drake of MIT – an outstanding teacher, counselor, and friend and a source of many insights into decision analysis and risk analyses and the interplay between them.

REFERENCES

Aliferis C., Tsamardinos I., & Statnikov A. (2003). Hiton a novel Markov Blanket algorithm for optimal variable selection. *AMIA Annual Symposium Procedings*; 21-5.

Andrieu C., de Freitas N., Doucet A., & Jordan M.I.(2003) An introduction to MCMC for machine learning. *Machine Learning*, 50, 5-43. http://citeseer.csail.mit.edu/andrieu03introduction.html

Bagley S.C., White H. & Golomb B.A. (2001). Logistic regression in the medical literature: standards for use and reporting, with particular attention to one medical domain. *J Clin Epidemiol*, *54* 979-85.

Burmaster, D.E & Anderson P.D(1994). Principles of good practice for the use of Monte Carlo techniques in human health and ecological risk assessments. *Risk Analysis 14* 77-81.

Buzby, J.C., Roberts, T., Jordan Lin C.T., & MacDonald, J.M. (1996). Bacterial Foodborne Disease: Medical Costs and Productivity Losses. *USDA, Economic Research Service, Agricultural Economics Report # 741*. (http://www.ers.usda.gov/publications/aer741/)

Byrd, D.M., & Cothern CR. (2000) Introduction to Risk Analysis: A Systematic Approach To Science-Based Decision Making. ABS Consulting.

Cassin M.H., Paoli G.M., & Lammerding A.M. (1998). Simulation modeling for microbial risk assessment. *J Food Prot. 61*1560-6.

Chaudhuri P., Lo W.D., Loh W.Y., Yang C.C.(1995) Generalized Regression Trees. *Statistica Sinica* 5, 641-666. http://www.stat.wisc.edu/~loh/treeprogs/guide/grapes.pdf

Chang K.C., & Tian Z., (2002). Efficient inference for mixed Bayesian networks. http://citeseer.nj.nec.com/cachedpage/570208/1

Cheng J, Greiner R, Kelly J, Bell D, Liu W. Learning Bayesian Networks from Data: An Information-Theory Based Approach. 2001. http://citeseer.ist.psu.edu/628344.html

Cheng J, Druzdzel. 2000. AIS-BN: An Adaptive Importance Sampling Algorithm for Evidential Reasoning in Large Bayesian Networks. http://citeseer.nj.nec.com/cheng00aisbn.html

Choi BC, Noseworthy AL. (1992) Classification, direction, and prevention of bias in epidemiologic research. *J Occup Med*. Mar;34(3):265-71.

Cox L.A. *Risk Analysis: Foundations, Models, and Methods.* Kluwer Academic Press, Boston, MA.

Deeks JJ, Dinnes J, D'Amico R, Sowden AJ, Sakarovitch C, Song F, Petticrew M, Altman DG; International Stroke Trial Collaborative Group; European Carotid Surgery Trial Collaborative Group. Evaluating non-randomised intervention studies. Health Technol Assess. 2003;7(27):iii-x, 1-173.

Druzdzel , MJ, 1994. Some Useful Properties of Probabilistic Knowledge Representations From the Point of View of Intelligent Systems. http://citeseer.nj.nec.com/druzdzel94some.html

FAO/WHO, 2001. Joint FAO/WHO Expert Consultation on Risk Assessment of Microbiological Hazards in Foods: Risk characterization of Salmonella spp in eggs and broiler chickens and Listeria monocytogenes in ready-to-eat foods. FAO Headquarters, Rome, Italy 30 April - 4 May 2001
http://www.who.int/foodsafety/publications/micro/en/may2001.pdf

FDA-CFSAN, 2003. FDA/Center for Food Safety and Applied Nutrition
USDA/Food Safety and Inspection Service Centers for Disease Control and Prevention. Quantitative Assessment of Relative Risk to Public Health from Foodborne *Listeria monocytogenes* Among Selected Categories of Ready-to-Eat Foods. Appendix 3. September 2003. http://vm.cfsan.fda.gov/~dms/lmr2-a3.html

Feldman RA. Confounding factors in observational and intervention studies.
Ital J Gastroenterol Hepatol. 1998 Oct;30 Suppl 3:S248-53.

Freedman DA. Graphical models for causation, and the identification problem. Eval Rev. 2004 Aug, 28(4):267-93.

Frey L, Fisher D, Tsamardinos I, Aliferis C, Statnikov A. 2003. Identifying Markov Blankets with Decision Tree Induction**.** http://citeseer.nj.nec.com/frey03identifying.html

Friedman CR, Hoekstra RM, Samuel M, Marcus R, Bender J, Shiferaw B, Reddy S, Ahuja SD, Helfrick DL, Hardnett F, Carter M, Anderson B, Tauxe RV; Emerging Infections Program FoodNet Working Group. Risk factors for sporadic *Campylobacter* infection in the United States: A case-control study in FoodNet sites. Clin Infect Dis. 2004 Apr 15;38 Suppl 3:S285-96.

Friedman C, Reddy S, Samual M, Marcus R, Bender J, Desai S, Shiferaw B, Helfrick D, Carter M, Anderson B, Hoekstra M, and the EIP Working Group. Risk Factors for Sporadic *Campylobacter* Infections in the United States: A Case-Control Study on FoodNet Sites. 2nd International Conference on Emerging Infectious Diseases. Atlanta, GA, July 2000. http://www.cdc.gov/foodnet/pub/iceid/2000/friedman_c.htm

Friedman N, Goldszmidt, M Learning Bayesian Networks With Local Structure. http://citeseer.ist.psu.edu/friedman96learning.html

Friedman N, Goldszmidt, M. Discretizing Continuous Attributes While Learning Bayesian Networks. 1996b. http://citeseer.ist.psu.edu/friedman96discretizing.html

Gelman A, Van Mechelen I, Verbeke G, Heitjan DF, Meulders M. Multiple imputation for model checking: completed-data plots with missing and latent data. Biometrics. 2005 Mar;61(1):74-85.

Green, PJ. 1995. Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. http://citeseer.nj.nec.com/green95reversible.html

Greenland S. Quantifying biases in causal models: classical confounding vs collider-stratification bias. Epidemiology. 2003 May;14(3):300-6.

Greenland S. Modeling and variable selection in epidemiologic analysis. Am J Public Health. 1989 Mar;79(3):340-9.

Greenland S, Brumback B. An overview of relations among causal modeling methods. Int J Epidemiol. 2002 Oct;31(5):1030-7.

Greenland S, Morgenstern H. Confounding in health research. Annu Rev Public Health. 2001;22:189-212.

Grimes DA, Schulz KF. Bias and causal associations in observational research. Lancet. 2002 Jan 19;359(9302):248-52.

Guatama T, Van Hulle MM. 2003. Surrogate-Based Test For Granger – Causality. http://citeseer.ist.psu.edu/588339.html

Gurm HS, Litaker DG. Framing procedural risks to patients: is 99% safe the same as a risk of 1 in 100? *Acad Med* 2000 Aug;75(8):840-842.

Haas C.N., J.B. Rose, and C.P. Gerba. 1999. *Quantitative Microbial Risk Assessment,* John Wiley & Sons, New York.

Haas CN., On modeling correlated random variables in risk assessment, Risk Anal 1999 Dec;19(6):1205-14

Hartemink AJ, Gifford DK, Jaakkola TS, Young RA. Using graphical models and genomic expression data to statistically validate models of genetic regulatory networks. Pac Symp Biocomput. 2001;:422-33.

Hazen, G, 2003. Multiattribute Structure for QALYS. http://fisher.osu.edu/~butler_267/DAPapers/WP030018.pdf

Henrion M. Propagation of uncertainty by probabilistic logic sampling in Bayes' networks. In J. F. Lemmer & L. N. Kanal (Eds.), *Uncertainty in Artificial Intelligence 2.* Amsterdam: North Holland, 1988.

Henrion M, Pradhan M, Del Favero B, Huang K, Provan G, O'Rorke P. Why is diagnosis using belief networks insensitive to imprecision in probabilities? 1996. http://citeseer.ist.psu.edu/henrion96why.html

Hoffman FO, Hammonds JS. Propagation of uncertainty in risk assessments: the need to distinguish between uncertainty due to lack of knowledge and uncertainty due to variability. Risk Anal. 1994 Oct;14(5):707-12.

Holcomb DL, Smith MA, Ware GO, Hung YC, Brackett RE, Doyle MP. Comparison of six dose-response models for use with food-borne pathogens. Risk Anal. 1999 Dec, 19(6):1091-100.

Ibrahim JG, Chen M-H, Lipsitz SR, Herring AH. Missing-data methods for generalized linear models: A comparative review. Journal of the American Statistical Association 2005 Mar;100(469):332-346.

Janson S. Coupling And Poisson Approximation. 1994. http://citeseer.ist.psu.edu/janson94coupling.html

Jardine C, Hrudey S, Shortreed J, Craig L, Krewski D, Furgal C, McColl S. Risk management frameworks for human health and environmental risks. J Toxicol Environ Health B Crit Rev. 2003 Nov-Dec;6(6):569-720.

Kang SH, Kodell RL, Chen JJ. Incorporating model uncertainties along with data uncertainties in microbial risk assessment. Regul Toxicol Pharmacol. 2000 Aug;32(1):68-72.

Keiding N, Budtz-Jorgensen E. The Precautionary Principle and statistical approaches to uncertainty. Int J Occup Med Environ Health. 2004;17(1):147-51.

Lange K. *Applied Probability.* Springer. New York. 2003

Lemon SC, Roy J, Clark MA, Friedmann PD, Rakowski W. Classification and regression tree analysis in public health: methodological review and comparison with logistic regression. Ann Behav Med. 2003 Dec;26(3):172-81.

Loh, 2002. Regression trees with unbiased variable selection and interaction. Statistica Sinica **12**(2002), 361-386. www.stat.sinica.edu.tw/statistica/ password.asp?vol=12&num=2&art=1

Maclure M. Multivariate refutation of aetiological hypotheses in non-experimental epidemiology. Int J Epidemiol. 1990 Dec;19(4):782-7.

Maclure M.Taxonomic axes of epidemiologic study designs: a refutationist perspective. J Clin Epidemiol. 1991;44(10):1045-53.

Mark DH. Interpreting the term selection bias in medical research. Fam Med. 1997 Feb;29(2):132-6.

Mather FJ, White LE, Langlois EC, Shorter CF, Swalm CM, Shaffer JG, Hartley WR. Statistical methods for linking health, exposure, and hazards. Environ Health Perspect. 2004 Oct;112(14):1440-5.
http://ehp.niehs.nih.gov/members/2004/7145/7145.html

Miloslavsky M, van der Laan M. Fitting of mixtures with unspecified number of components using cross validation distance estimate. Computational Statistics & Data Analysis 2003 Jan;41(3-4): 413-428.

Miyamoto JM.. Quality-Adjusted Life Years (QALY) Utility Models under Expected Utility and Rank Dependent Utility Assumptions. J Math Psychol. 1999 Jun;43(2):201-237.

Ng KL, Hamby DM.. Fundamentals for establishing a risk communication program. *Health Phys* 1997 Sep;73(3):473-482.

Parsons DJ, Orton TG, D'Souza J, Moore A, Jones R, Dodd CE. A comparison of three modeling approaches for quantitative risk assessment using the case study of *Salmonella* spp in poultry meat. Int J Food Microbiol. 2005 Jan 15;98(1):35-51.

Patton DE. The ABCs of risk assessment. EPA JOURNAL 1993 19(1):10-15
http://www.bethel.edu/~kisrob/hon301k/readings/risk/RiskEPA/riskepa1.html

Phillips CV, Goodman KJ. The missed lessons of Sir Austin Bradford Hill. Epidemiol Perspect Innov. 2004 Oct 04;1(1):3.

Rai SN, Krewski D. Uncertainty and variability analysis in multiplicative risk models. Risk Anal. 1998 Feb;18(1):37-45.

Richardson S, Green PJ, 1997. On Bayesian analysis of mixtures with an unknown number of components. http://citeseer.ist.psu.edu/richardson97bayesian.html

Robinson, DA. 1981. “Infective Dose of *Campylobacter* jejuni in Milk”, *British Medical Journal*, 282, 1584.

Romano JP, Wolf M. Exact and approximate stepdown methods for multiple hypothesis testing. Journal of the American Statistical Association 2005 Mar;100(469):94-108.

Savitz DA, Greenland S, Stolley PD, Kelsey JL. Scientific standards of criticism: a reaction to "Scientific standards in epidemiologic studies of the menace of daily life," by A.R. Feinstein. Epidemiology. 1990 Jan;1(1):78-83.

Schafer JL. *Analysis of Incomplete Multivariate Data.* New York: Chapman and Hall;1997

Shipley, B. (2000). *Cause and Correlation in Biology. A User’s Guide to Path Analysis, Structural Equations and Causal Inference.* Cambridge University Press. http://callisto.si.usherb.ca:8080/bshipley/my%20book.htm

Sonis J. A closer look at confounding. Fam Med. 1998 Sep, 30(8):584-8.

Stephens M. Bayesian analysis of mixture models with an unknown number of components—an alternative to reversible jump methods Ann. Statist. 28, no. 1 (2000), 40–74

Swanson NR, Ozyildirim A, Pisu M. 2001. A comparison of alternative causality and predictive accuracy tests in the presence of integrated and co-integrated economic variables. http://citeseer.nj.nec.com/swanson01comparison.html

Thompson KM, Burmaster DE, Crouch EA. Monte Carlo techniques for quantitative uncertainty analysis in public health risk assessments. Risk Anal. 1992 Mar;12(1): 53-63.

Tsamardinos I, Aliferis C, Statnikov A. 2003. Time and Sample Efficient Discovery of Markov Blankets and Direct Causal Relations. http://citeseer.nj.nec.com/tsamardinos03time.html

Viallefont V, Raftery AE, Richardson S. Variable selection and Bayesian model averaging in case-control studies. Stat Med. 2001 Nov 15;20(21):3215-30.

Vose, DJ. *Risk Analysis: A Quantitative Guide,* 2nd Ed. John Wiley & Sons. 2000.

Wang D, Zhang W, Bakhai A. Comparison of Bayesian model averaging and stepwise methods for model selection in logistic regression. Stat Med. 2004 Nov 30;23(22):3451-67.

Weed DL. Interpreting epidemiological evidence: how meta-analysis and causal inference methods are related. Int J Epidemiol. 2000 Jun;29(3):387-90.

WHO/FAO, 2002. Risk assessments of *Salmonella* in eggs and broiler chickens - Interpretative Summary and Full Report. http://www.who.int/foodsafety/publications/micro/Salmonella/en/

Williamson J. *Bayesian Nets and Causality Philosophical and Computational Foundations.* Oxford University Press. 2005.

Yokota F, Thompson KM. Value of information analysis in environmental health risk management decisions: past, present, and future.Risk Anal. 2004 Jun, 24(3):635-50.