

DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES
CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

POLICY FIRST OR SCIENCE FIRST:
TWO APPROACHES IN THE MANAGEMENT OF RISK

Talbot Page

Forthcoming as "Economics and Risk Assessment,"
in The Impact on Marine Coastal Waters of the
Ocean Disposal of Municipal Wastewater and Its
Constituents, edited by Edward Myers. Chapter XI
Cambridge, Mass.: MIT Press.



SOCIAL SCIENCE WORKING PAPER 473

March 1983

ABSTRACT

This paper compares two approaches toward handling of uncertainty in public policy decisions. In one approach, "science first," scientific questions of fact are addressed first. Once it is concluded that there is a high probability that there is a problem (for example that a chemical is a carcinogen) then the analysis proceeds to value questions, where the costs and benefits of possible remedies are weighed. In the alternative approach, "policy first," value questions are addressed from the beginning. The potential costs and benefits of wrong decisions are considered from the start, along with judgments about the probabilities of false positives and false negatives. This analysis leads to the identification of the scientific uncertainties most critical for the public policy decision; it suggests the most important areas for research; and it guides decisions of what to do in the meantime, for precautionary actions.

The discussion is undertaken in terms of a particular public policy problem, management of municipal wastewater disposal in coastal waters.

uncertainties the alternatives are basically few and simple, and the weight of evidence leans strongly to one side.

POLICY FIRST OR SCIENCE FIRST:
TWO APPROACHES IN THE MANAGEMENT OF RISK

INTRODUCTION

In this chapter the focus shifts from science to policy. The preceding chapters pull together a large amount of scientific information related to wastewater and sludge disposal in coastal waters, but at the same time they suggest that what is known is but a small fraction of what needs to be known for confidence in making some of the most important policy decisions. We know little about the effects of toxic chemicals, and little about the importance for toxics of wastewater and sludge disposal in coastal waters relative to other transport mechanisms such as air pollution and groundwater contamination. There are large gaps in our knowledge of the more traditional materials as well — the natural organics, nutrients, and metals. We can measure changes in worm populations where sludge is discharged, but we know little of the long term implications of such disturbances.

The nature of uncertainty varies from one aspect of the management problem to another, from toxics to pathogens to the other materials. In some areas, such as toxics, uncertainty is very great and decision making resembles gambling. In some other areas, the appropriate alternatives are more clearcut, because in spite of the remaining

The purpose of this chapter is to consider how variations in uncertainties associated with coastal disposal of wastewater and sludge can be taken into account in the formation of policy and decision. The starting point for the chapter is the realization that because of existing uncertainty, the policy maker is condemned to make decisions under uncertainty. A decision to postpone some action until more is learned is a decision under uncertainty, just as much as a decision to take precautionary action in the meantime. A decision to promote further research, and what type of research, is still a decision under uncertainty.

Yet decisions are made "in the meantime." We have built ocean outfalls and treatment plants, and we have existing and emerging policies concerning the level of treatment, pretreatment, disposal and source control. How such policies can be understood and assessed and, if possible, improved upon is the subject of this and the next two chapters. This chapter attempts to build a framework for policy analysis. The next chapter considers the legal institutions which provide means for implementing policy options as well as constraining them. The following chapter offers an evaluation of specific policy choices.

UNCERTAINTY AND THE POLICY APPROACH

The best place to begin a policy analysis is with questions of resource commitment. How and where to commit resources are the "ultimate" policy questions for agencies and legislatures concerned with wastewater and sludge disposal into coastal waters. Besides this, questions of resource commitment are an important starting point for another reason.

While there is much to be said for accumulating knowledge for its own sake, from the point of view of policy formation — to put the matter in perhaps a too utilitarian fashion — the wealth of scientific information in the preceding chapters is useful only to the degree that it helps resolve questions of where and how to commit resources. To put the matter a little more positively, beginning with questions of resource commitment helps structure scientific information for decision purposes. By itself, scientific inquiry is likely to be an open-ended process, with investigation of each interesting question leading to three or four more interesting questions. To structure this information for decision purposes, we start with a list of policy questions which have to be decided at the current time, one way or another (by default if not by active choice). We begin with questions of resource commitment such as these, list A:

A1. How much regional variation should there be in the treatment of wastewater and sludge? Previous policy in the 1972 FWPCA was toward uniform treatment. Section 301(h) of the 1977 amendments allows for some regional variation, but how much should there be?

A2. How advanced should be the treatment? Generally speaking, the more advanced the treatment the more sludge. Where do we stop?

A3. What should be the balance between ocean disposal and disposal into other media? How much wastewater should be treated and used for groundwater recharge and irrigation, if any? How much sludge to landfills and incinerators? Present law, in the Ocean Dumping Act, bans ocean discharge of sludge after 1981.

A4. What should be the balance between treatment at the sewage plant and pre-treatment by industrial firms before disposal into the municipal sewage system? This question applies to metals and toxic chemicals, and differently in each case.

A5. How much source control should there be? This question applies especially to toxic chemicals. The only alternatives to ultimate disposal somewhere in the environment are chemical transformations to harmless molecules or restriction of production in the first place. Some chemicals may be sufficiently hazardous that they are best controlled at their source.

A6. What research and monitoring should be sponsored by government agencies?

While these questions will help organize the subsequent analysis, it is not possible to answer them directly. They depend to a large measure on judgements as to the seriousness of the effects of wastewater and sludge disposal. The principal questions of effects can be briefly enumerated in list B:

B1. How bad is the toxics problem? Obviously the response to this and other effects questions will vary with the location and particular situation.

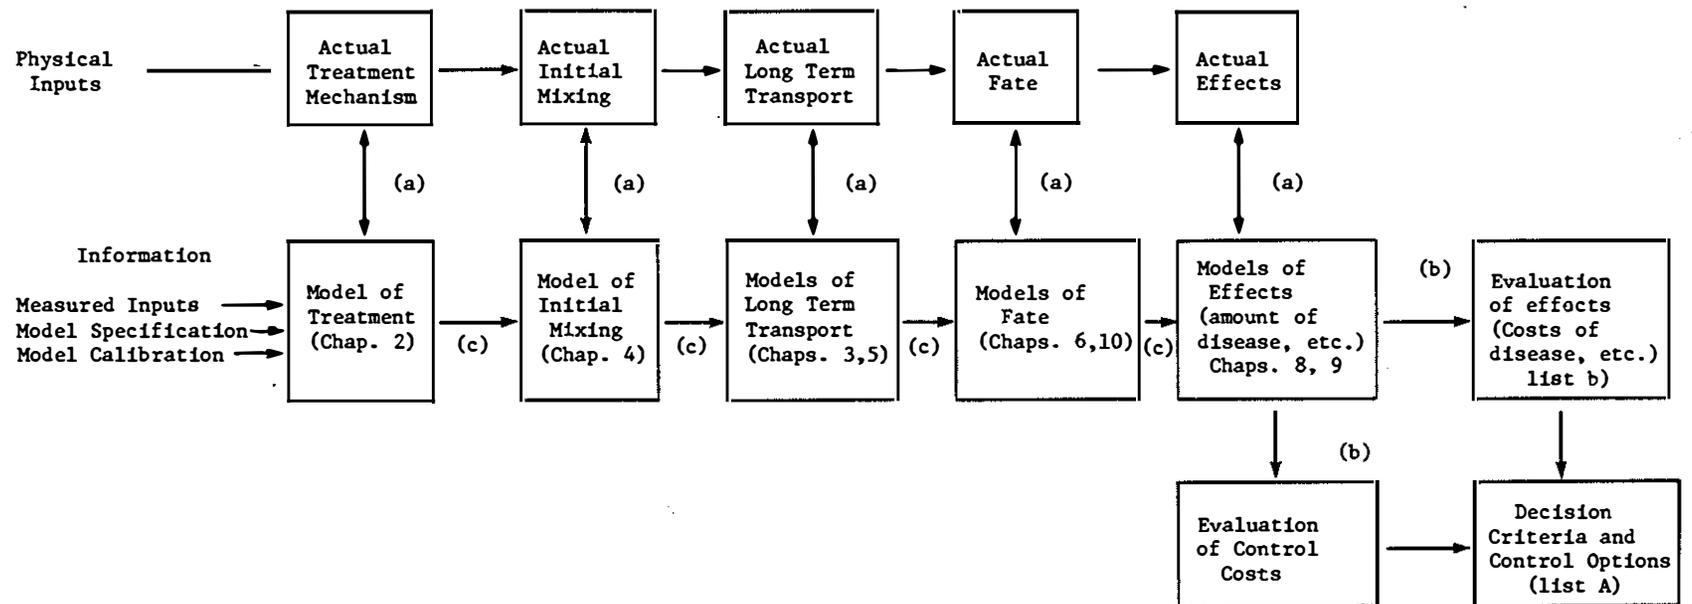
B2. How bad is the problem with pathogens and communicable disease?

B3. How much damage is there to recreation and use of coastal water; how much aesthetic effect is there?

B4. How severe is the stress to marine populations?

These questions in turn depend upon interpreting the existing scien-

Figure 1



tific information.

The policy analyst and the scientist tend to work in opposite directions. To organize and focus his inquiry, the policy analyst begins with the ultimate, practical questions of decision (list A) and works back to effects (list B) and then to what is known about marine transport, chemical and biological systems. The scientist tends to work forward, beginning with the physical inputs to the coastal waters, to model the causal flows as they move forward in time.¹ In this chapter the treatment of uncertainty (and correspondingly the value of information) is derived from the policy analyst's approach. Both approaches are needed; they contribute to each other as sketched in Figure 1.

In terms of Figure 1 the science approach, followed in much of the book, begins in the upper left corner and works down toward the lower right corner. Actual mechanisms cannot be understood directly, but scientists try to make their models correspond to the actual mechanisms as closely as possible — knowledge for its own sake. Much attention is devoted to tightening the correspondences (a) and developing detailed submodels.

Cost-benefit analysis focuses on the evaluation of effects and the costs of control (b). There are few cost-benefit studies of wastewater disposal, and most of the attention appears to be on the quantification of recreational benefits from waste treatment (Freeman, 1979a) and the engineering costs of conventional sewage treatment

(Mueller, Chapter 2 of this volume). A narrowly focused cost-benefit analysis accepts the predicted effects as more or less given and concentrates on the problems of evaluation (Freeman, 1979b).

Although concerned with the evaluation of decisions, this chapter does not attempt a narrowly focused cost-benefit analysis. A broader approach for analysis of wastewater and sludge disposal is adopted because of the dominating role of uncertainty. How we "allocate" or treat existing uncertainties and the steps taken toward their resolution have major, sometimes determining, effects on actual decisions. For example, DDT in wastewater discharges into the Southern California Bight were not regulated until large volumes of DDT were discovered in the wastewater (about 1970). Once discovered, regulation by source controls did not wait for a quantified cost-benefit study of the damage to marine ecosystems from this discharge. In a similar way, whether or not a chemical is banned depends more on the strength of evidence of its carcinogenicity and potency than on a fine-tuned argument as to the value of life. The basic question to be addressed is how evaluations of the costs and benefits of effects and possible control measures are to be brought together with the uncertainties that pervade the prediction of effects and understanding of marine mechanisms.

Given the complexity of marine mechanisms and the enormous number of uncertainties involved, it is fair to ask at the outset what might be gained from a policy approach that attempts to work back from potential decisions in the lower right corner of Figure 1 toward the

phenomena of the upper left. Three types of benefit can be distinguished.

First in a few cases the analysis might be cut and dried — effects and control possibilities are well enough understood so that the expected net benefits of alternative decisions can be estimated quantitatively, and a single course of action selected as offering the highest net benefits.

Second, in other cases the uncertainties are sufficiently pervasive so that explicit calculation of expected net benefits from alternative possible decisions is not very helpful. Nonetheless, "islands of relative certainty" rise above the general sea of uncertainty. These islands are not built upon underwater mountains of evidence, like normal islands, but on broad understanding of what appear to be the largest and most important effects, in their qualitative evaluation of costs and benefits. In such cases, the existing evidence is not conclusive in a scientific sense, but has been accumulating in favor of the relative certainty, and it appears that further resolution of uncertainty will strengthen the evidence. For example, concern with toxic chemical discharge into ocean water existed a decade ago (Brooks, 1971). Research in the intervening period has strengthened our concern, and it appears likely that future research will further strengthen it. Several islands of relative certainty with their policy implications are tentatively identified in this chapter and discussed in more detail in Chapter 13.

Third, and perhaps most important, the policy approach, even when not yielding specific decisions, may offer some guidance as to the process of policy formation.

In other words, how one thinks about the various uncertainties associated with effects has a large impact on the control measures adopted. What is needed is a framework flexible enough to encompass the range of uncertainties associated with wastewater and sludge discharge. This range can be described in terms of the materials discharged, divided into four classes.

A Spectrum of Uncertainty

Class I includes nutrients and natural organic material in the form of suspended solids, ammonia and other natural oxygen-demanding materials. These materials naturally cycle through ecosystems in large quantities. Class II includes pathogens: bacteria and viruses. Class III includes heavy metals like lead and cadmium in concentrations far greater than natural systems. Class IV includes toxic chemicals which, by affecting the genetic code (genotoxic), may cause cancer, teratogenic effects, and other diseases. A large number of synthetic organics and radioactive materials fall into this class.

As a crude generalization, uncertainty increases from Class I to IV. We know most about the mechanisms and effects of the natural materials in Class I. For example, models can predict with some confidence the effects of discharge of BOD (biological oxygen demand) on

dissolved oxygen concentrations. The diffusion and die-off rates for pathogens have also been well studied, as the control of pathogens has been a primary concern in sewage treatment. However, there are important uncertainties still remaining in our understanding of pathogens, as noted in Chapter 8. One such uncertainty is how well coliform, an indicator bacterium, matches the die-off rates for the pathogens of real concern, both in the treatment process and in the marine environment.

The presence of heavy metals in sewage became a concern when it was found that large slugs of them could temporarily upset the sewage treatment process. Also, until the 1970s it was thought that the addition of metals to the background levels in seawater was small and there was little or no bioaccumulation of them in the food chain. But it is now known that waste discharges of metals are of a higher fraction of the background levels than previously thought (Morel and Schiff, Chapter 6 of this volume). Moreover, we now know that some metals can be biologically activated, for example methyl mercury. Once activated, some become toxic chemicals. Some bioconcentrate up the food chain, like a number of other toxic chemicals. Methylation of mercury came as a surprise, and there are important remaining uncertainties as to which metals, under what conditions, can become bioactivated.

But probably the largest uncertainties are associated with the enormous number of industrially produced chemicals that find their way into wastewater and sludge. We know little about the long term toxicity of all but a relative handful of the 75,000 chemicals in active

commerce, many of which are new chemicals, synthesized after World War II.² We do not know how many of these chemicals, in what quantities, find their way to municipal wastewater systems and thence to coastal waters. We do not know the relative importance of wastewater systems as a conduit compared with other forms of transport and dispersion. Of the relatively small number of chemicals that have been tested adequately (perhaps several hundred for cancer and fewer for other long term effects such as teratogenicity), we have only rudimentary knowledge of toxic potency, with estimates often ranging over several orders of magnitude. For these largely unknown and uncharacterized chemicals the key questions for coastal disposal are: (1) how much is being discharged? (2) how toxic? (3) how persistent? (4) how likely to bioaccumulate, ultimately as a pathway to human exposure? and (5) what effect on marine populations?

Several factors help explain this generalization of increasing uncertainty from Class I to IV, as crude as it is:

Numbers. In Class I there are only a few aggregate indexes of concern, such as BOD and SS; in Class II a relatively small number of important diseases of concern; in Class III, though a small number of metals, a larger and undetermined number of reaction products; in Class IV an enormously larger number of candidate chemicals to begin with, often at very low concentrations (ppb and below).

Pathways. Knowledge about pathways decreases as we move from Class I to Class IV. Class I materials move along well known natural

pathways, while Class IV pollutants have long half lives and pathways of many complex steps.

Latency. Some of the effects of Class I materials are almost immediate; some of the effects of pathogens occur within a few days of discharge; some of the effects of bioactivated metals may take several years to manifest themselves; the usual latency for carcinogens is 20 to 40 years from the time of exposure (not including the time of bioaccumulation); and for mutagens the latency may be a generation or more.

Visibility. Somewhat associated with the other factors, the effects of Class I materials are typically more readily apparent than those of Class IV.

Zero-Infinity Dilemmas

The costs and benefits associated with a control alternative also vary over a wide range. At one end of the range the underlying effects may be reasonably well understood and the costs and benefits of controlling the effect quantified with some confidence, as with smells and other simple aesthetics. In such cases uncertainty plays a minor role, and the control activity is undertaken to the point where the incremental costs are equal to the incremental benefits. Such cases can be called cases of classical pollution, and are most amenable to classical cost-benefit analysis, where uncertainty plays a minor role.³

At the other end of the range (Class IV) an effect is not well understood, nor the underlying mechanisms. Nonetheless, it is useful

to compare the potential environmental cost with its precautionary cost of control. As the term has been used in the area of nuclear safety, a "zero-infinity dilemma" is one where there is some small assessed probability of a catastrophic cost. Thirty years ago methylization of mercury appeared highly unlikely; yet the cost of control of mercury discharge (basically source control) was small compared with the damage done by the bioactivated chemical. Those observations lead to a second, crude generalization — control problems for Class I pollutants often tend to have a classical pollution character; the pollutants in Class IV a zero-infinity character. For one thing, Class I pollutants often affect recreational and aesthetic values, while Class IV pollutants are associated with dreaded disease (but so are some of the pathogens of Class II). For another, many of the effects associated with Class I pollutants are readily reversible, whereas the effects of Class III and IV are more irreversible.⁴

Several characteristics taken together, which define a notion of environmental risk,⁵ apply to Class IV pollutants. The relatively large amount of ignorance, which leads to pervasive uncertainties all along the modeling process, has already been mentioned. So have relative irreversibility, which not only makes mistakes in the direction of insufficient precaution costly, but also long in time to correct; latency; and the "zero-infinity" characteristic. These characteristics tend to make pollutants of environmental risk difficult to manage, in establishing institutions and incentives leading to the "proper" amount of control.

Two other characteristics add to the difficulty. These characteristics tend to apply to all pollutants, not just ones of environmental risk, but they are important characteristics of the latter as well. First, the costs are external. The costs of Class IV pollutants are widely diffused in time and location; the costs are borne not by the generators of the costs but by the public at large. Second, benefits are internal. The benefits accrue to the private firms generating the risks, in the form of lower operating costs. Sewerage agencies who do not take control measures also benefit in lower operating costs. How these characteristics of internal benefit and external cost complicate the control process is illustrated by a comparison of Class IV pollutants with Class III pollutants. The characteristic of external costs applies less strongly to Class III pollutants. The basic reason why metals came into control, by pretreatment, was that some of the costs were borne internally by the treatment plant, rather than externally by the public at large. Thus it was to the advantage of the operators of treatment plants, who were accountable for the internal costs of the plants, to control metals. Metals were affecting plant operations and raising treatment costs. A similar incentive for control does not exist for Class IV pollutants, which do not directly increase the costs of operating a treatment plant. In contrast with the classical pollutants of Class I, pollutants of Class IV tend to have more of the environmental risk characteristic.

MAXIMIZING NET BENEFITS

The goal in this section is to apply a decision analysis which is flexible enough to encompass the spectrum of uncertainty and the range of balance to extreme imbalance of costs and benefits. The analysis is driven by a maximization of expected net benefits, tailored to the situation where information is incomplete but can be improved through research or testing (Weiner et al., 1979; Pauker and Kassirer, 1980; Weinstein, 1979; Page, in press). Examples of decisions which fit into this framework include the following. The decision to extend the Orange County outfall from its original design length; the decision (as yet unmade) whether or not to increase the capacity of the treatment system to better accommodate storm water overflows in San Francisco; and the decision to ban dieldrin (an example of control by source reduction).

Hypotheses and the Costs of Being Wrong

We begin with a hypothesis about the existence or non-existence of some adverse environmental effect. One such hypothesis is that some particular chemical is a human carcinogen, which when released at a certain rate to seawater at some location will bioconcentrate to 10 ppm in certain fish. The evidence may be fragmentary and weak for this hypothesis and the likelihood of it may seem low. Another hypothesis is that with an extension of an outfall to eight kilometers into deep water in a particular region the shoreline coliform count will exceed 10/ml less than 5 percent of the time. So much may be known about the short term diffusion around the outfall, and the other mechanics of the

problem, that this hypothesis might be held true (or false) with a much higher degree of confidence.

Generally, hypotheses can be stated in one way and then restated in the opposite way. For the sake of definiteness, subsequent hypotheses will be stated about the existence of some adverse environmental effect. By setting up our hypotheses as ones hypothesizing the existence of adverse environmental effects, we are able to interpret consistently the cost of a false positive as equal in magnitude to the resource cost of precautionary control. Note that the cost of a false positive is borne only when we act to mitigate a potential environmental harm, which happens to be nonexistent. The cost of a false positive is the cost of one type of wrong decision. Because we are dealing in relative costs we can count the cost of precautionary control as one unit.

The cost of a false negative arises when we reject the hypothesis of an adverse environmental effect, but the effect exists. We can count the full environmental cost of this harm as $D+1$ units. If the correct decision had been made the $D+1$ cost of environmental harm could have been avoided, but at the resources cost of one unit, for the precautionary action. Thus the net saving by undertaking the precautionary action is the difference, or D units. Thus we count D as the cost of a false negative. For Class III and IV pollutants the cost of a false negative may be several fold the cost of a false positive; in other words D may be "much" higher than one. Alternatively, D is negative when the cost of control is higher than the environmental harm

prevented. The cost of a false negative is the cost of the second type of wrong decision.

It may seem more comprehensive and balanced to consider four cases: the cost of precautionary action when the environmental hypothesis is true (one unit); the cost of no action when the environmental hypothesis is false (zero units); the cost of not taking precautionary action when the environmental hypothesis is true ($D+1$ units); and the cost of taking precautionary action when the environmental hypothesis is not true (one unit). This explicitly takes into account the direct costs of right decisions as well as the costs of wrong decisions. But in this chapter we are measuring the costs of wrong decisions in the opportunity sense, as the cost of the wrong decision relative to the cost that could have been avoided if the right decision had been taken. Thus the cost of a false positive is (one minus zero) or one, and the cost of a false negative is ($D+1$ minus one) or D . In this way the costs associated with right decisions are implicitly taken into account, and the analysis made more compact. Whether all the cases are analyzed explicitly or only the wrong decisions are analyzed in an opportunity cost sense, the strategies of minimum expected cost are the same and the approaches equivalent.

Likelihood of Hypotheses, Existing Information

After consideration of all the existing information, not all hypotheses appear equally likely. For an analysis of expected net benefits, some evaluation of likelihoods is needed. And here lies one

of the principal differences in the approaches sketched in Figure 1. Traditionally many scientists prefer not to evaluate the relative likelihoods of hypotheses in terms of quantitative, subjective probability assignments. The usual reason given for this reluctance is that too little is known about the truth or falsehood of the hypotheses to provide the grounds for numerical statements of likelihood. If pressed for quantification of some effect, instead of approaching the matter by developing subjective probability assignments for the underlying hypotheses, a different approach is taken. "Most likely" or "conservative" hypotheses are adopted; point estimates are calculated as though the hypotheses were true; and then a "safety factor" or "margin of error" is applied at the end. Decision theorists counterargue that this approach precludes the examination of minor hypotheses, which may have important implications for precautionary control (it makes the decision tree too narrow); provides no basis for determining how large should be the factor of safety; and concentrates on the propagation of measurement error while neglecting error in the specification of models where the most critical uncertainties often lie (National Academy of Sciences, 1980; Leamer, 1978).

For our purposes the matter is not either/or. Full quantification of subjective probability assignments of the likelihoods of hypotheses is not necessary, any more than is the full quantification of costs and benefits of the various effects and control alternatives. The analysis proceeds as though such quantifications could be made, but then is interpreted, as much as possible, qualitatively, in terms of

more or less likely, more or less costly. In fact, there is no point in full quantification of assessed likelihoods of various hypotheses, when it is not possible to quantify fully the costs and benefits of the alternative effects. We note in passing, though, that there is a growing literature on why some sort of evaluation of the relative likelihoods of hypotheses is necessary, how such evaluations can be done, and how accurate and consistent they are (Raiffa, 1970; Fischhoff, in press; Morgan, 1979; Savage, 1971; Hogarth, 1975).

Here for the exposition we will focus on just one hypothesis at a time and write the assessed odds against it as $N:1$. For an environmental risk, N is definitionally greater than 1, the hypothesis of harm being considered less likely than the benign alternative, perhaps many fold less likely (N several fold bigger than one). For a classic pollutant, N is typically less than one, with the hypothesis more likely than not. Similarly, for a classical pollutant the costs of wrong decisions, either way, are often roughly comparable (especially for the marginal decision), and even when the adverse hypothesis is true the net saving from a precautionary decision, D , is "near" zero, perhaps well between one and minus one. But for an environmental risk, the net saving from a precautionary decision, D , is large compared with its cost, still counted as one unit. In other words, by definition, for classic pollutants both N and D are small, for ones of environmental risk both N and D are large.

New Information, Probability of Error

In their general setting, classical and environmental risk pollutants both require evaluating existing information, defining research issues, and "updating" the existing information from research and statistically interpretable tests of the data. Although for a given decision, with its researchable questions, there are likely to be many research issues and many tests of them, we can think of the many tests as one composite test, which when undertaken will either lead us to accept or reject the hypothesis in question.

With pervasive uncertainties the new information is unable to resolve the question with certainty. For a given empirical test, the test may erroneously suggest that the hypothesis is true when it is not. This outcome is a false positive.⁶ The probability of a false positive is known as the significance level of the test, and is customarily denoted α . For well understood tests this probability is easily quantifiable, and customarily tests are designed so that α is 5 percent, an "acceptably low" probability of a false positive. For new information not gained through a statistically designed test there is still some probability that the information will indicate the hypothesis is true when it is not, but this probability is not easily quantifiable (unless the new information is so overwhelming we might assign a zero probability of a false positive).

As for the other type of error, the test may erroneously suggest that the hypothesis is false when it is true. This outcome is known as a false negative. The probability of a false negative is customarily denoted β . For well understood tests this probability is also

directly quantifiable, although its calculation depends upon the magnitude of the effect (for the calculation of α the magnitude of the effect is assumed zero) and thus is sometimes more difficult. Again, for information gained outside the structure of a formal statistical test, there will still be some probability of new information suggesting that the adverse environmental effect is absent when it exists (a false negative), but this probability is not easily quantifiable.

At this point we can state another important difference between the two approaches sketched in Figure 1. In the science-first approach it is customary to first set α to some predefined level, often 5 percent. Then for a given research and testing budget a statistical test is chosen to minimize β subject to α equal to the preset level. Notice in the selection of the test and α and β that evaluation of the relative costs and benefits plays no role. In the policy approach, the relative costs and benefits play a role as important as α and β . Maximizing net benefits is equivalent to minimizing the expected costs of wrong decisions, which is done by minimizing the sum of the probabilities of false positives and false negatives, weighted by the costs of false positives and false negatives. This equivalence is noted and used in the section The Maximization, below.

Together α and β characterize the information content of a statistical test or proposed new research. It would be reassuring if for a given decision we could set up the relevant hypothesis and empirical test of it so that both α and β were small, both equal to or less than 5 percent. Although this appears possible with some classical

pollutants, for which information is more complete, it appears unrealistic for typical environmental risk pollutants, for which information is highly incomplete. As we shall see more clearly later, when numerical calculations are presented, the difference between the science-first and the policy-first approaches to uncertainty in new information matters less for classical pollutants than for pollutants of environmental risk.

Cost of New Information

Finally, to complete the framework of expected cost, we need to compare the cost of acquiring new information with its value. The value is measured by its effect in decreasing the risk of too much precautionary behavior, and the risk of too little. For a particular decision the cost of new information from additional research and testing is counted as T units, where a unit is the cost of a false positive.

The Maximization

Once the decision under consideration is specified, the relevant hypothesis stated, the existing information incorporated into an assessment of the odds or likelihood of the hypothesis, and new research specified which would help resolve the uncertainty, there are basically three decisions possible.

First, one concludes that the research would not provide information sufficient to justify its costs and the decision maker acts on the basis of accepting the hypothesis as true (that the harmful

environmental effect exists). The higher the cost of a false negative the more appropriate this decision.

Second, one concludes that the research would not provide information sufficiently valuable to justify its cost, and the decision maker acts on the basis of rejecting the hypothesis. The higher the cost of a false positive (unnecessary precautionary action) relative to the cost of a false negative, the more appropriate this decision.

Third, the additional research is undertaken, the hypothesis accepted or rejected, and precautionary action taken or not taken on the basis of the new information combined with the existing information.

It is possible for the first decision to be sometimes appropriate even in the extreme case when the resource cost of gathering additional information is zero. Such situations can arise when the cost of a false negative is high. Even if the resource cost of gathering new information is zero, there is a risk of a false negative as long as β is greater than zero. By deciding to act precautionarily the potential cost of a false negative is avoided at the cost of a false positive, which occurs whenever the hypothesis is false. Since the assessed probability of the hypothesis being false is $N/(N+1)$ and the cost of a false positive is valued at one unit, the expected cost of this decision is just⁷

$$N/(N+1). \quad (1)$$

At the other extreme the cost of precautionary control is high

in comparison with the cost of a false negative. Gathering new information and relying on it entails the risk of a false positive, which can be avoided altogether if the new information is not sought and the hypothesis is rejected on the basis of existing information. In this case the cost of a false negative is borne whenever the hypothesis is in fact true. With the probability of the latter event assessed at $1/(N+1)$ and the cost of a false negative valued at D , the expected cost is

$$D/(N+1). \quad (2)$$

In between the extremes we gather the new information and condition action on the basis of the combined new and old information. We also bear the resource cost of the test, which we can count in units T , relative to the cost of control (the cost of a false positive). The expected cost of this decision is

$$D\beta/(N+1) + N\alpha/(N+1) + T. \quad (3)$$

In principle the approach of cost minimization directs us to select the smallest of (1), (2) and (3). In selecting (3) it is possible to consider tests of different cost (T), and with various probabilities of false positives and false negatives. For a given cost of the test, α and β "trade off"; that is, if we insist on a lower α , then we must live with a higher β , and vice versa. It is possible to bring both down simultaneously, by increasing the cost of the test or research project, but there are limits to the gain in test information even with very expensive tests.

In practice, it is hard to quantify all these variables and

trade-offs. Nonetheless some quantifications are possible, as we illustrate below. Moreover, the framework of expected cost helps guide qualitative inferences which will also be discussed below. The geometry of this framework is outlined in the Appendix.

Levels of Uncertainty

A basic purpose in setting up this expected net benefit is to provide for essential roles for uncertainty in the analysis. Uncertainty can be analyzed at several levels: uncertainty concerning existing information, as to the likelihood of the hypothesis, or uncertainty as to the proper assessment of N ; uncertainty as to the completeness of new information, or uncertainty as to α and β and the power of new research to resolve the existing uncertainties; and uncertainty as to the magnitudes of the potential harms and the costs of precautionary controls. For some very well understood effects in Class I pollutants the first two types of uncertainty may not be important, and may drop out of the analysis. To model the non-probabilistic case, set $N=0$. In this case there is no need to acquire new information about environmental effects and the only question is whether $D>0$ or $D<0$. If $D>0$, the cost of environmental harm is greater than the resource cost of control, and the preventive action is worth the cost. If $D<0$, the cost of prevention is more than the cost of the environmental harm, and the control is not worth its cost. Thus in this limiting case attention shifts from the probabilistic and statistical aspects of decision making to the deterministic ones estimating the engineering costs of con-

trol and measuring the environmental harm (quality of recreation diminished and so on). Much of standard cost-benefit analysis is focused in this direction.

ILLUSTRATIONS

A few examples may help illustrate some of the foregoing ideas. The first example is rather simple and its common sense solution does not require the formalization of expected value maximization of the last section. It serves more to illustrate some of the preceding definitions. In the second example the formalization of expected benefit maximization contributes to a clearer understanding of acceptable risk and its use in setting standards for pollution control. The third example suggests the likely asymmetries for Class IV pollutants.

1. Design of Treatment Facility

In the design of an ocean outfall, one of the basic engineering decisions is its length from shore. The analysis involves a trade-off between marginal environmental benefits of greater distance from shore (and greater depth) and the increased cost of construction and operation. If the chosen length is fairly long, so that shoreline impacts (like coliform counts) are already predicted to be low, the incremental benefits from a small increase in length may be small and hard to measure. But the incremental costs may also be small, because the pipe laying equipment is already at the working site, the difficult surf-zone section has to be completed anyway, and the incremental length is

the easiest part to build. Thus installing more pipe may cost only one half the average unit cost. In contrast, remobilizing later to extend the pipe length (if it turns out to be necessary) may cost three times the average cost now or six times the present incremental (marginal) cost because of the need to mobilize the equipment again and rearrange the diffuser section.

In this case the designer may decide as a precaution to extend the outfall a little further because it is relatively cheap insurance in case the water quality does not come out as favorably as predicted. In other words, the cost of a false positive equals the incremental cost of extending the outfall. If it turns out that the extra length is needed to protect the water quality, this is a true positive. (In this model water quality is "adequately" protected by minimizing the sum of environmental costs and management costs.) If it turns out that the extra length is not needed, this is a false positive. The cost of a false positive represents the money that could have been saved had there been perfect information from the beginning.

From the point of view of the sewerage agency concerned only with internal costs and meeting mandated standards, the cost of a false negative can be defined directly. If the outfall were not extended and has to be extended later, this would be a false negative. The cost of a false negative is the amount that could have been saved had there been perfect information and the outfall been built right, at the beginning. The cost of a false negative is five units, counting as one unit the incremental cost of extension at the original time of

construction and six units as the incremental cost of extension later.
(Net saving is $D = 6-1 = 5$.)⁸

However, the agency, whether it is local, state or federal, may be interested in more than minimizing the expected construction and operating costs, subject to meeting existing water quality standards. It may be concerned with how well the standards reflect the true potential environmental damage from waste discharge. And in principle, the cost of a false negative is defined in terms of the potential environmental damage from insufficient precautionary action, rather than the cost of remedy if remedy is later required. Nonetheless, the cost of remedy compared with the cost of original construction is relevant. Estimating the potential environmental cost is likely to be much more difficult and uncertain than estimating the incremental costs of construction. And thus the ratio of incremental construction costs (five in this case), can be viewed as the cost of insurance against the largely unquantified environmental risks, where the higher the ratio the lower the cost of insurance.

In a case like this, the decision to extend the outfall clearly depends on how much confidence the engineer has in the water quality analysis. Since it is not a precise art, it is attractive to run the risk of overbuilding the outfall when D (defined in terms of internal costs) is as high as five. On the West Coast (see Table ?? in Chapter 4) outfalls have been made very long. For example, the Orange County Outfall extends 27,400 ft (8352 m) off the Orange County coast. The outfall was made probably 5,000 ft longer than necessary⁹

to meet the regulatory requirements, but it was felt that the assurance of high performance and other benefits made it an attractive expenditure. Other factors entered into the decision also, such as the tendency for the requirements to get tighter during the lifetime of the outfall, which may be as long as 75 years.

A second design example suggests another pattern of potential costs and benefits. The problem is what to do about combined sewer overflows, such as is faced in many large eastern cities and the City of San Francisco on the West Coast. Because of the large volume of street runoff going into the sewers during rainstorms, the normal dry-weather treatment plants are completely inadequate to handle the suddenly increased flow when it rains. Therefore, overflows of untreated wastewater mixed with street runoff may occur in coastal waters, but this occurs only a few percent of the time during the year (e.g., four percent of the time in San Francisco where the rainy period is confined to the winter months). Because of the transitory nature of these overflows, they do not appear to be very damaging ecologically, and are often regarded primarily as aesthetic nuisances. Coliform counts may exceed requirements for a day or two during and after the rain.

But in fundamental contrast with the previous example, the precautionary cost of control is roughly comparable to the cost of remedy, if control is postponed. That is, later control appears to be no more expensive in real dollars than proceeding at present with schemes of large storage vaults, pumping plants and greatly increased treatment

plants to handle this increased motif.

This leaves us to compare the potential costs of environmental harm with precautionary costs of control. The cost of precautionary control, for stormwater runoff, seems to be very high (e.g., about \$1.5 billion capital cost alone for the City of San Francisco). If indeed stormwater runoff constitutes primarily a problem of transitory, aesthetic nature, the cost of a false negative will be low relative to the cost of a false positive, and may even be less than the cost of a false positive. To the extent that stormwater runoff poses a toxics problem, the question is how much the proposed treatment would control toxics. If little, then there is little precautionary benefit for toxics, and the cost of a false negative, for this type of control is retrofitting an existing system, it appears that D may be close to 1 or even negative.

In the case of a stormwater overflow, the immediate effects are well understood because the overflows have already been ongoing for many years, and continuation of the present method of operation is not very risky in terms of encountering large unknowns, for the short term effects.

A long-run solution for cities with combined sewers is to replace them gradually over a long period of time with a system of

separate sewers, realizing that sooner or later sewers have to be replaced on something like a hundred-year cycle.

In this second example, even though the environmental effects are still uncertain, there is less motivation to proceed than there was in the first example, because D appears to be smaller in the second example. As we have seen, D is by no means the only consideration, but it is useful to distinguish among different patterns of costs of false positives and false negatives, in different decision problems. It is also worth noting that the Environmental Protection Agency has required some cities to adopt very expensive management programs for controlling wastewater overflows, such as in San Francisco, so that this analysis of potential costs is of direct practical importance.

2. Acceptable Risk

Health standards for water quality and hence effluent discharge depend directly on a concept of acceptable risk. The science-first approach leads toward one notion of acceptable risk, and the policy-first approach another. The two approaches treat both uncertainty and evaluation differently, as will be shown in terms of a statistical example. One conclusion to be drawn from the example is that these differences are less important to their policy implications for Class I and II pollutants than for Class IV pollutants.

In the science-first approach attention is focused on the role of a false positive. First it must be shown that, with a high degree

of confidence, some harmful effect is associated with a pathogen or chemical. Once that is done, an ample margin of safety is allowed to cover for the uncertainties associated with the magnitude of effect and pathways to exposure. The final result is an "acceptable daily intake" or a "tolerance limit." Generally, in the determination of risk, the cost of a false negative is not a consideration, nor its probability, nor the cost of a false positive. The analysis turns upon the probability of a false positive, and then the safety factor, which is more or less arbitrarily chosen. In a sense, the uncertainties which arise at each stage of the analysis are saved up until the end, where they blossom in the margin of safety.

"Acceptable daily intakes" (ADI) and "tolerance limits" should be interpreted with caution. First, as already noted, of the many thousands of chemicals of concern, only a small fraction have been tested. If an ADI or "tolerance limit" does not exist for a chemical, this does not imply that the chemical is safe in any exposure, although the chemical is often treated as though this were true. Second, some of the ADIs and tolerance limits are primarily based on tests for acute toxicity and then projected into limits for chronic toxicity (or genetic toxicity). But the relationship between acute toxicity and genetic toxicity can vary widely. Third, the validity of a large number of the tests is now uncertain. For one example, the largest testing firm in the world, Industrial Bio-Test, went bankrupt amid charges of fraud and incompetent testing, leaving in its wake some 600 tolerances with questionable legitimacy. Fourth, the interactions or

synergisms along chemicals in combination may be far more important than the chemicals acting separately. Generally the ADIs and tolerance limits are based on tests of chemicals in isolation. Yet there are cases where a non-carcinogen potentiates a carcinogen a thousand-fold. These sorts of interactions will be missed entirely in many of the tests leading to ADI and tolerances. But these possible interactions are clearly important for risk assessment of wastewater and sludge disposal, partly because of the large numbers of chemicals in the waste stream and partly because of the large number of chemicals in the receiving waters. Thus for all these reasons, it is important not to take ADIs and tolerances at face value as though they provided sufficient protection with virtual certainty, because of their safety factors.

We consider here one source of uncertainty which arises in the course of testing. For testing chemicals of suspected carcinogenicity, the most respected and widely used test is some form of the National Cancer Institute bioassay. It is useful to discuss this test here for several reasons: the probabilities of the false positives and false negatives are well defined and can be calculated explicitly; the nature of the trade-off between α and β becomes clearer; and the difference between uncertainty within a model and uncertainty as to model specification is illustrated. And it will illustrate why the difference between the science approach and policy approach is more important for Class IV pollutants than for Class I.

The NCI bioassay is a more or less standard test, involving 50

rodents (usually mice or rats) in a control group, and 50 in each treated group. Often there are two treatments, at different doses of the chemical under investigation. The two sexes are typically treated in separate groups; often two strains, or two different species, are used; and about twenty sites on the animal are studied for pathological response over the lifetime experiment. The Fisher exact test is usually used to interpret the number of cancers in the treated and control groups. For many of the most important test species and strains, the background level of cancer in the controls is about 2 percent for most sites (background rates for breast cancer in women and lung cancer for both men and women are several times this).

If a chemical is a potent carcinogen for which there is little other toxic effect, the Fisher exact test is very likely to detect the carcinogen (the false negative probability β is low). Suppose, for example, that a chemical has no acutely toxic effects and can be fed to animals to elevate the risk of cancer to be 15 times larger than its original background rate of 2 percent, which is the same rate as in the controls. In this case the expected percent of cancers in the treated group is 32 percent ($15 \times 2 + 2$). For such a potent carcinogen there is only a 3 percent chance of a false negative ($\beta = .03$). This calculation assumes that the Fisher exact test is set in the traditional way with the probability of a false positive α equal to 5 percent.

However, when the chemical has acute toxicity as well as long term cancer effects, it may not be possible to give high enough doses to achieve an excess cancer risk of 15 fold. Suppose we are dealing

with the same carcinogen but it has acutely toxic effects, and we have to cut down the dose to one third or one quarter its previous level. At the lower dose level, the excess risk of this carcinogen is much smaller, perhaps 5 fold increase over the background level, with an excess risk of 10 percent. The chance of a false negative in this case is 54 percent.

This example suggests why it is so important to conduct analyses of statistical power, when dealing with toxic chemicals and other problems of environmental risk.¹⁰ In many cases the probability of detection (β is the probability of non-detection), depends very sensitively on the magnitude of the effect being investigated, as we have just seen. A similar sensitivity of β with respect to the magnitude of the environmental effect has been shown in a study (Harris, Page and Reiches, 1977) on cancer and drinking water (this study is discussed in the next section). In testing, a negative result can mean either that there is no effect or that the effect exists but is undetected. There is simply no way to tell which inference is more appropriate unless β is calculated, as a function of the magnitude of the effect of concern. In other words, there is virtually no meaning in a negative result without calculation of β (a power analysis). In spite of this fact few studies of environmental effects in coastal waters -- or elsewhere -- undertake analyses of statistical power.

For risks to human populations small increases over background rates can be very important. For example, it is not uncommon for an excess risk of 10 per million to be considered unacceptably high.

(Just where one draws the line depends on cold-blooded balancing of the benefits of the chemical and magnitude of effect, along with consideration of the ethics in the imposition of grave harms.) Because detectability is such an important problem, in the investigation of potentially toxic chemicals, the National Cancer Institute, the Environmental Protection Agency, and the National Toxicological Program require maximally tolerated doses, as well as lower doses, in long term bioassays, to decrease the probability of false negatives. But there are, of course, limitations to the sizes of doses which can be given, often because of acute toxicity, and we are not able to achieve low β s for existing bioassays, for many of the effects of concern.

Masking of chronic effects by acute effects is especially important when large numbers of people are to be subjected to very low doses. This of course may be the case for chemicals discharged into ocean waters and subject to dispersal and bioaccumulation. There is some evidence that chronic effects of genetic toxicity become linear with very low doses, while acute effects exhibit "safe thresholds." If true, then, as the dose is decreased but the number of people exposed increases in proportion, the importance of the chronic effect increases relative to the acute effect.

The question also arises as to whether low α s can actually be achieved, due to multiple testing. Traditionally, for a given chemical under test the probability of a false positive is set for one site and one dose. Even if there is no effect at all, it appears that with 20 sites and 2 doses there is an excellent chance of a false positive

somewhere in the forty experiments. Salzberg (1977) raised this question and calculated that the chance of a false positive for the composite test would be about 65 percent, if the chemical were considered to be a carcinogen upon a positive for one or more sites (the calculation assumes some statistical dependence among sites). The calculation led to a re-examination of the uncertainties associated with the specification of the statistical model itself. The Fisher exact test presumes no knowledge about the cancer incidence in the controls. Thus the assumption of the model assumes more uncertainty than in fact exists. The historical rates of cancer in the controls is well known, and this provides information on the probability of cancer of the controls. When this information on the background rate of cancer in the controls is taken into account, it is found that the true probability of a false positive is much lower than it would be if there were no information as to the spontaneous cancer rate in the controls. Fears, Tarone and Chou (1977) recalculated the probability of a false positive for a single site and dose when this information is taken into account and then found upper bounds on α for the whole experiment of 20 sites and 2 doses. These upper bounds were between 4 and 10 percent. The recalculations do not affect β .

The first point to be made here is that analyzing the uncertainties in the specification of models is highly important, and at times may be even more important than analyzing the uncertainties within specified models. The second point is that in presetting α to 5 percent there may be an acceptable trade-off between α and β , for

potent carcinogens without competing toxicity, for acute effects of chemicals, and for many pathogens. For these cases it may be possible to achieve α and β both equal to or less than 5 percent. But for weak carcinogens with wide exposure, or for carcinogens with competing toxicity, presetting α to 5 percent is not acceptable without further calculation. It is not acceptable because it may imply a β of 50 percent or more. If the cost of a false negative is high compared with the cost of a false positive, then a decision approach suggests that β be decreased even at the price of increasing α . How much β should be decreased depends on the properties of the test, in particular how fast β can be decreased for a given increase in α ; it depends as well on the relative costs of wrong decisions. (See the Appendix for the relationship.)

The science-first approach begins with mechanisms and effects and avoids judgements about the costs and benefits associated with effects. For this reason it is compatible with presetting α to 5 percent and then minimizing β . The procedure avoids considering the costs and benefits of effects. For Class I and Class II pollutants the procedure may lead to β also 5 percent or less; for Class IV pollutants the procedure may lead to very high β . The policy-first approach begins with consideration of the costs of wrong decisions. For Class I and Class II pollutants trading off α and β to minimize the expected cost of wrong decisions may lead to a very similar choice of α and β as in the science-first approach. But for Class IV pollutants β tends to be larger for a given α and the minimum expected cost trade-off of α

and β is likely to diverge more from the α and β defined by the science-first approach. This means that the two approaches are likely to differ most for Class IV pollutants in their implications as to which precautionary decisions are worthwhile.

To summarize: the science-first approach focuses on the probability of a false positive and on proving the existence of an effect. In the policy-first approach all four ingredients — the cost and probability of a false positive and the cost and probability of a false negative — are symmetric and equally important for a concept of acceptable risk. Where the cost of a false negative is likely to be higher than the cost of a false positive, the probability of a false negative (β) is more critically sensitive for the decision than the probability of a false positive (α). When both α and β can be made small, perhaps both 5 percent or less, this difference in approach is not likely to matter very much. But when there is more uncertainty in the new information (for example, the sum of α plus β is 50 percent or more), and when there is considerable uncertainty as to the likelihood of the underlying hypothesis, then the two approaches are likely to diverge in their implications as to which effects are "positive" and which to be taken seriously for decision purposes.

Comparison with Drinking Water Carcinogenicity

There are now some thirty or more studies of the potential cancer hazard from chemicals in drinking water (for a review see Crump and Guess, 1980). In contrast there appear to be few studies of the

potential cancer hazard from chemicals discharged into salt water. A comparison of the two potential hazards suggests some parallels as well as obvious differences. It also provides an illustration of asymmetries typical in environmental risks.

In 1974 there were 64 organic chemicals identified in drinking water; by 1981 there were over 700. These 700 represent, by weight, only about 15 percent of the organic fraction of drinking water. Some of these chemicals come from agricultural runoff, some fall out from the air, and many come from industrial discharges through municipal sewage systems. Presumably the same chemicals are to be found in wastewater and sludge discharges into coastal waters, but concentrations are likely to be larger in wastewater and sludge, because in drinking water the chemicals are identified "at the tap" rather than at the discharge pipe.

In addition, a large (but unknown) number of halogenated chemicals are synthesized by the chlorination process, acting on organic precursors. In the case of chlorinated drinking water, a principal feedstock appears to be naturally occurring organics in soil humus; in the case of sewage chlorination, Class I pollutants. In a mass balance of chlorination in drinking water, only about 10 percent of the total chlorine content has been identified in product chemicals. Brominated species are formed as well, due to chlorine's greater chemical activity and the presence of bromide ions. Although there appears to be less attention paid to the characterization of halogenated compounds in treated wastewater, the same situation presumably exists in the latter

case, but in intensified form, for two reasons. First there is far higher organic content in wastewater and second, larger concentrations of chlorine are used in wastewater treatment.

The potential human health effect has been investigated by means of two conceptually different methods. In one method -- an epidemiologic method -- human populations have been studied, comparing cancer rates in people with high exposure to industrialized and chlorinated water with cancer rates with low exposure. In the second method animal bioassays are used on single chemicals and the results extrapolated to estimates of human risk. Epidemiologic investigation is now in its third stage. In the first stage, whole populations were studied, with aggregate measures of exposure (for one of the first attempts at this see Page et al., 1976). These studies were primarily useful in constructing hypotheses. For example, they pinpointed cancers of the gastrointestinal tract as of special concern and they suggested that, at least from present cancers resulting from exposures 30 to 40 years ago, chlorination might have been a more important factor than the aggregate presence of synthetic organics of that time.

In the second stage case control studies were conducted. These studies characterized exposures much more precisely, on an individual-by-individual basis. While a number of questions about dose-response relationships and possibly confounding variables remain, the case control studies have tended to confirm and sharpen the evidence for a carcinogenic risk from drinking water (Crump and Guess, 1980). One of the main problems with the existing case control studies is that they are

too small in scale to detect excess risks of 5 to 20 percent over the background as statistically significant at the 5 percent level ($\alpha = 0.05$). These are the excesses suggested by both the ecologic and the case control studies. In the third stage, larger case control studies are being undertaken to increase the statistical power of the tests. But even at this third stage, with sample sizes of 3000, the trade-off of α and β is unfavorable. To detect an increase of risk of 15 percent over background, with traditional probability of false positive set at 5 percent, there is still an 80 percent probability of false negative ($\alpha = 0.05$, $\beta = 0.80$).

In the second method, chemicals are identified in the water, along with their concentrations. Carcinogenic potency is taken from animal studies such as the National Cancer Institute's bioassays.¹¹ This effect is then extrapolated from the high dose of the experiment to the low dose in the water and extrapolated across the species lines from rodents to humans, taking into account differing metabolisms, body weight and lifetimes. The early extrapolation studies focused entirely on chloroform. Later studies included extrapolations on all the chemicals identified in drinking water as carcinogens for which animal studies exist, but chloroform remains the principal chemical of concern for drinking water.

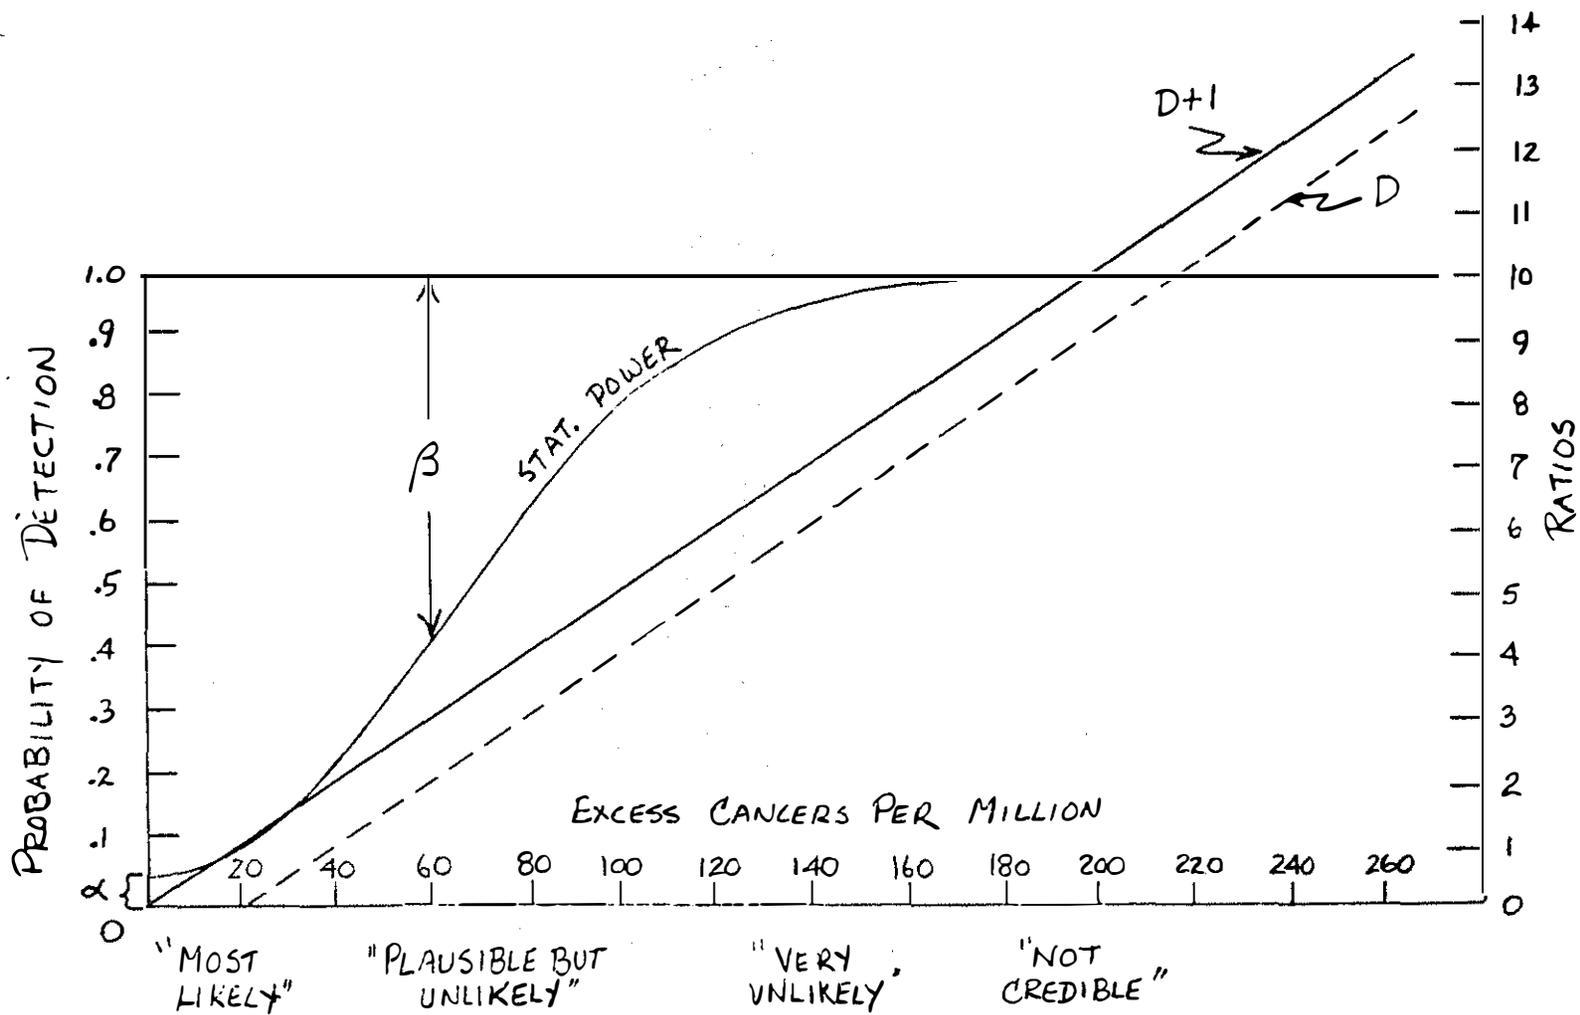
Quantitative estimates arising from two independent data bases provide a check of one on the other. Both suggest a carcinogenic risk. Both approaches are highly uncertain in their quantitative estimates of excess risk and neither proves beyond a doubt that a non-negligible

risk exists. But most interesting, the two approaches yield risk estimates that are within about an order of magnitude of each other, roughly in the range of 0.5 to 5 percent excess over background. Given the uncertainties associated with each method, probably at least an order of magnitude or two for each, the two methods agree within their range of resolution. The rough agreement between the two methods is much more suggestive than the implications of either method taken by itself. (See Crump and Guess, 1980, for a comparison of the two approaches.)

A second observation is that the risk estimated from the animal studies is considerably lower than that from the epidemiological studies. This could be due to chance, or uncertainty in the specification of the models. It could be due to the fact that the extrapolations only take into account the small fraction of chemicals yet identified. Or it could be due to the neglect of chemical interactions in the extrapolation approach. (There are tests for toxicity of wastewater taken as a whole but these are almost always for acute effects.) Because the comparison of risk estimates in drinking water suggests that for longer term effects like cancer synergisms may be important, such possible chemical interactions deserve much more consideration in the setting of standards for coastal waters.

Analysis of the implications of one of the regression equations in an ecologic study of drinking water suggests asymmetries which may be typical for environmental risks associated with ocean discharge of potentially toxic chemicals. (It appears that there are few, if any,

FIGURE 2
PROBABILITY OF DETECTION (STATISTICAL POWER)
AND RATIOS OF COSTS AND BENEFITS



Source: Page, Bruser, in Scientific Basis for Health and Safety Regulation,
Brookings, In Press.

NOTES FOR FIGURE 2

Statistical power, which is read on left scale, is the probability of detection as a function of the magnitude of the effect. One minus the power is β , the probability of non-detection. The probability of (erroneous) detection, when there is no effect, is α .

D is the ratio of the cost of a false negative to the cost of a false positive. D - 1 is the ratio of the benefit of precautionary control to the cost of control. Both are functions of the magnitude of effect, and they are read on the right scale.

The regression equation analyzed is

$$M = 49 - 12(R) - 1.8(I) + 62(W)$$

(9.9) (-4.4) (-3.8) (5.0)

where

M Gastrointestinal Cancer White Males
 R Urbanization
 I Median Family Income
 W Percent Mississippi Water
 t Values in Parentheses

The judgements of likelihood in quotes are qualified versions of N, the subjective odds against the environmental hypothesis. These were informal judgements, made by commentators, at the time of the original research, and never quantified.

The annual costs of control, by granular activated carbon, are estimated for a city of 1,193,000; trihalomethane concentration 250 ppb; GAC 90 percent effective removed; benefit per cancer prevented 500,000; steady state comparison.

corresponding analyses of risk for coastal discharge, so we are left with the fresh water comparison.) Figure 2 shows a regression relating parish (county) cancer rates of the gastrointestinal tract, for white males, in Louisiana, to the percent of the parish drinking water from the Mississippi and other contaminated surface water compared with less contaminated ground water (W), occupational exposure (I), and urbanization of the parish (R).¹² Even though the drinking water variable is highly significant (indicating that a complete switch from ground to surface water is associated with 60 extra cancers per million population per year), the statistical power associated with this point estimate is very low. Under the assumption that the underlying model is correctly specified -- an assumption requiring examination because models are almost always misspecified in some ways -- the regression has the following interpretation. If the number of excess cancers is 140 per million, then there is about a 95 percent chance that the regression will find a drinking water effect, by producing a coefficient for the drinking water variable positive and statistically significant at the traditional level ($\beta = 0.05$ or $p < 0.05$). But if the excess mortality is 62 per million, which happens to be the point estimate of the regression, then there is only about a 40 percent probability of detection. The basic reason for low statistical power is that there are large numbers of people at risk, but the risk for each is small (the "binomial variance" problem). The "binomial variance problem" would also apply to any effort to quantify the risks from chemicals discharged in wastewater and sludge in coastal waters. In addition, the problem would apply in intensified form, due to the greater

dispersal and greater numbers of people, and other species, ultimately exposed if there is bioaccumulation. The limitation of statistical power from binomial variance for this regression equation is graphed on Figure 2. As would be expected, the larger the actual effect, the greater is the statistical power.

The cost of precautionary control, computed for a city of 1.2 million considering drinking water filtration by granular activated carbon, was estimated by EPA to be \$10.3 million annually.¹³ The annual benefit of precautionary action depends linearly on the magnitude of the cancer risk, the effectiveness of the treatment, and the value attributed to each cancer prevented. Thus the ratio of benefit-to-cost is a linear function of the size of the carcinogenic effect and is graphed on Figure 2. It can be seen that if there were 100 or more excess cancers per million then the chance of detection would be about 80 percent ($\beta = 0.20$); moreover, with this large an effect the benefits of control would be roughly five times the costs. However, the principal area of interest is not at 100 excess cancers per million. In accordance with the existing information at the time, the likelihood of such an extreme effect would be considered extremely low. The principal area of interest is where the statistical power is low -- in the range of 0-3 excess cancers, which at the time was considered increasingly unlikely at the high end, but still plausibly credible. But where the initial assessment of likelihood becomes higher, the statistical power diminishes, and the cost-benefit ratio becomes smaller.

The problem can be compared with searching for a needle in a

haystack. First of all, the needle may not be in the haystack at all (the low probability of the adverse hypothesis, or high N). But if the needle is in the haystack, it is hard to find (low statistical power in the area of interest, or high β). But if the needle is present and we miss it, we pay a high price (a high cost of the effect compared with the cost of precautionary action), or high D . Clearly in such a case it is important to analyze carefully the sources of uncertainty and to search for statistical models and tests that are sufficiently powerful to be relied upon, especially in the interpretation of negative results.

IMPLICATIONS FOR THE PROCESS OF POLICY FORMATION

The policy approach begins with questions with decision consequences, and then works back to the health and environmental effects of coastal discharge, and thence to the scientific models which predict the effects. In this approach, the value of information is the flip side of the cost of uncertainty. With uncertainty there is the risk of a wrong decision, from either a false positive or a false negative. The value of some new research or new piece of information is the reduction in the expected cost of wrong decisions. Even though our understanding of biological and chemical systems is too fragmentary to follow this approach in a fully quantitative way, it still suggests specific directions for the process of policy formation, some of which are discussed here.

Linkage

In order to evaluate the consequences of some decision, it is necessary to link the outputs of one submodel to the input of another from waste treatment (Chapter 2) to the initial mixing (Chapter 4) to the effects (Chapters 7,8,9) to the control options (Chapters 12,13). This book is a step toward linkage, as it brings together large amounts of relevant material from a wide variety of sources. The policy approach outlined in this chapter suggests continued and further attention on the linkage of the submodels. In terms of Figure 1, this means further attention on links (c) and especially (b). This attention is in addition to (a), which is sometimes the consuming focus of independent scientific studies, in their concern for submodel development and validation. Dominance of the science-first approach may help explain the relatively minor discussion in this volume of two subjects — fish tumors and the mussel watch program. Both are mentioned as important. But each could deserve a whole chapter discussing their meaning. Because of the pervasive uncertainties these chapters would have to be speculative, defining questions and establishing hypotheses to be investigated by existing knowledge and proposed research. In the policy first approach these chapters would come at or near the beginning and would serve to the search for links (b), (c), and (a) of the following chapters. For a substantive but basically descriptive discussion of shellfish toxicity see National Academy of Sciences (1980); of fish tumors, Kraybill et al. (1977).

From the science-first approach it is perfectly reasonable to postpone speculative evaluation of potential harms when the basic facts

themselves are speculative. The science-first approach concentrates first on the determination of the facts, leaving the evaluations until later when then they are less speculative. But a policy-first approach starts with speculative questions of evaluation.

There are several reasons for concern about fish tumors and other biological indicators. (1) In recent work by Ames and by Messelson, it is becoming apparent that different species react similarly to carcinogens, and they are likely to yield similar rankings of potency, in orders of magnitude. Besides the natural repugnance in eating cancerous meat from fish, tumor-bearing fish are likely to contain carcinogens dangerous to humans just as they proved to be dangerous to the fish. It is believed, for example, that a principal route of human exposure to PCBs is from fish consumption, especially fish from particularly contaminated bodies of fresh water. (2) It is possible that toxic chemicals may disrupt fish and other aquatic populations, affecting aggregate food resources. (3) Fish and other aquatic organisms, including plants, may provide an "early warning system" for later ecosystem disruptions. Early warning systems are particularly important in the investigation of potential hazards to coastal discharge of wastewater and sludge, due to possible long dormancies of materials deposited in bottom sediments. Two decades ago it was thought that mercury discharges would safely remain on the bottom. Methylization came as a surprise. A decade ago it was thought that ground water contamination would be a minor problem due to natural filtering in the soil. In the 1980s serious groundwater contamination may become

another surprise. It is useful to anticipate and investigate potential surprises before their causative factors become irreversibly beyond control. (4) Fish and other organisms allow for the study of the effects of chemicals in combination.

Beginning with the speculative questions as to the meaning of the potential effects on human health and on marine populations helps to formulate hypotheses and research agendas to make the evaluations less speculative. In doing so, it forces us to link the various submodels and effects models together. Key areas of uncertainty are identified, for which resolution would establish the links more strongly. The point is that beginning with the evaluative questions helps focus the scientific activity along lines most useful for decision purposes. The point is not that the existing evidence of fish and shellfish contamination already proves that there is an enormous problem requiring immediate remedial action. When the effects are better understood, it may turn out that contamination problems are small, requiring only a little control here and there; or it may turn out otherwise.

Combining Science and Policy Approaches

The "ultimate" policy questions listed at the beginning of the chapter do not spring from a vacuum. In large part, the existing science suggests the policy questions worth worrying about. Thus the process of policy formation needs to encourage both approaches simultaneously, one beginning with science questions, the other with policy questions.

It is sometimes suggested that the work of the policy analysts and decision makers be kept separate from the work of the scientists, perhaps because the questions of evaluation might taint the science. A division of labor is fine, but only insofar as the divided labors fit together in a common framework. The one offered here is that of minimizing expected cost. In this framework it is not possible to define a positive test result without making evaluations. This is because α and β trade off. Nor is it possible to define a notion of acceptable risk without at least qualitative cost and benefit evaluations. The assessment of N is a scientific judgement; the assessment of D an economic one. Yet the two play entirely symmetric roles in parts of the minimum expected cost calculation, as illustrated in the Appendix. This observation is not surprising because in a definition of expected value, N plays the role of a probability, D the value of a consequence (expected values are probabilities times outcome values). In the absence of observations, past definitions of a positive test have been attempted on "scientific grounds alone" and there have been similar attempts to define acceptable risk, and hence standards, as a scientific matter.

Taking these observations into account means that the communication links between the scientists and the policy analysts need to be considerably improved. In the past, it has sometimes been considered sufficient to report a test positive or negative at some significance level, usually $\alpha = 5$ percent. This has given way to reporting p values and confidence intervals. It is necessary to take the next step to report β values, for different magnitudes of effect of concern, as a

routine matter (confidence intervals yield different information from β values). It is necessary for the scientist to report further information on uncertainty as well. The α and β values and the results of the test are computed under the assumption that the underlying statistical and physical model is correct. Generally, there are uncertainties as to the specification of the model, and these uncertainties must be reported as well. A test of whether or not enough information is being communicated from the scientist to the decision maker is met when the decision maker can himself define what is a positive test. Going the other way, the policy maker has the responsibility of communicating to the scientist which areas of uncertainty would have the highest payoff in resolution. There are a countless number of interesting scientific questions to pursue, and the scientist needs information as to which are the most important for decision purposes. As a part of this communication, the costs of wrong decisions need to be evaluated, in at least a rough and perhaps qualitative way, and communicated to the scientist.

Surprises are more likely when we act as though existing models were true, when in fact there is considerable uncertainty associated with most model specifications. Thus better communication of the uncertainty of model specification helps guard against surprises. Other potential surprises — besides the classic example of methylation of mercury — might be identified if specification uncertainties were more intensively characterized and communicated.

A vital need which becomes apparent from the policy-first

approach is the need to generalize. There are simply too many chemicals to consider each on a one-by-one basis. Generalization is difficult when the underlying phenomena are not clearly understood; however, scientists have responded to this need in a creative and useful way. Chapter 6 summarizes and extends generalizations which help carry us beyond a one-by-one approach to chemicals. Particularly useful, for decision purposes, are generalizations based on an effect of concern, such as the octanol-water partition coefficient, which is a predictor of bioaccumulation. A great deal of work is now being undertaken on generalizations which would help predict which chemicals are likely to be genotoxic. Another important effort in generalization, one which so far has received little resource commitment, is to investigate the general principles by which chemicals are likely to act synergistically.

IMPLICATIONS FOR SPECIFIC POLICIES

There will be further discussion of specific policies in Chapter 13, but some islands of relative certainty can be tentatively identified here. In searching for a "relative certainty" it is useful to look for places where the evidence, uncertain as it is, has tended to strengthen in a particular direction over time.

More on Toxics

Evidence has increased in the last two decades that toxic chemicals should be taken more seriously as a potential problem. This chapter suggests that traditional approaches to standard setting and

acceptable risk are less appropriate to Class IV pollutants than to the others, and thus the traditional approach is most in need of rethinking in this area.

Redirection of Secondary Treatment

To some extent requirements for secondary treatment were established to control toxic substances. However, it is becoming increasingly clear that the most effective way of controlling toxic chemicals is by pretreatment, production controls, and source reduction. To some extent requirements for advanced treatment arose in response to the goal of minimizing discharge for all classes of pollutants into ocean waters. But in the past decade the costs of other forms of disposal of the increased volume of sludge from advanced treatment have been better understood. Incineration of sludge and land fill are seen to be increasingly more expensive and the concomitant environmental problems more apparent, thus leading to a rethinking of the reasons for and costs of advanced treatment. One otherwise attractive alternative, land use of wastewater and sludge for irrigation, soil conditioning and fertilization, largely depends on the control of Class II, III and IV pollutants. And again, the only effective control of Class III and IV pollutants appears to be pretreatment or source control.

More Trade-offs in the Receiving Media

Legislation to control wastes was enacted, in large part, on a case-by-case basis, with ocean and river disposal considered separately

from land disposal, and separately from incineration (air disposal). In dealing with the problem of waste treatment and disposal compartmentally, it was often required that the waste, in each case, go elsewhere. But in considering the problem as a whole, it is clear that it all goes somewhere, except for source reduction or treatment that transforms problem materials. In the past decade there has been increasing interest in comparing and trading off the costs of disposal in one medium versus another. This is true for all four classes of pollutants and further resolutions of uncertainty are likely to strengthen the case for managing wastes on a multi-media basis.

More Regionalization

The area of least uncertainty, in modeling the life cycle of wastewater discharges, is in the diffusion and transport of materials. A great deal is known about the way plumes move and how basins flush, or don't flush (Chapters 3,4,5). Regional differences in mixing and transport are well understood and have reached the point of fine tuning. It is useful to take advantage of our knowledge of regional differences in designing treatment and discharge systems.

APPENDIX

In minimizing the expected cost of wrong decisions, we have three alternative strategies: gather no further information but take precautionary action on the basis of existing information, with expected cost

$$N/(N+1); \quad (1)$$

gather no further information and take no precautionary action on the basis of existing information, with expected cost

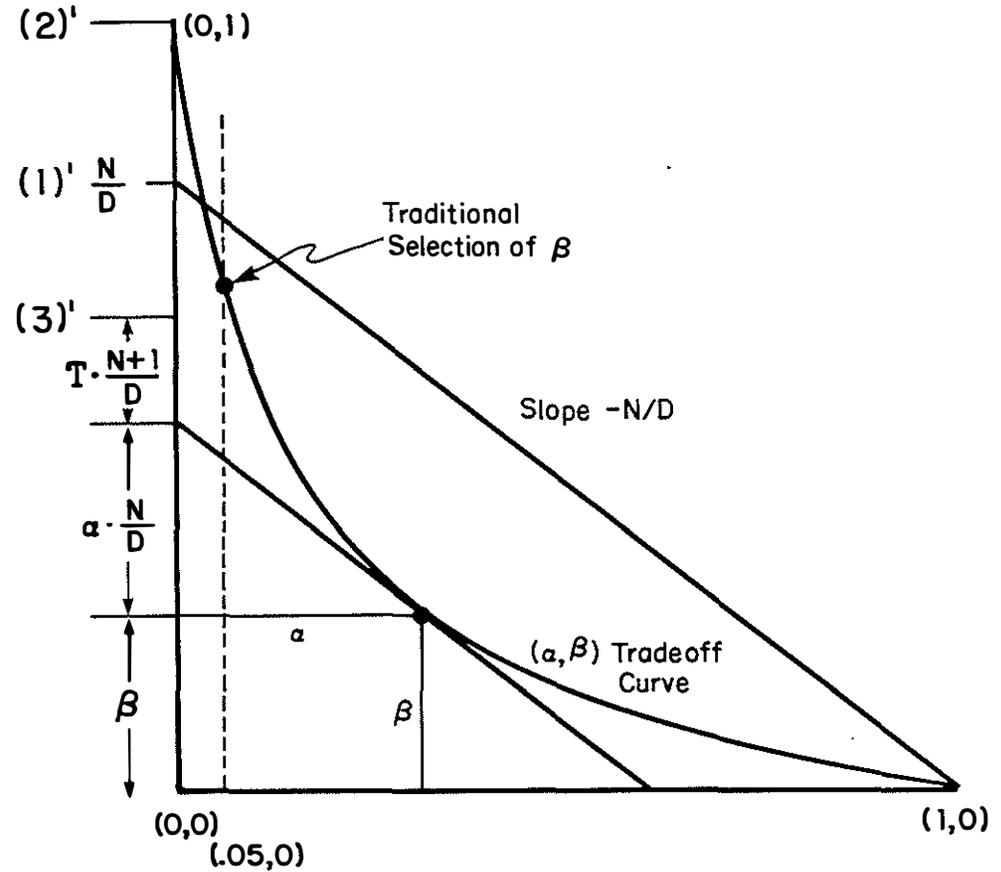
$$D/(N+1); \quad (2)$$

and gather further information, and act on the basis of that, in conjunction with existing information

$$D\beta/(N+1) + N\alpha/(N+1) + T \quad (3)$$

In considering (3) we pick the most advantageous trade-off of (α, β) , the pair leading to minimum (3). For a given cost of the test the trade-off opportunities are shown in Figure 3. By increasing the cost of the test (T), the whole trade-off curve can be pushed in toward the origin. The best scale of the test, or research program, is achieved when the gains from pushing the trade-off curve inward are just offset by the increased cost of gathering the new information (increased T).

FIGURE 3
GEOMETRY OF COST MINIMIZATION



For expository purposes, we are discussing a simple hypothesis and a simple alternative (a hypothesis with a single environmental effect, or no effect at all). In this case the trade-off curve defines the "most powerful test" for each significance level. In the more complicated case of many hypotheses and many alternatives the approach of minimizing expected cost is directly generalizable.

To show the relative rankings of the three strategies, and the best choice for the (α, β) trade-off, we change the scale by multiplying each of the expected costs in (1), (2), and (3) by $(N+1)/D$. We seek the minimum of (1), (2), and (3) by seeking the minimum of

$$N/D \quad (1)'$$

$$1 \quad (2)'$$

$$\beta + \frac{N}{D}\alpha + T(N+1)/D \quad (3)'$$

By drawing the line segment from (1.0) with slope $-N/D$, (1)' shows up on the y-axis as marked (1)'. For comparison (1)' is just a unit up on this same axis. To compare the best (3)' with the (1)' and (2)', draw the tangent with slope $-N/D$ to the trade-off curve. The α and β at this tangency comprise the cost minimizing trade-off of α and β for the given test cost T . Where this tangent hits the y-axis add $T(N+1)/D$, to locate (3)'. Of the three strategies the one lowest on the y-axis is the least cost minimizing alternative.

It can be readily seen that any other choice of α and β other than the one picked out by the tangent with slope $-N/D$ would have a

higher expected cost than for (3)' defined by the tangency solution (the tangent leads to the lowest (3)' on the y-axis). In particular note that the traditional way of selecting the (α, β) trade-off is very unlikely to pick out the cost minimizing combination. In the traditional selection of the (α, β) trade-off is chosen to be 5 percent and β is chosen as whatever is compatible with this preselection. In other words where vertical dashed line through x-axis at 0.05 crosses the (α, β) trade-off curve, the trade-off of α and β is chosen.

The geometry of Figure 3 shows how the factors discussed in the Chapter combine to determine the expected cost: the existing information translated into a statement of the level of suspicion about the adverse hypothesis (N); the relative costs of precaution and the potential risk (D); and the sharpness of the test or additional information in resolving uncertainty, summarized in the (α, β) trade-off curve. Figure 3 shows that in one part of the cost minimization N and D play, entirely symmetric roles, and in another part they play different roles. Figure 3 suggests that it is not worth obtaining fine tuned information on D if we are very vague as to N and vice-versa. Moreover, Figure 3 shows the difference between the approach of cost minimization and the one of classical hypothesis testing.

Footnotes

1. An engineering approach which attempts to optimize over a whole system, like the policy approach, may also begin with general questions of costs and benefits (list A) and move toward specific ones of modeling. The problem, of course, is to resolve questions like those in list A sufficiently to define an objective function which can be optimized. When the objective function is narrowly and exogeneously prescribed, for example, "design the least expensive tertiary treatment plant meeting the following specifications . . .", the approach becomes suboptimization.

2. The inventory of the Toxic Substances Control Act (Section 8(b)) contains about 65,000 chemicals in active commercial use. This inventory does not list a number of pesticides and other chemicals controlled by other laws. The estimate of 10,000 additional chemicals itself is highly uncertain, and indeed the number on the inventory is not definite, due to problems of listing. Thus there is considerable uncertainty in the number of chemicals in active commerce, let alone the number and mass flows of chemicals in municipal sewage systems.

3. Even for the "simple" cases quantitative cost-benefit analysis is difficult in practice, and sometimes there are legal obstacles to trading off control options, across media. (See Freeman [1979a] for a review of cost-benefit studies of water pollution control.)

4. The effects of a Class II pathogen producing cholera are more reversible than the effects of a Class IV chemical producing cancer, in the sense that the effects are longer term for the latter. Once cholera is controlled, the full effects of disease reduction are almost immediate. But once PCBs are controlled it will take decades, perhaps centuries, before the original position is regained. Of course, for the individual who dies of either cholera or cancer, the effects are irreversible in either case.

5. These characteristics are: strong uncertainty as to mechanism, relatively modest cost of precautionary action in comparison with the potential severity of the environmental harm, low assessed probability of harm, irreversibility of potential harm, latency of harm, external costs, internal benefits of risk, and wide dispersal of risk (collective risk). See Page (1978).

6. A test false positive can be distinguished from the false positive of precautionary control in the previous discussion. But when tests are followed one mistake leads to the other and the two concepts of false positive become linked.

7. Expected cost is the probability of harm times the potential magnitude of the harm. When there are several possible severities and probabilities these products are summed, as in the third case below.

8. There are several complicating factors. First, it can be asked: Is not the cost of remedy overstated, due to postponement of construction and the discounting of future construction costs? Ordinarily one would

say ye , but the matter is complicated by the fact that the D should "really" be defined in terms of potential environmental costs, and in some sense the remedy costs are a proxy. Second, some of the potential environmental costs will begin to accrue right away, under a false negative, so that these should not be discounted. Third, the problem of discounting becomes deeper when considerations of intertemporal equity are taken into account (see Page, 1980; and Ferejohn, 1978). And fourth, there is the problem of establishing incentive compatibility.

Suppose for the sake of illustration that the incremental environmental harm, if the extension is not built, is three units. Because the cost of extension at the time of original construction is only one unit, the extension is worth undertaking. However, there is an incentive for the sewerage agency to minimize construction costs, subject to meeting predicted water quality constraints. This situation produces an incentive for the agency to predict low environmental impacts. Once the outfall is built too small, the incremental cost calculation goes the other way. Now it costs six units to control three units of environmental damage, and the agency can legitimately argue that the incremental construction is no longer worth the cost.

Thus an agency with an incentive to minimize its internal construction and operation costs has an incentive to underbuild at the time of original construction. To offset this incentive, the agency might be told beforehand that remedy is likely to be required. The state or federal agency may develop the following policy. It announces to the sewerage agency before there is any construction that if it turns

out that the environmental costs are 3 units, the sewer agency will be required to retrofit, even at the cost of 6 units. With this policy, there is greater incentive for the agency to build the outfall longer in the first place, for a net saving of 3-1 or 2 units. (See, for discussion of rules of remedy, in another context but still for expected cost minimization, Kronman and Posner (1979).) The cost of remedy would be a consideration in setting up such a policy rule. Establishing a policy rule to minimize expected costs across the whole decision process, as in this case, is an example of the "principle-agent problem."

9. Personal communication, N.H. Brooks, February 1981.

10. "Statistical power" and "probability of false negatives" are equivalent concepts. The first is the probability of detection, the second the probability of non-detection. They are both functions of the magnitude of effect being investigated, and for each magnitude of effect they sum to one. Statistical power is $1-\beta$.

11. One definition of potency is the daily dose in mg per kg of body weight leading to cancer in 50 percent of the treated animals, corrected for competing mortality and cancer in the controls.

12. The statistical analysis is discussed in more detail in Harris et al. (1977).

13. This cost-benefit analysis is discussed in more detail in Page et al. (in press).

References

- Brooks, N.H. (1971) Statement at hearings, Public Hearing regarding Proposed General Principles and Provisions for Discharge to Ocean Waters, California Water Resources Control Board, San Diego, Dec. 2, 1971.
- Crump, K. and Guess, H. (1980) "Drinking Water and Cancer: Review of Recent Findings and Assessment of Risks," prepared for the Council on Environmental Quality.
- Fears, T.R., Tarone, R.E. and Chu, K.C. (1977) "Error Rates for Carcinogenicity Screens," Cancer Research, Vol. 37.
- Ferejohn, J.A. and Page, T. (1978) "On the Foundations of Intertemporal Choice," Journal of Agricultural Economics.
- Fischhoff, B. (in press) "Lay Foibles and Expert Fables in Judgements about Risk," in Progress in Resource Management and Environment Planning, Vol. 3, T. O'Riordan and R.K. Turner, Eds. Chichester: Wiley.
- Freeman, A.M. (1979a) "The Benefits of Air and Water Pollution Control: A Review and Synthesis of Recent Estimates," prepared for the Council on Environmental Quality, Washington, D.C.
- Freeman, A.M. (1979b) The Benefits of Environmental Improvement: Theory and Practice, Baltimore: Johns Hopkins Press.
- Harris, R., Page, T., and Reiches, N. (1977) "Carcinogenic Hazards of Organic Chemicals in Drinking Water," in Origins of Human Cancer, Cold Spring Harbor, N.Y.: Cold Spring Harbor Laboratory.
- Hogarth, R. (1975) "Cognitive Processes and the Assessment of Subjective Probability Distributions," Journal of the American Statistical Association, Vol. 70, No. 350.
- Kraybill, H.F., Dawe, C.J., Harshbarger, J.C. and Tardiff, R.G. (eds.) (1977) Aquatic Pollutants and Biologic Effects with Emphasis on Neoplasia, Annals of the New York Academy of Sciences, Vol. 298.
- Kronman, A.T. and Posner, R. (1979) The Economics of Contract Law. Boston: Little, Brown and Co.
- Leamer, E. (1978) Specification Searches: Ad Hoc Inference with Nonexperimental Data. New York: Wiley.
- Morel, F. and Schiff, S.L. (in this volume) Biogeochemical Transformation of Municipal Waste in Coastal Waters, this volume (Chapter VI, pp. 93-4 in ms.).
- Morgan, G., Henrion, M. and Morris, S. (1979) "Expert Judgements for Policy Analysis," Report of a Workshop at Brookhaven National Laboratory, July 8-11, 1979.
- National Academy of Sciences (1980a) Regulating Pesticides, Washington, D.C., p. 8-9.
- National Academy of Sciences (1980b) The International Mussel Watch, Washington, D.C.
- Page, T. (1978) "A Generic View of Toxic Chemicals and Similar Risks," Ecology Law Quarterly, Vol. 7, No. 2, 207-244.
- Page, T. (1980) "A Kantian Perspective on the Social Rate of Discount," Horst Siebert (ed.), Erschopfbare Ressourcen, Berlin: Duncker and Humblot.
- Page, T. (in press) "A Framework for Unreasonable Risk in the Toxic Substances Act," Proceedings of the New York Academy of Sciences.
- Page, T., Harris, R. and Epstein, S. (1976) "Drinking Water and Cancer Mortality in Louisiana," Science, Vol. 193, No. 4247.
- Page, T., Harris, R. and Bruser, J. (in press) "Removal of Carcinogens from Drinking Water: a Cost-Benefit Analysis," in Scientific Basis for Health and Safety Regulation, Washington, D.C.: Brookings Institution.
- Pauker, S. and Kassirer, J. (1980) "The Threshold Approach to Clinical Decision Making," New England Journal of Medicine, Vol. 302, No. 20.
- Raiffa, H. (1970) Decision Analysis: Introductory Lectures on Choices Under Uncertainty, Reading, Mass: Addison-Wesley Publishing Co.
- Salzburg, D.S. (1977) "Use of Statistics When Examining Lifetime Studies in Rodents to Detect Carcinogenicity," J. Toxicology and Environmental Health, Vol. 3.
- Savage, L.J. (1971) "Elicitation of Personal Probabilities and Expectations," Journal of the American Statistical Association, Vol. 66, No. 366.

Weiner, D.A. et al. (1979) "Exercise Stress Testing," New England Journal of Medicine, Vol. 301, No. 5.

Weinstein, M. (1979) "Decision Making for Toxic Substances Control: Cost-Effective Information Development for the Control of Environmental Carcinogens," Public Policy, Vol. 27, No. 3.