4. COMPARATIVE RISK FRAMEWORK METHODOLOGY

4.1. INTRODUCTION

This chapter describes a comparative risk framework methodology (CRFM) which combines the NAS risk assessment paradigm with cost-effectiveness analysis approach (NAS, 1983; Haddix et al, 1996; Gold et al., 1996). The resulting approach provides a systematic way to assess interdependent environmental health risks and compare the impact of alternative interventions on those risks. The purpose of this chapter is to describe the key steps in the methodology and to identify central issues which must be addressed in developing an analysis of this type. This is not intended to be a prescriptive or exhaustive discussion of the methodology. The reader should bear in mind that the methodology was developed for application to drinking water treatment systems. Other issues may need to be addressed for other types of applications.

4.1.1. NAS 1983 Risk Assessment Paradigm. To assess risks posed by many different chemical pollutants and classes of chemical pollutants, U.S. EPA has relied on the four step process of risk assessment. This four step approach, known as the risk assessment paradigm, was originally presented by the National Academy of Sciences in 1983 (NAS, 1983). In the paradigm, the nature of the problem and the potential toxic effects are first defined and the data needs are determined in the *hazard identification* step (see Figure 4-1). The next step, *dose-response assessment*, consists of quantifying the relationship between the adverse health effect and exposure to or dose of a toxic agent, and characterizing the overall scientific confidence in these data. The *exposure assessment* step consists of delineating the pathways and routes of exposure that relate a source to an exposure event, identifying exposed individuals, quantifying

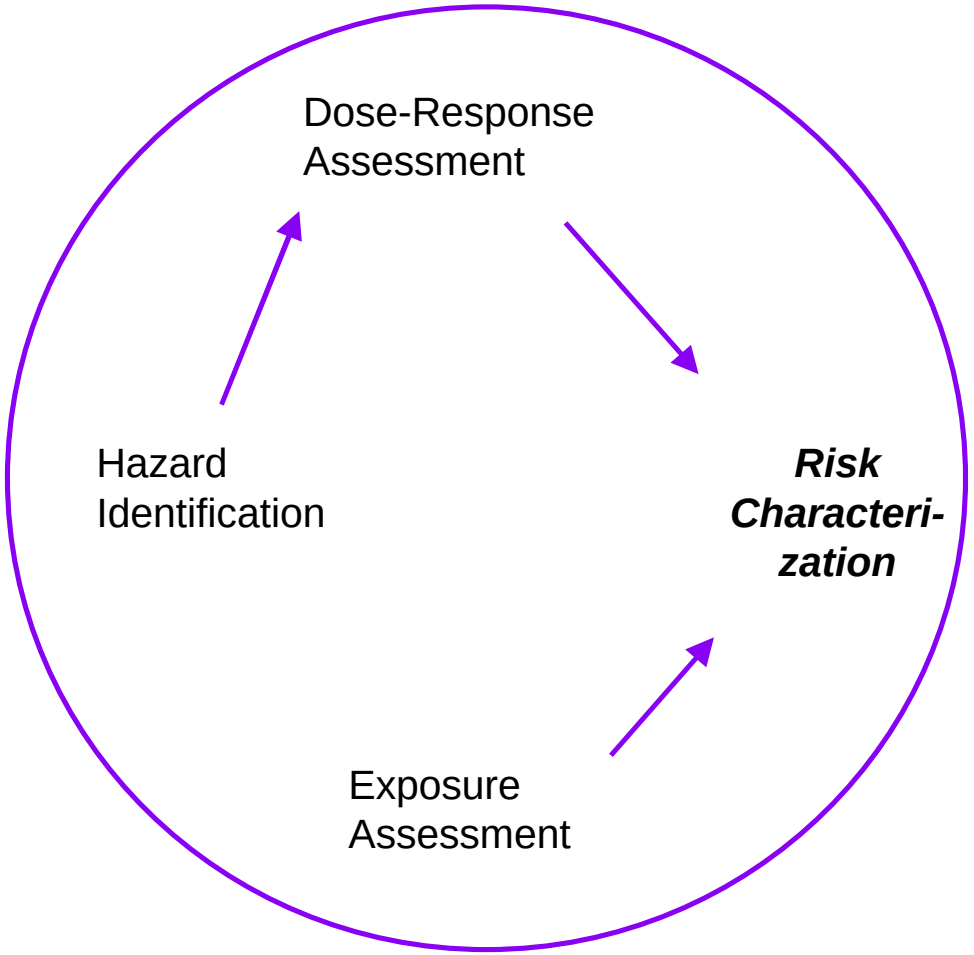

Figure 4-1. This figure illustrates the risk assessment paradigm as described in 1983 by NAS. Figure adapted from NAS, 1983

concentrations in exposure media, quantifying the magnitude and duration of the exposure, and characterizing the overall scientific confidence in these data, e.g., identifying sources of uncertainty. Finally, in *risk characterization*, the major scientific evidence and "bottom-line" results from hazard identification, dose-response assessment, and exposure assessment are evaluated and integrated. This is translated into an overall conclusion about the likelihood of an adverse health effect occurring due to the presence of a potential hazard (NAS, 1983; NRC, 1994).

This general approach to risk assessment has been codified in the U.S. Federal Government including the U.S. Environmental Protection Agency (U.S. EPA, 1989, 1992a). Use of the risk assessment paradigm is a well-accepted approach for assessing health risks from chemicals, particularly those whose presence has no known positive human health impact, e.g., particulates, dioxins or mercury. The risk assessment paradigm has been adapted to assess risks posed to ecological systems, and risk assessment approaches for non-chemical entities such as pathogenic microorganisms are also being developed.

Risk assessments have been used to estimate plausible levels of risk that might result from exposure to environmental pollutants. These assessments are frequently used to guide regulatory decisions regarding the level of exposure that results in 'acceptable' health risks. When the exposure level exceeds the acceptable level, action is taken to reduce the concentration of the pollutant(s). While this use of risk assessment serves as a starting point, additional approaches are needed to help decision makers determine how to best manage risks. This is particularly true when the alternative strategies or interventions for reducing a health risk may carry costs or health risks of their own.

4.1.2. Comparing Risks. From a traditional risk assessment perspective, a major impediment to successfully undertaking comparative risk analysis has been the need to express all the health outcomes in terms of a common measure. It is difficult to compare disparate health outcomes such as cancer, which is severe and generally occurs once late in life, and gastrointestinal infection, which generally produces only mild disease that may occur repeatedly throughout life. Because of this, these risks are typically considered in isolation, ignoring the potential tradeoffs between risks. However, methods are available to compare the risk of a case of cancer to the risk of a case of diarrhea caused by gastrointestinal illness.

One way to facilitate a comparison is to simply identify characteristics associated with these different health endpoints that provide information about the risk in a clear and direct way. This type of approach has been recommended by Graham and Weiner (1995) and builds on initial suggestions of Lave (1981). In order to facilitate a credible comparison of the risks, Graham and Weiner (1995) recommend focusing on the following specific descriptive characteristics of the outcomes of concern: 1) the magnitude, or probability, of the risk; 2) the size of the population affected; 3) the certainty in the risk estimates; 4) the type of adverse outcome; 5) the population distribution of the risk; and 6) the timing of the risk. In addition, it is important to consider the duration and severity of the health outcome.

While describing each risk individually can provide important insights, several approaches allow us to compare the risks directly using a common metric. To address the problem of finding a common outcome measure that allows for the comparison of disparate risks, mortality and incidence rates can be converted into indicators of life lengthening or life shortening. These types of measures have been commonly used in medical decision making, public health, and economics

(Putnam and Graham, 1993; Gardner and Sandborn, 1990), and are usually expressed as lives or life-years lost or gained.

Expressing different health status changes in terms of lives or life years lost allows for the comparison of different mortality risks associated with alternative interventions. However their usefulness is limited when comparing mortality risks (e.g., cancer death) to morbidity (diarrheal illness).

Alternately, these measures can be adjusted (i.e. weighted) by subjective quality of life estimates to generate Years of Potential Life Lost (YPLL) (CDC, 1982; Gardner and Sandborn, 1990), Disability Adjusted Life Years (DALYs) (Murray, 1994), Quality Adjusted Life Years (QALYs) (Putnam and Graham, 1993), or Years of Healthy Life (YHL) (Erickson, 1995). The weighting or adjustment factors reflect some social valuation of the outcome characteristics which may include a decreased length of life from mortality and a decreased quality of life from morbidity associated with a particular condition. The use of these outcome measures in the evaluation of public health and medical interventions suggests their potential usefulness to environmental health interventions such as drinking water disinfection decisions.

4.1.3. Comparing Public Health Interventions. During the past years several efforts have been made to develop systematic approaches to assess a wide range of public health and medical interventions in this type of framework, generally referred to as cost-effectiveness analysis. In an effort to improve the impact of limited public health resources, these methods have been applied to a wide range of interventions including immunization delivery, sexually transmitted disease prevention, and injury prevention. In policy analysis these tools are used to identify and quantify

the potential health and economic outcomes of alternative health interventions. These results can be used to compare and choose among those alternatives.

Haddix et al (1996) provide a systematic description of the application of cost-effectiveness analysis and related methods to the evaluation of public health interventions, as implemented at the Centers for Disease Control and Prevention (CDC). Similarly, Gold et al (1996) present the findings of an multi-disciplinary panel appointed by the U.S. Public Health Service to define a set of 'best practices' when using cost-effectiveness to facilitate resource allocation decisions regarding health and medical interventions. This section draws on the recommendations of these authors to suggest how cost-effectiveness analysis can be applied to comparative risk assessment of environmental public health interventions. This is not intended to be an exhaustive discussion of the approaches, and the reader is referred to the original publications for a more complete discussion of their application.

This approach may also be useful as a way of gauging the full impact of making more stringent regulatory guidance levels. Any decision to mandate lower DBP levels will result in changes in the way drinking water is treated. Possible changes include alternate treatment technologies (e.g., membrane filtration), disinfectants (e.g., ozone, chloramine, etc.) and/or disinfectant concentrations, and time of treatment. The CRFM permits an evaluation of all the potentially negative impacts related to these changes, including microbial risk and cost considerations, and contrasts them to the potential benefits that might result from the mandated DBP reduction. Although cost-effectiveness analyses and related methods can be very useful, it is important to remember that they are meant to augment, not replace, the decision making process. The final choice of water treatment option will be based on many community and purveyor

considerations such as budget, public opinion, time and personnel resources, and regulatory constraints.

4.2. COST EFFECTIVENESS ANALYSIS IN PUBLIC HEALTH

The CRFM incorporates many of the principles of cost-effectiveness analysis (CEA) as it has been applied to the assessment of alternative public health interventions. The purpose of a cost-effectiveness analysis (CEA) is to systematically and quantitatively assess the expected outcomes and resource costs of alternative interventions in a way which facilitates decision making by different stakeholders. Cost-effectiveness analyses of public health interventions typically rely on information from a wide range of sources, including epidemiology, clinical microbiology, clinical trials, and intervention studies, to estimate the expected health outcomes of a series of alternative programs, policies, or practices. The approach can be expanded to include information from other disciplines (e.g., engineering and risk assessment) regarding the expected impact of alternative strategies. This is then combined with an economic analysis of the costs of the intervention to provide decision makers and others with an estimate of the combined impact of the alternatives. This section summarizes the analytical tools commonly used in cost-effectiveness analyses of public health interventions (Haddix et al., 1996; Gold et al., 1996).

4.2.1. Analytical Methods. While a variety of terms and labels are used by different authors, two basic analytical approaches have been used to assess alternative public health interventions in this type of framework. In both approaches the costs of alternative interventions (expressed in dollars) are compared to the associated health outcomes. The approaches differ depending on whether health outcomes are expressed in natural units (e.g., cancer cases averted), quality adjusted health utility measures (e.g., QALYs), or dollar equivalents. In this document, studies

which compare resource costs to health outcomes (cost per case averted or cost per QALY), are referred to as cost-effectiveness analyses. The second approach, expressing intervention costs and outcomes in dollars, is referred to as cost-benefit analysis. The CRFM and case study use cost-effectiveness analysis methods. However, both CEA and CBA can be used to systematically compare the costs of interventions with their effect on morbidity and mortality from multiple risks, and either method could be used in the CRFM. .

4.2.2. Cost-Effectiveness Analysis (CEA). In CEA using natural units for health outcomes, the net cost of the intervention is compared directly to the disease events averted, e.g. costs per case of cancer, death, or injury. These are most useful when comparing a series of interventions designed to affect a single health outcome. For example, this may be useful when comparing three different strategies for reducing new cases of HIV among teenagers. In more complicated studies, where more than one health endpoint is affected by the interventions, reporting disaggregated numbers of events (along with other outcome metrics) can add clarity and transparency to an analysis.

Other cost-effectiveness analyses build quality-of-life adjustments into the calculations, in order to express all health outcomes in terms of a single metric. Benefits are expressed in terms of the number of life-years saved, adjusted for any change in quality associated with the health outcomes linked to each intervention. The most common outcome measure for this type of analysis is the QALY. Methods for calculating QALYs are presented below.

A cost effectiveness analysis designed to compare alternative interventions to a baseline is called an incremental analysis. An incremental analysis considers the *additional* cost of an intervention compared to the baseline, and the *additional* benefits associated with the intervention

compared to the baseline. The resulting cost-effectiveness ratio (cost per QALY or cost per case averted) reflect this incremental benefit and cost compared to the baseline, rather than the total costs and benefits associated with the intervention.

The current case study is an incremental CEA that uses a QALY utility measure. This provides a common metric for the disparate health outcomes related to drinking water treatment. It also allows an evaluation of the trade-offs between morbidity and mortality and incorporates the preferences that society has for different types of morbidity, which are all key components of any decision concerning public health interventions. The major disadvantage of this approach is that QALYs represent subjective estimates that vary among individuals and populations (Farnham, 1996).

4.2.3. Cost-Benefit Analysis (CBA). Cost-benefit analysis expresses the consequences of an intervention totally in monetary terms. All programmatic costs and benefits (lives saved, etc.) are converted to dollar values that reflect the estimated impact on societal resources. Various types of cost-benefit analyses have been performed previously for treatment technologies and treatment goals (Clark and Adams, 1993; Clark et al., 1993; Regli et al., 1993; Craun, 1996; U.S. EPA, 1994), these analyses provide useful data on the monetary costs and benefits associated with controlling either DBPs or pathogens.

A traditional cost-benefit analysis could be developed based on the current case study. The purpose of using a CBA for the current problem would be to assess and compare the monetary value of each of the costs associated with alternative drinking water treatment choices. These include the implementation and annual operating costs associated with each treatment technology, the costs associated with lost aesthetic water quality, and the health-related costs

associated with the predicted effects of exposure to drinking water contaminants. These health-related costs include the costs of medical care, lost income, and pain and suffering associated with gastrointestinal disease, cancer, adverse reproductive outcomes, or other diseases that might be caused by drinking water contamination.

Some costs have a clearly defined monetary value (e.g., medical costs, lost productivity (income) costs, and technology implementation costs) because they are directly purchased with dollars (i.e., they are marketable goods). Others, such as the value of avoiding pain and suffering are not traded in the market place and are less easily quantified. The results of a CBA are expressed most usefully as net benefits (i.e., program benefits minus program costs). The advantage of this is that the net benefits can then be used to compare interventions with other expenditures, including those whose purpose is the accrual of non-health related benefits (e.g., educational expenditures).

4.3. COMPARATIVE RISK FRAMEWORK METHODOLOGY

4.3.1. CRFM Overview. The framework presented and applied in this document represents the integration of cost-effectiveness analysis as applied to public health interventions with the 1983 NAS Risk Assessment Paradigm. The CRFM provides a systematic approach to assessing the expected impact of specific treatment alternatives on microbial and chemical risks, and comparing these disparate risks to the costs associated with the alternatives. Figure 4-2 provides an overview of the CRFM. Each of the elements of the figure represents a step in the process of combining risk assessment and cost-effectiveness analysis in this framework, and is briefly described below.

Figure 4-2. Comparative Risk Assessment Framework Overview

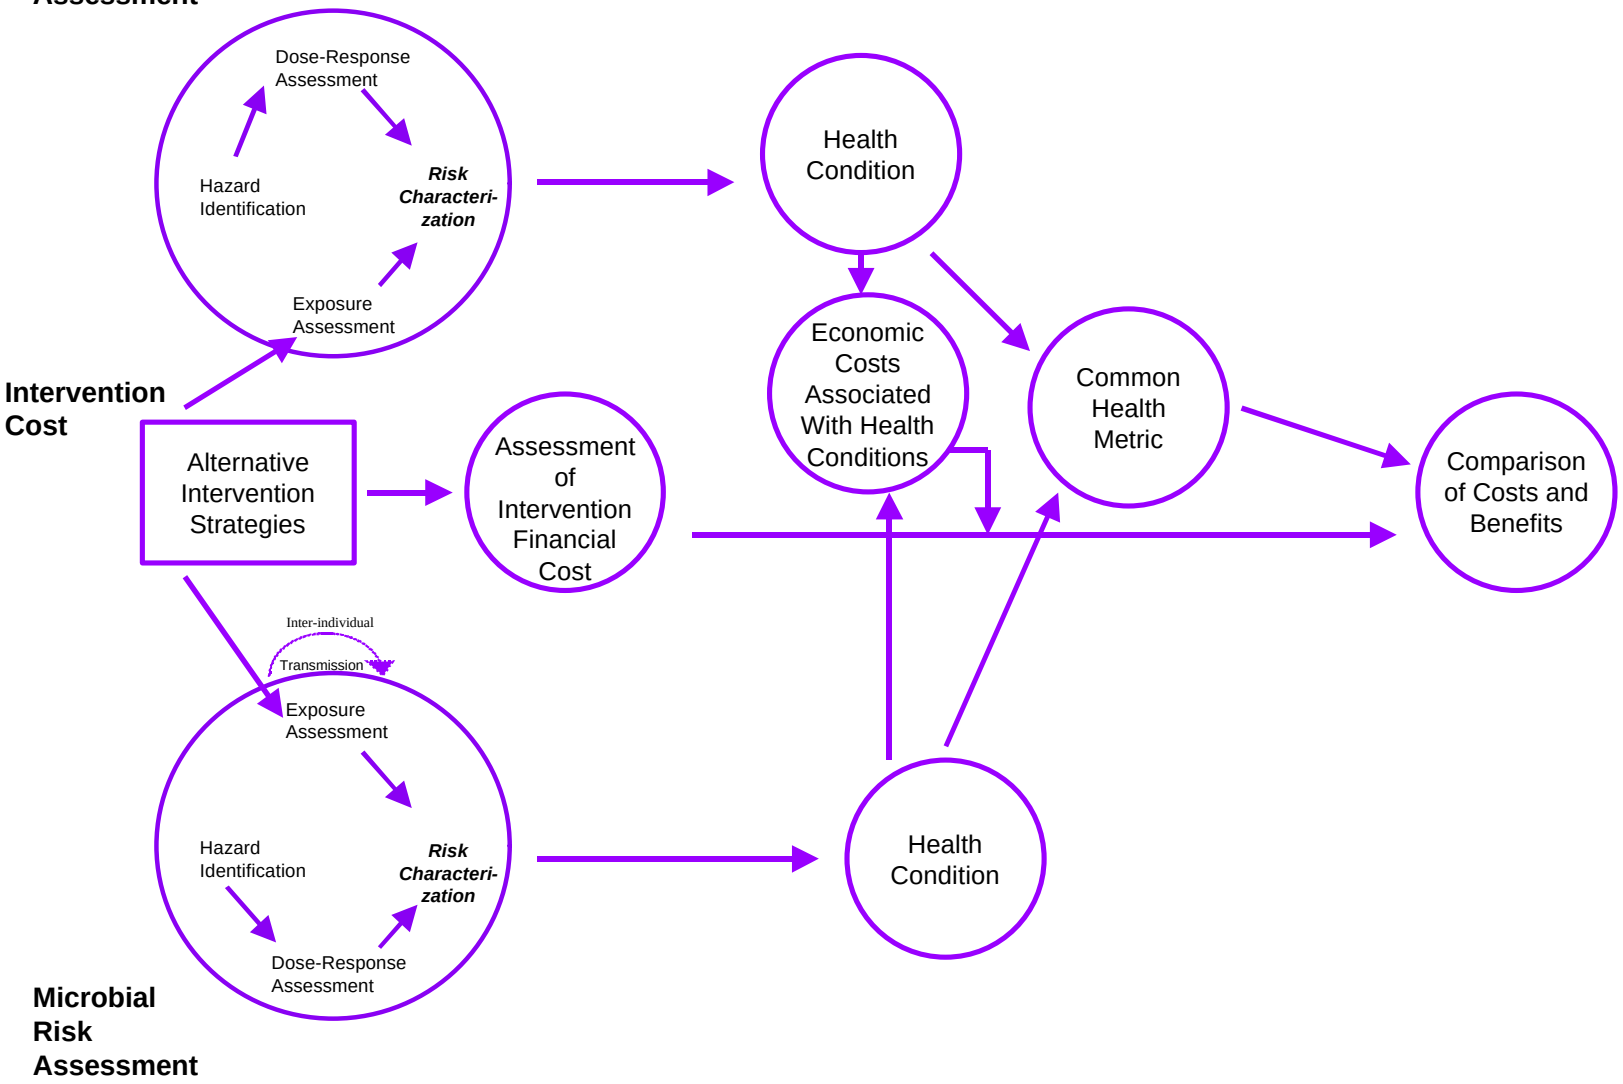

The starting point for the CRFM is the 1983 NAS Risk Assessment Paradigm (Figure 4-1), where exposure and dose-response assessment are combined to characterize chemical and microbial risks (circles in Figure 4-2 labeled “Chemical Risk Assessment” and “Microbial Risk Assessment”). The assessment of these risks requires a consideration of microbial and DBP concentrations in finished water, tap water intake rates, and dose-response assessment for the microbial and chemical agents.

To compare chemical and microbial risks, the effects estimated in the risk characterization must be expressed as human health conditions (circles labeled “Health Condition”). For microbial risks this requires a description of the potential types or stages of microbial illness that may result from exposure, and the specific symptoms associated with each. For risks associated with DBPs, specific cancer, reproductive, and developmental outcomes must be identified as the most likely or relevant human health conditions.

The range of potential human health conditions can then be converted into a common health metric. In this framework, QALYs are used to capture changes in the length and quality of life associated with the different health conditions (circle labeled “Common Health Metric”). In addition, the different health conditions can also result in economic costs associated with medical treatment and lost productivity (circle labeled “Economic Costs Associated with Health Conditions”).

The CRFM assesses microbial and chemical risks in relation to specific treatment alternatives. As a result, an analysis must consider the estimated impact of baseline and alternative strategies on the different risks under consideration. In this case, treatments affect contaminant concentrations, exposures, and ultimately the type and number of in expected health

outcomes. In addition to affecting the different health risks, the alternative treatment strategies also require different resource expenditures for implementation. In the final stage of the framework, alternative strategies are compared by assessing their expected impact on health (QALYs) and economic outcomes (circle labeled "Comparison of Costs and Benefits").

The remainder of this chapter describes the elements of the CRFM. In each of the subsequent sections, one element or step of the framework (as shown in Figure 4-2) is described in greater detail. For each step, the discussion includes the identification of specific components that must be included in any application of the methodology, an identification of issues that arise in when applying this approach to the drinking water example, and suggestions of methodological approaches for addressing these issues. Figure 4-3 provides a graphical representation of the framework which highlights key issues that arise in applying the framework to the drinking water treatment example. Note that Figure 4-3 omits assessment of economic costs associated with various health effects because this step is omitted from the case study described in Chapters 5 and 6. This step could be incorporated into a more complete analysis. The discussion of these issues is directed particularly at the application of the framework to water treatment decisions, but may be applicable to other comparative environmental health risk problems.

4.3.2. Problem Formulation. In order to develop an application of the CRFM, a number of issues must be addressed or defined in advance. These initial steps are similar to those needed for any cost-effectiveness analysis of a public health intervention, and are drawn from Gold et al (1996) and Haddix et al (1996). In addition to the general steps in framing an analysis, the

Figure 4-3. Issues Addressed in Comparative Risk Assessment Framework

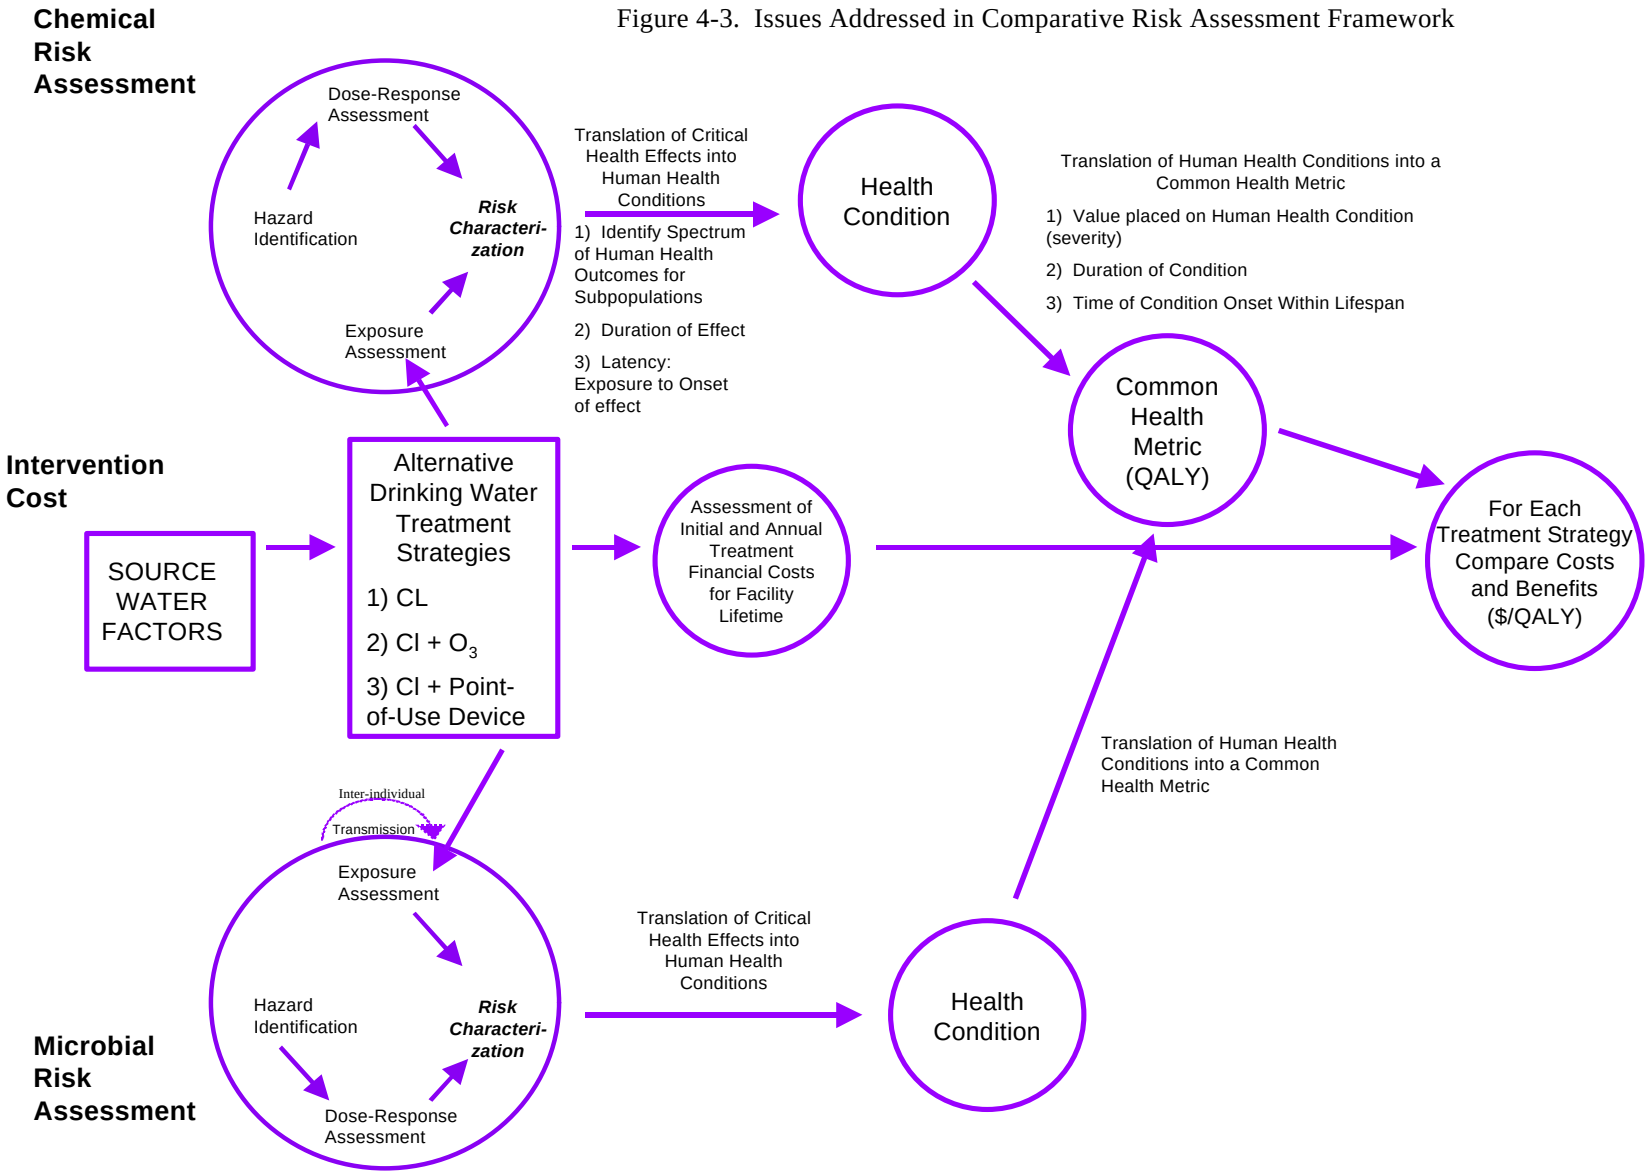

discussion raises specific issues which arise in the application of these methods to environmental public health interventions, particularly those with comparative risks.

4.3.4.1. Study Question or Objective — The objective of any cost-effectiveness analysis of public health interventions is to identify and quantify the expected outcomes associated with alternative interventions, in order to clarify the choices faced by the decision makers. The problem or question to be analyzed must be operationally defined so that the comparison is clearly and unambiguously identified. The study question takes a form such as "what is the cost-effectiveness of the new treatment technology Y, compared to standard treatment technology X?" For example, the current case-study evaluates the cost per QALY of adding a pre-ozonation step to a traditional post-chlorination treatment train. Later, following input from all target audiences, the approach could be expanded to evaluate additional treatment options or health outcomes.

4.3.4.2. Audience - Primary, Secondary — The audiences for an analysis are generally considered to be the consumers of the study results. The primary audience for an application of the CRFM is individuals or organizations who will directly use the study results in making decisions regarding the alternative interventions under consideration. Secondary audiences may include parties who will be affected by these decisions, in particular individuals whose health is affected. The analysis must be designed and presented with clarity and transparency to ensure that these audiences can understand the potential implications for them.

For the current case study, the primary audience would be the policy and program decision makers who would choose or recommend the treatment technologies that most efficiently balance the costs and the risks from DBPs and waterborne pathogens for a hypothetical water system under consideration. Secondary audiences would include the stakeholders that have

vested interests in issues related to drinking water treatment. Most important among these are the groups whose health is directly affected by the water treatment decisions. Although other audiences (the public, other scientists, etc.) may also have interest, the decision makers are the primary focus for any analytic results.

4.3.4.3. Perspective - Primary and Secondary —The perspective of the analysis must be specified because it determines which costs and benefits are included in the analysis. Farnham et al. (1996) explain that most economic analyses take the societal perspective, analyzing all benefits of a program no matter who receives them, and all costs of a program, no matter who pays them. Since the purpose of the analysis is to estimate and clarify the expected relevant outcomes of a decision for a particular audience, the perspective of the analysis should be selected to match the needs of the intended audiences. The societal perspective is considered appropriate for most public health studies because the goal of the research is to analyze the allocation of societal resources among competing activities. Therefore, the societal perspective has been taken for the current comparative risk approach to drinking water treatment (i.e., risks to the general population, including sensitive or resistant subgroups).

While the societal perspective is appropriate for the primary audience of the case study, there are secondary audiences which also need to understand the impact of alternative interventions on them. The analysis should allow them to understand the options and outcomes from their perspective. These secondary perspectives are likely to be particularly important in comparative risk problems where different risks affect different sub-populations. In these cases different perspectives may lead to different results.

Presenting disaggregated results (based on who receives the benefits or bears the costs) can provide secondary audiences with the necessary information to assess the expected relevant outcomes from their perspective. The accompanying case study presents these disaggregated results as well.

4.3.4.4. Time Horizon — The relevant time frame and analytical horizon for the analysis must be defined. The intervention time frame is the specified period in which the intervention strategies are actually applied. The analytical horizon refers to the period over which the costs and benefits of the health outcomes, which occur as a result of the intervention, are considered. The intervention time frame and analytical horizon must be long enough to account for several important factors, including seasonal variations in costs or outcomes, and any latency period between exposure and disease. The two periods do not have to be the same.

For the drinking water analysis, both the treatment technology costs and the resulting changes in microbial or chemical exposure levels occur over the expected lifetime of the treatment plant(s). The latency period for the resulting benefits (or harms) of the these exposures differ for each of the health endpoints. Diarrheal illness caused by microbial infection occurs very close to the time of exposure. However the cancer and developmental effects may occur or continue well after the actual period of exposure.

4.4. TRANSLATING RISK CHARACTERIZATION TO HEALTH CONDITIONS

This section, and each of the subsequent ones, addresses one step in the process of combining risk assessment and cost-effectiveness analysis of drinking water treatment decisions in the CRFM. The starting point is an assessment of the separate microbial and chemical risks as described in the 1983 NAS Risk Assessment Paradigm, which would include a consideration of

chemical and microbial concentrations in drinking water, tap water consumption, and dose-response relationships for the different contaminants. Each step (as shown in Figure 4-3) adds an additional layer to each of the separate microbial and chemical risk assessments, eventually resulting in a comparative analysis of the impact of alternative strategies on both risks.

4.4.1. Issues in Translating Risk to Health Conditions. A number of specific issues arise in translating risk assessment results into specific health conditions (see Figure 4-3). For microbial risks this entails identifying and describing the relevant types and stages of illness. For chemical risks this requires the conversion of critical health effects information into likely or representative human health conditions. For both types of risks, the following central issues must be addressed.

Identify the spectrum of human health outcomes - This requires identifying the specific health outcomes (e.g., likely reproductive effects and cancer endpoints) and the specific populations at risk (or disproportionately at risk) for the outcomes. For each health outcome there is a spectrum of specific health conditions which might result (e.g. mild illness, severe illness, death). Each of the these conditions must be defined to characterize the distribution of potential outcomes.

Duration of outcome - The duration of each type or stage of health condition must be estimated. Reversible effects, such as mild microbial illness or temporary infertility may last for relatively short periods. Others may affect the quality or length of an individual's life expectancy. The severity of a condition is a combination of its magnitude and duration.

Latency of outcome - The time between exposure and the likely onset of a health event can differ greatly between conditions. The analysis must determine whether conditions are likely

to occur in the same time period as the exposure. If the event is likely to occur in a later period, the length of the latency period must be estimated.

4.4.2. Selecting Outcomes. The principal health outcomes that are included in an analysis are determined by the intended effect of the interventions under consideration. However interventions may set off a cascade of secondary health outcomes (some beneficial and others not), which may be included in the analysis. These may include the adverse effects associated with the intervention, or reducing the secondary transmission of an infectious disease as a result of the intervention.

It may not be feasible to directly incorporate all potential health consequences of the interventions being considered. The analysis must, however, include those outcomes which are *relevant* for the decision being made, and those outcomes which may be borne by a particular subset of the population. Potential outcomes that are not included should be identified along with the reason for their exclusion.

In the case study presented here, microbial and chemical exposures may result in a variety of reproductive, developmental, cancer, and microbial disease endpoints. For practical reasons, the analysis focuses on those specific endpoints or conditions which are likely to account for the majority of the mortality and morbidity. They are temporary infertility, developmental defects, bladder, rectal, and colon cancer among the DBP-induced endpoints, and *Cryptosporidium* induced illness and mortality among the microbial endpoints.

4.5. VALUING HEALTH OUTCOMES USING A COMMON HEALTH METRIC

Once all relevant health conditions have been identified and characterized, they must be expressed using a common metric to facilitate their comparison. In this case, Quality Adjusted

Life Years (QALYs) are used as the common metric. The section below describes the method for calculating QALYs based on information about the specific health conditions. In order to develop these estimates, the following issues must be addressed for each of the conditions: 1) determine the value or weight associated with the health condition, 2) account for the duration of the condition, 3) account for the timing of the event within the lifetime of the individual, and 4) account for the latency period delaying the manifestation of the effect following exposure.

A number of survey instruments have been developed to capture changes in health-related quality of life, along with changes in chronic morbidity or mortality, as expressed as QALYs (Patrick and Erickson, 1993; Torrance, 1987). Among the methods used to measure health status are the Quality of Well-Being Scale, the EuroQual, and the Health Utilities Index. These indices combine objective and subjective factors to create a univariate scale which measures overall well-being. These measures provide a systematic way to assess changes in health-related quality of life, and compare it to other sources of morbidity or mortality (Hadorn and Hays, 1991). The assignment of utility scores to various outcomes makes it possible to compare those outcomes using a common metric. The current case study assigns each adverse health outcome a QALY "cost" equal to the number of years in *perfect* health an individual would be willing to sacrifice in order to avoid that adverse health effect.

In practice, the "cost" of an adverse health condition has two components. The first component is the extent to which the health condition compromises an individual's quality of life. This component is quantified by assessing the number of years in perfect health that the individual finds equally desirable to living the remainder of that individual's expected life-span with the health condition under consideration. As the health conditions we are evaluating are *less*

desirable than perfect health, individuals are indifferent between spending the remainder of their lives with the health condition under consideration and living for some *shorter* period of time in perfect health. The life years in perfect health that the individual is willing to sacrifice in exchange for elimination of the health condition under consideration is the "cost" of the health condition in QALYs. The second component is the value of the life years lost if the health condition under consideration shortens an individual's life-span. Specifically, the "cost" of this component is the number of lost life years multiplied by the QALY value of each year lost.

For example, consider a 25 year-old individual expected to live until age 75. Her health at present and in the future is not perfect, but is typical, meaning that she suffers from a range of mild ailments (colds, etc.) typical of the general population. Suppose that she is indifferent between living the remaining 50 years of her life in "typical" health, and living for 47.5 years in an ideal state of "perfect" health. This preference implies that the value she places on each year of life is 47.5 QALYs ÷ 50 years, or 0.95 QALYs per year.

Now suppose that this 25 year-old individual contracts a serious illness. As a result of the illness, the individual suffers from permanent symptoms that decrease the value of each year of life from 0.95 QALYs to 0.85 QALYs. That is, the individual is now willing to sacrifice even more of her anticipated life span in exchange for perfect health; specifically, she is indifferent between living "X" years with these symptoms and 0.85 × "X" years in perfect health. In addition, the illness shortens the individual's life-expectancy by 10 years so that she is now expected to die at age 65 instead of at age 75.

The cost of the illness is the difference between the value of the individual's remaining life years (from her own perspective) *prior* to the illness and the value of her remaining life years *after*

contracting the illness. Prior to the illness, she placed a value of 0.95 QALYs on each of her remaining 50 life years, yielding a total value of 0.95 QALYs per year × 50 years, or 47.5 QALYs. She now places a value of 0.85 QALYs on each of her remaining 40 life years, yielding a total value of 0.85 QALYs per year × 40 years, or 34 QALYs. Hence, the cost of the illness is 47.5 QALYs - 34 QALYs, or 13.5 QALYs.

This result can also be calculated by considering the quality of life component and change in life expectancy component separately. The cost of the illness due to *shortened life expectancy* is the number of years of life lost (10) multiplied by the value each life year would have had in the absence of the illness under consideration (0.95 QALYs), yielding 0.95 QALYs per year × 10 years, or 9.5 QALYs. The cost of the illness under consideration due to the *decreased quality of life* is the number of years the individual will live with this illness (40) multiplied by the lost value of each remaining life year (0.95 QALYs - 0.85 QALYs = 0.10 QALYs). Hence, the cost of the illness due to decreased quality of life is 40 years × 0.10 QALYs per year, or 4.0 QALYs. The sum of these two cost components is 13.5 QALYs, as calculated in the preceding paragraph.

Table 4-1 summarizes the example just discussed. Note that the cost of the illness is the difference between the two total QALY values in the last row of the table.

For the purpose of this case study, it is assumed that QALYs have the properties of a "utility measure." As such, QALY scores can be assigned to uncertain outcomes by calculating the expected QALY loss associated with that outcome. For example, suppose one wishes to compare two scenarios. In the first, there is a 90% chance that an individual will live the remaining 50 years of his or her life in perfect health. There is also a 10% chance that the individual will contract a condition with a cost of 5 QALYs. Therefore, the expected cost for

TABLE 4-1

Summary of Calculated Cost of Illness in QALYs

Age Range	Life Value Without Illness			Life Value With Illness		
	Value of Each Year (QALYs/Yr)	Number of Years	Total Value (QALYs)	Value of Each Year (QALYs/Yr)	Number of Years	Total Value (QALYs)
25-65 65-75	0.95 0.95	40 10	38 9.5	0.85 0 (dead)	40 10	34 0
Total			47.5			34

this scenario is (90% × 0 lost QALYs) + (10% × 5 lost QALYs), or 0.5 QALYs. The second scenario has only one possible outcome. Specifically, the individual is certain to contract a condition with a cost of 1 QALY. That is, the cost of the second scenario is 1 QALY. In this case, the first scenario has a lower cost, and is therefore preferable.

Strictly speaking, the assumption that QALYs are a utility measure is only valid if it is assumed that individuals are "risk neutral." However, in reality, the relative desirability of alternative scenarios may not be consistent with their expected QALY scores. Specifically, the QALY score for a scenario does not, in general, reflect the degree of an individual's risk aversity. For example, the following two possibilities have the same QALY score:

- A 99% chance of 10 years in perfect health followed by painless death and a 1% chance of instant and painless death. The value of this option is 99% × 10 QALYs + 1% × 0 QALYs, or 9.9 QALYs.
- A 100% chance of 9.9 years in perfect health, followed by painless death. The value of this option is 9.9 QALYs.

Although the two options have the same values, a "risk averse" individual may prefer the second option in order to avoid the 1% chance of instant death. It is for much the same reason that individuals purchase fire insurance; they prefer a certain small loss every year (the premium) to a large probability of no loss (no loss due to payment of a policy premium), and a small probability of a large loss (destruction of the property without insurance).

Despite this limitation, health economists use QALYs extensively to compare the desirability of various policies or options that affect health. One reason is that it is not difficult for many individuals to grasp the time-tradeoff concept underlying the QALY, making it easier to gather information representing the true preferences of a larger number of individuals. Assessing "true" utility values for health conditions requires a greater sophistication on the part of a surveyed individual since that individual must choose between a certain health outcome and a lottery in which there is a probability p that he or she will live in perfect health, and a probability of 1-p of instant and painless death. (The value of p that makes the individual indifferent to the two options is the utility of the health condition.)

4.6. ASSESSING THE IMPACT OF ALTERNATIVE INTERVENTIONS

4.6.1. Identifying Interventions - Baseline and Alternatives. The strategies being evaluated must be clearly identified and contrasted with the baseline strategy, which is the strategy that best represents current practice. The choice of baseline strongly influences the relative benefit of the alternative strategies, and it must be selected carefully to ensure that the resulting comparisons are realistic and relevant given the study objectives. In particular, the results of the analysis are relevant or applicable to the extent that the baseline strategy accurately reflects the current status of a particular area or community.

In selecting alternative interventions, the broadest practical range of realistic options should be considered. The nature of the alternatives is likely to be influenced by the baseline that is chosen. In some cases alternatives may add or eliminate discrete elements of the baseline strategy, such as adding an element to an existing water treatment train. In other cases, the alternatives may represent changes in a continuous attribute of the baseline, such as changing the level of pre-ozonation or contact time in a water treatment train. The current case study compares a baseline chlorination treatment train to two alternatives: baseline plus pre-ozonation and baseline plus point-of-use filters.

4.6.2. Issues in Assessing Intervention Impact. Several factors will affect the estimated impact of baseline and alternative interventions on health outcomes. In developing an analysis in this framework, these issues must be considered in assessing intervention effectiveness.

- Source water concentration of microbial agents for the particular water system
- Treatment efficacy of alternative interventions for microbial agents
- DBP concentrations for the specific alternatives under consideration under local conditions
- Size of the total population and specific susceptible subpopulations within the community being served

4.6.3. Using Tools from Decision Analysis to Estimate Intervention Outcomes. Any cost effectiveness analysis requires a realistic estimate of the expected outcomes of the baseline strategy and the alternatives under consideration. Studies of the cost effectiveness of medical and public health interventions have relied on a wide range of techniques to estimate these expected outcomes. In the case of some clinical interventions and a limited number of population-based interventions, randomized controlled trials allow the analyst to directly link a given intervention to

a set of empirically estimated outcomes. However, more often it is necessary to compile information from a range of sources in order to create a series of indirect links between a particular intervention and a set of expected outcomes which can be estimated.

A number of tools are available for structuring and combining quantitative estimates of these indirect links in order to estimate the expected outcomes of an intervention. Cost effectiveness analyses of public health interventions often employ decision analysis methods. A discussion of this use of the decision analysis tools, in particular decision trees, is presented here as one potential method for structuring the assessment of an intervention's effectiveness.

The underlying theory of decision analysis is that of expected utility, defined as the value of a preferred outcome when viewed from a particular perspective. This theory assumes that, if given a choice among different options, the option with the highest expected utility should be selected. Decision analysis is a technique for helping decision makers identify the option with the greatest expected utility from the perspective of a particular individual, society or community (Snider et al., 1996). The expected utility of any particular option is calculated from estimates of the probability of the different possible outcomes and the strength of the preference for each outcome. This preference, expressed in terms of QALYs, dollars, or other metrics, represents the weight by which the different health outcomes are multiplied. It is a reflection of how a particular group values a particular outcome, and is the means by which the subjective value system of the group can be incorporated into the analysis.

There are several different approaches to decision analysis, but one of the most widely used ones is a decision tree model. In this method, the structural elements (i.e., decisions, chance events, and final outcomes) are organized into a tree that clearly specifies the sequence of events.

The next two subsections of this document provide a description of the process of creating a decision tree for use in a cost-effectiveness study of a public health intervention. Section 4.6.3.1 describes the process using the simple hypothetical intervention, demonstrating how probabilities are assigned to the different events and how the expected utility for each option is calculated. This section may be skipped by those who are familiar with decision analysis. Section 4.6.3.2 then summarizes the benefits of using the decision tree, along with other associated tools from decision analysis, to structure a comparative risk problem.

4.6.3.1. Example Trees for Public Health Decisions — The objective of constructing a decision tree for any decision analysis is to structure the analysis and create a graphic representation of how all the possible options and choices associated with a particular problem relate to the possible outcomes. The trees are constructed from left to right, using a standard convention. The options to be compared constitute the main 'branches' of the tree. These arise from a 'decision node', which is depicted graphically as a square and indicates a choice under the control of the decision maker (Petitti, 1994; Haddix et al., 1996). The example in Figure 4-4 depicts a basic decision tree with three options for controlling an unspecified disease. The choices include: current practice, Alternative A directed at a high risk sub-population, and Alternative B directed at the entire population.

The final outcomes along each branch of the tree will be determined by a series of chance events. These events are not under the control of the decision maker and are depicted in the tree by circles referred to as 'chance nodes'. The two chance nodes depicted in Figure 4-4 (actual probabilities not shown) represent the likelihood that individuals belong to different portions of the population, the likelihood of developing the disease, and the likelihood of suffering adverse

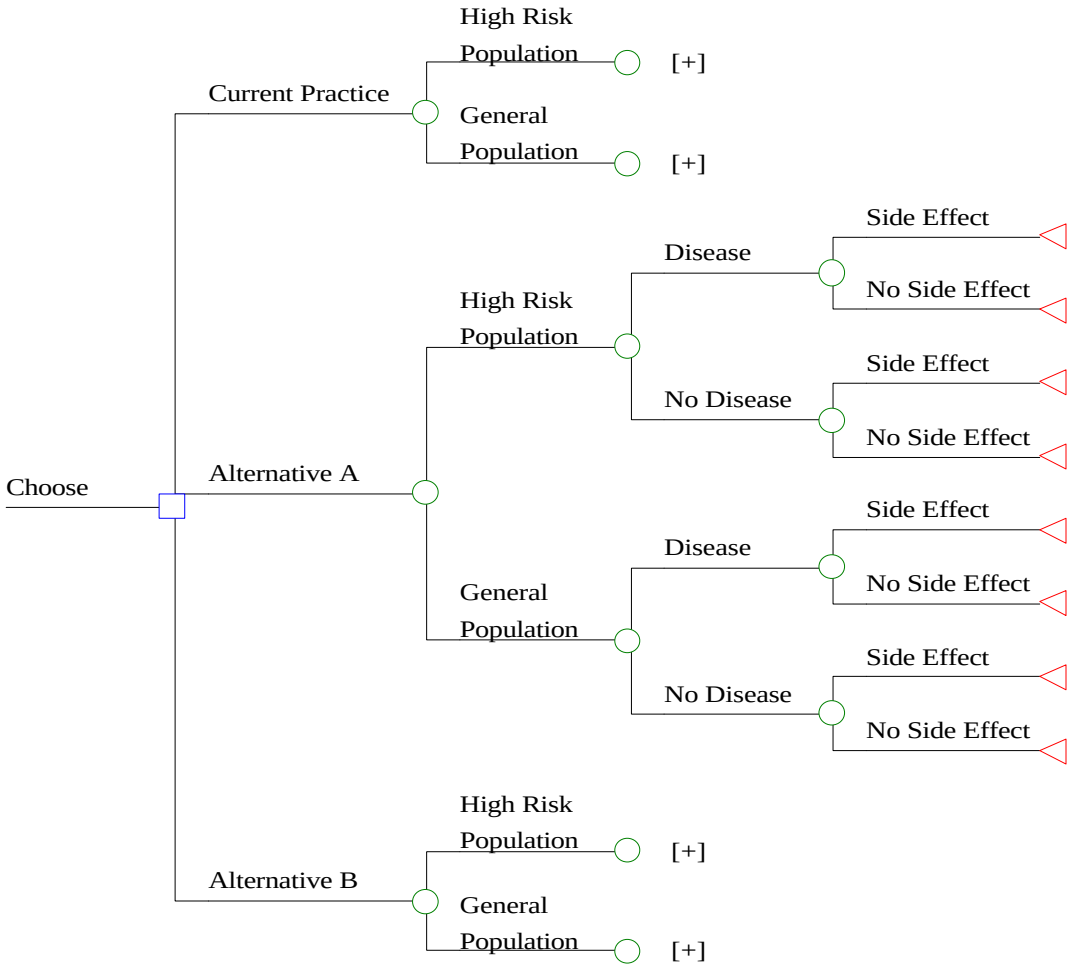

FIGURE 4-4

Decision Tree for Sample Public Health Intervention

effects as a result of the intervention. All possible outcomes at each chance node must be listed and clearly delineated so that a numerical probability can be attached to each event. Note that the events listed at each node must be both exhaustive and mutually exclusive, i.e. the probabilities must sum to 1.00 (Petitti, 1994; Haddix et al., 1996).

The probabilities used in decision analyses are usually derived from the literature, not generated *de novo* for each analysis. When insufficient published data are available, expert opinion is frequently sought to develop the best estimates (Petitti, 1994). For analysis of many public health interventions, the information for these indirect links may come from sources such as epidemiological studies, surveillance data, intervention studies, or clinical trials. For interventions addressing environmental health risks, engineering studies and data used in related risk assessments would also be essential.

For each of the three alternatives, the tree follows two distinct segments of the population separately, based on their risk status. The probability assigned to each of the branches would depend on the proportion of the target population defined as 'high risk'. Depending on the alternative (current practice, A, or B), some portion of each risk group are expected to develop the disease. The probability of developing the condition would presumably be greater in the high risk populations for all alternatives. The likelihood of disease presumably would also be reduced among the target populations for the intervention (i.e., high risk individuals in Alternative A and all individuals in Alternative B).

The decision tree also shows a final set of chance nodes depicting the likelihood of developing an adverse effect associated with the hypothetical intervention. The likelihood of suffering this adverse effect may be the same for all groups receiving the intervention.

For each of the decision options considered, the tree follows the entire population to four possible outcomes: 1) disease and adverse effect, 2) disease no adverse effect, 3) no disease and adverse effect, and 4) no disease and no adverse effect. The product of all the probabilities associated with any outcome, when summed within a decision option, represents the chance that this particular outcome will occur if that option is selected.

To compare the outcomes associated with each decision option, one needs to assign weights that reflect the value of each outcome. This process of assigning values to the outcomes permits each of these disparate outcomes to be compared on a common scale. This common metric may be dollars or a health utility measure such as QALYs. Other methods are available for developing more generalized utility measures. These are used to calculate the expected value of each of the decision alternatives.

The current case study presents a more realistic and detailed decision analysis of the water treatment problem. The decision tree includes 1) a baseline treatment alternative, 2) adding ozone to a standard treatment option that includes post-chlorination, a realistic decision that may be faced by many water purveyors and communities, and 3) providing in-home filters to immune compromised individuals within the population. The analysis considers illness and death resulting from microbial infection, potential reproductive risks, developmental risks, and cancer risks posed by DBP exposure. The decision option probabilities have been based on available published data or on the opinions of experts.

4.6.3.2. Uses of Decision Trees — Decision trees provides a logical and transparent framework for structuring the problem and summarizing the available data. The process of developing a decision tree provides a focus to the decision making process and lays out a path

that clearly describes the possible health outcomes and relative utilities associated with decision options. The process can also help to identify key parameters or probabilities for which there are not sufficient data. Research can then be directed to better define the uncertain variables which are most influential. The decision tree also provides a graphical representation of the alternatives under consideration and their consequences, which can help clarify the nature of the decision to different audiences.

The final outcomes or numerical summary measures generated by arithmetically combining the probabilities and outcome weights in a decision tree provide valuable assistance in making complex decisions. However, the utility of decision analytic tools to formulate, structure, and describe a problem often outweigh the utility of the calculated expected outcome. In some cases a decision tree is developed in the early stages of an analysis to define and characterize the problem, while actual calculations of expected outcomes are carried with spreadsheet or other software. In particular, when the analysis addresses parameter uncertainty or variability other computational methods are used, such as Monte Carlo analysis.

4.7. ASSESSING THE FINANCIAL COST OF ALTERNATIVES

4.7.1. Identifying Costs Components. An essential component of any cost-effectiveness study is an estimation of the financial costs associated with implementing the disease prevention strategy under study. These costs are generally considered to fall into three categories: program costs, costs to participants in the program (where appropriate), and costs associated with the side effects from the intervention. In addition to the costs of implementing a specific intervention, the analysis can also consider the costs associated with the health outcomes resulting from each alternative.

That is, to the extent that an intervention reduces the occurrence of a particular outcome, it may also reduce the medical costs and lost productivity associated with the condition.

The case study estimates costs from the perspective of the water provider, and includes only the implementation and maintenance costs associated with alternate drinking water treatment strategies. These are provided in the Chapter 5. The medical treatment costs and productivity losses associated with the different health conditions are not included.

Characteristics of a particular treatment system are likely to affect cost estimates developed for specific applications of this framework. The following issues should be considered.

- Water system size and population served by specific alternatives
- Current installed treatment facility characteristics
- Size of susceptible subpopulations

4.8. COMPARING RISKS AND COSTS

4.8.1. Summary Measures. The principle outcome measure in any cost-effectiveness study is the cost-effectiveness ratio, defined to be the incremental cost per QALY for the alternative interventions. The choice of which outcome measures to report is driven by the need to provide the identified audiences with a clear understanding of the expected consequences of each of the potential alternatives. For each alternative, total costs and the expected number of each type of health event should be reported (e.g., cases of cancer, cancer deaths, cases diarrheal illness). Total number of QALYs can also be reported for each alternative, based on the QALY weights of the outcomes. Providing both the number of events and their associated QALY costs, makes the analysis more transparent, understandable, and reproducible. In addition to reporting the total costs and health outcomes for each alternative, the analysis should report the incremental costs

and benefits of each alternative compared to the baseline. In some cases, it may be helpful to make additional comparisons between the other alternatives.

4.8.2. Time Preferences. Economic evaluations of public health interventions need to account for the fact that the associated costs and benefits often occur in different periods of time. Public health interventions generate benefits and costs that continue into the future, and individual decision makers and society generally weigh benefits and costs that they receive today more heavily than those they might receive in the future. Cost-effectiveness and cost-benefit analyses use discounting to convert all future health and economic costs and benefits to a current equivalent or present value, based on the selected 'discount rate'.

Discounting is intended to capture two types of preferences: the decision makers preference to defer a cost to a future period rather than the present, and their preference to receive a health benefit now rather than later. The first type of preference, which weights current expenditures more than future ones, reflects the fact that current investment in an intervention by private or public actors displaces some other expenditure. If the resources were not spent on the intervention, it could be placed in a safe investment with a positive real rate of return (such as government bonds), allowing the original amount (plus the return on the investment) to be spent in a future period. The Congressional Budget Office (CBO) identified 2% as an appropriate discount rate, based on the real (interest free) rate of return on long-term government bonds (Hartman, 1990). The Government Accounting Office recommends a rate of 2.7 to 3.0%, based on the return of government bonds for greater than 1 year (US OMB 1996b).

Just as future costs are discounted to a present value in a cost-effectiveness analyses, so are the health benefits associated with an intervention. Future health benefits are discounted

based on the social rate of time preference for consumption. To the extent they would prefer to receive a health improvement in this period, as opposed to a future period. Lind (1990) suggests that the appropriate social discount rate in these cases is 1% to 3%. Freeman (1993) recommends 2% to 3%. In some cases, the current costs of an intervention result in current consumption (rather than) investment being deferred to a future period. In those cases, the discount rate should be based on the social discount rate or time preference for consumption, as in the case of health benefits.

While these two rates (social rate of time preference and rate of return on investments) are theoretically distinct, in practice a single discount rate is usually used for discounting both future costs and benefits. In part this is done for practical reasons. Moore and Viscusi (1990) estimated that the financial rate of return on investment and the social rate of time preference for environmental health effects are both approximately 2%.

Gold et al. (1996) recommend a discount rate of 3% for discounting both costs and health benefits in cost-effectiveness studies of public health interventions. Haddix et al (1996) suggest that a rate of 3% to 5% would be appropriate. Both note that reporting results based on a range of discount rates is recommended, in order to insure comparability between studies.

4.8.3. Distributional Effects. Although most cost effectiveness analyses of public health interventions take a societal perspective and look at the aggregate impact on society, the costs and benefits may not be uniformly distributed within society. Applications of the CRFM must clearly describe the ways in which different groups are affected by the alternatives under consideration.

Distributional effects can take many forms. In some cases, the financial costs of an intervention may be borne by society as a whole, while the benefits accrue to an identifiable subset of the population. In others, Alternative A may benefit one group, while Alternative B primarily benefits another. In cases of comparative risks, an alternative may create a health benefit for one group while inducing a health risk for another. While aggregate measures provide a way to compare health outcomes among different populations, they cannot resolve the political and ethical issues involved in trading health risks between groups. Instead, the purpose of the analysis is to identify and quantify these tradeoffs for the decision makers and the affected groups.

In the case study presented here, individuals with immune-suppressive disorders, including AIDS, are more likely to contract waterborne protozoal disease, and are much more likely to die from it, than the general population (Hoxie, 1996). These relatively few individuals benefit the most from enhancements that increase the efficacy of existing drinking water treatment, although the costs are distributed throughout society. The relative impact of microbial and DBP contamination, and therefore the decisions relating to treatment technology choices, may differ depending on the proportion of immune-suppressed individuals within any particular population.

Other populations, such as children, the elderly, and heavy consumers of tap water, may also be more likely to suffer from microbial illness, and benefit more from additional treatment. In addition, the developmental and reproductive effects associated with DBP exposures do not equally affect all individuals.

4.8.4. Accounting for Uncertainty and Variability. For any analysis using the CRFM, the sources of uncertainty must be identified and all assumptions used in the analysis must be explicitly stated. These include assumptions regarding the efficacy of specific treatment options,

as well as uncertainty about the input parameters and model structure. This is particularly true in the case of environmental health interventions, where there is often substantial uncertainty regarding exposures and their potential effects. This section is not intended to provide a comprehensive review of the methods for uncertainty analysis. The reader is referred to the existing body of literature on this subject for a more complete discussion (US EPA, 1997; Burmaster and Anderson, 1994, Hoffman and Hammonds, 1994).

Each step in the comparative risk framework outlined here is associated with potential sources of uncertainty and variability, although the magnitude of the uncertainty and variability associated with the different types of risks may be very different. In a comparative risk setting these must be addressed explicitly and consistantly for all types of risks considered. Using protective assumptions to account for uncertainty or variability associated with one type of risk, may lead to the underestimation of the relative importance of another category of risk, and lead to a decision which reduces rather than improves public health.

The case study presented here demonstrates some of the sources of uncertainty and variability that may be shared by other comparative environmental health risk problems. Variability arises in part from parameters which take on different values for individuals within a given population. Potential sources of variability in this case study include differences in water consumption rates, and susceptibility to particular health effects. In addition, there is likely to be variability in microbial and chemical concentrations in water over time due to changes in temperature and other factors, although these differences may not translate into differences in long term risk among members of the population. If the case study were intended to represent a

group or class of treatment facilities, then additional inter-system variability would be introduced as a result of differences in source water microbial concentrations and treatment efficacy.

Other parameters in the case study are uncertain, meaning their precise value is not known. Among the uncertain parameters are the mean concentrations of microbial agents and DBPs, due to limited systematic sampling and limitations in the analytical methods for measuring them. The dose-response relationships for the known and unknown microbial agents and DBPs in the case study are also uncertain.

Along with these sources of variability and uncertainty in the estimation of health risks, a cost-effectiveness analysis introduces additional sources of uncertainty and variability. First, financial costs associated with different treatment options can be uncertain. The CEA also requires the conversion of health risk information to a QALY cost associated with these risks. For chemical risks, and to a lesser extent for microbial risks, limitations in the existing data result in uncertainty in estimating the nature, severity, and duration of the actual health conditions individuals may experience. In addition there is both uncertainty and variability in assigning QALY weights to the different conditions. Specifically, there are likely to be real differences in individuals' preferences for the different health states, as well as uncertainty in the actual QALY weights estimated.

Several methods are available to formally and systematically investigate the magnitude and impact of uncertainty and variability on a cost-effectiveness analysis. Sensitivity analysis involves recalculating the outcomes using reasonable, alternative input values, to quantify their impact on results, such as the cost-effectiveness ratio.

Probabilistic or Monte Carlo techniques provide a method for incorporating and quantifying the impact of multiple sources of uncertainty and variability simultaneously. In probabilistic analyses, uncertain and variable parameters are characterized with distributions rather than fixed point estimates. Values drawn from these input parameter distributions are used in repeated calculations of the expected outcomes, creating a distribution of the outcomes and a representation of the uncertainty. In addition to capturing the overall impact of variability or uncertainty, probabilistic methods can be combined with regression analysis to estimate the contribution of individual input parameters to the overall uncertainty in the expected outcomes. The current case study uses Monte Carlo techniques to quantify the overall impact of input parameter uncertainty on the predicted costs, health benefits, and cost-effectiveness ratio. Variability is characterized using separate simulations for specific populations of concern.

4.9. CONCLUSIONS

Combining risk assessment methods with cost-effectiveness analysis creates a comparative risk framework which can be used to assess the combined impact of water treatment alternatives on microbial and chemical risks. The stages of the CRFM (Figure 4-2) build on the assessment of separate risks as laid out in the 1983 NAS Risk Assessment Pardigm, translate the expected health outcomes into a common health metric (QALYs), and compare those benefits to the costs of alternatives. The method for conducting this type of analysis is outlined in this chapter and includes the following steps.

- Assess microbial and DBP mixtures risks separately
- Translate risk characterization into expected human health effects
- Express expected health effects in terms of a common health metric (e.g., QALYs)

- Assess the impact of alternative interventions on risks, individual health outcomes, and total health burden (QALY cost)
- Assess the initial and annual financial costs of alternatives
- Compare the financial costs and combined health benefits of alternatives

Subsequent chapters provide detailed examples for each of these steps and illustrate how they are combined into a cost-effectiveness analysis of drinking water treatment alternatives.