

September 14, 1983

ECONOMIC BENEFITS FROM CONTROL OF MAJOR ENVIRONMENTAL EPISODES

by

Mark F. Sharefkin
Resources for the Future
1755 Massachusetts Avenue, N.W.
Washington, D.C. 20036

EPA Cooperative Agreement CR 807 901 010

Project Officer:

Dr. James Hibbs
Office of Research and Development
U.S. Environmental Protection Agency
401 M Street, S.W.
Washington, D.C. 20036

OFFICE OF RESEARCH AND DEVELOPMENT
U.S. ENVIRONMENTAL PROTECTION AGENCY
WASHINGTON, D.C. 20460

DISCLAIMER

Although prepared with EPA funding, this report has neither been reviewed nor approved by the U.S. Environmental Protection Agency for publication as an EPA report. The contents do not necessarily reflect the views or policies of the U.S. Environmental Protection Agency, nor does mention of trade names or commercial products constitute endorsement or recommendation for use.

ABSTRACT

This study is a reconnaissance of methods for estimating the economic benefits from preventing major environmental episodes. The relevant methods are drawn from the theory of the cost of risk-bearing and from decision theory more generally. The reconnaissance is conducted in three case studies. Those studies span environmental pollution episodes in three transport media: surface water, groundwater, and air. The three case study episodes are Kepone contamination of the James River (surface water), contamination of the Cohansey aquifer by Price's Landfill (groundwater), and the Chemical Control incident (air). For each case study, the principal damage categories are identified, and data assembled from which those damage categories are estimated. In many cases the estimated damages are surprisingly large and suggest the cost-effectiveness of measures for prevention of such episodes. In the body of the report, standard expected utility theory concepts are used for estimating the cost of bearing the risks associated with major environmental episodes. But because both theory and experiment now suggest that those concepts may systematically underestimate those costs, an appendix is devoted to the statement and analysis of alternative measures.

CONTENTS

ABSTRACT	i
Chapter 1. Introduction	1
Chapter 2. Assessing and Valuing Risks Posed by Major Environmental Episodes: A Methodological Overview	17
Chapter 3. Kepone Contamination of the James River	40
Chapter 4. The Benefits of Avoiding a Kepone-Type Incident	57
Chapter 5. Contamination of the Cohansey Aquifer by Price's Pit	78
Chapter 6. The Benefits of Avoiding a Price's Pit-Type Incident	92
Chapter 7. The Chemical Control Incident	109
Chapter 8. The Benefits of Avoiding a Chemical Control-Type Incident	125
Chapter 9. Summary, Conclusions, and Research Needs	139
Appendix A. Other Candidate Case Studies Considered	151
Appendix B. Kepone Case Study Data Appendix	170
Appendix C. Price Case Study Data Appendix	191
Appendix D. Chemical Control Case Study Appendix	211
Appendix E. The Cost of Risk-Bearing and Anxiety	222
References	248

CHAPTER 1

INTRODUCTION

THE LEGACY OF CHEMICALIZATION

The mix of goods available to the American consumer in the 1980s would be almost unrecognizable to a consumer of the 1940s, and the same can be said for plant and equipment available to businessmen and farmers. And to a large extent, those differences were made possible by the remarkable inventiveness of the American and European chemical industries in the post-World War II period. Synthetic organic chemistry, the scientific basis for that inventiveness, was alive and well in the nineteenth century. But only after the war did that industry rise to its present prominence, during a period which saw the longest sustained expansion of the Western economies in recorded economic history.

That expansion came to an end in the stagflationary late 1970s and, coincidentally, so did any residual belief that the chemicalization of the economy was an unalloyed benefit. The skepticism dates from at least the 1950s) when Rachael Carson's ecological fantasy Silent Spring (1962) contributed to the stirrings of the environmental movement; Carson painted a picture of a poisoned world stilled by nondegradable, environmentally persistent pesticides

Three decades have passed since the publication of Silent Spring. In the early 1980s, the legacy of chemicalization has come to dominate the environmental news, and on some evenings all the news. That legacy has provoked public concerns of a kind that local, state, and federal government officials had long associated with race, communism, and public morality. And focusing and highlighting those concerns have been a series of incidents, seemingly unending and inexhaustible. The pattern is depressingly familiar. Suddenly, often serendipitously, hazardous substance-contamination of some community's drinking water supply, or building materials supply, or road surfaces, is discovered. Public officials are uncertain of the risk involved, and their uncertainty is perceived, and amplified, by the affected individuals in those communities. The site, and some citizens, appear in national press and national television news accounts of the incident. The incident fades from public view, except for an occasional item about the progress, or lack of progress, of cleanup efforts, or epidemiological studies, or suits seeking financial compensation. And then, one day, the next incident hits the news.

Among the most significant policy responses to this gathering storm of incidents were two major pieces of federal legislation. The Resource Conservation and Recovery Act, passed in 1980, created a system for monitoring flows of hazardous materials through the economy, with the aim of ending the era of malign neglect and careless dispersal. And the Comprehensive Emergency Response, Compensation and Liability Act (or "Superfund"), passed in 1980, levied taxes on every barrel of crude oil received at U.S. refineries, on petroleum product imports into the United States, and on about forty specified chemicals. The proceeds were to support cleanup operations at the more dangerously contaminated waste disposal and plant sites. The Superfund was never to exceed \$1.6 billion, and a sunset provision set 1985 as an expiration date for the tax provisions of the Act.

Both measures were passed by Congress after bitter legislative battles and compromises, and both were hailed as environmental victories upon passage. It is a sad tribute to the longevity of legislative aspirations and illusions that both are now, in 1982, widely perceived as inadequate and seriously flawed. The search for better and more effective policies will go on over the next five or ten years, and will be conducted against a backdrop of general economic disarray and fiscal stringency. New policies will have to pass a tough cost-effectiveness standard, and debate over those policies will, inevitably, be punctuated by new hazardous substance incidents. This book aims at informing that debate in one particular way: by using evidence from some of those incidents, or episodes, to help us gauge what we stand to gain from policies aimed at preventing such episodes. The debate may be more reasoned if all parties can agree on about how much is at stake.

SOME MAJOR ENVIRONMENTAL EPISODES

We will refer to the incidents we have described as major environmental episodes. The term is meant to distinguish those incidents from more routine kinds of environmental pollution, and the distinction seems warranted by the special characteristics of those incidents.

What are those characteristics? Perhaps they are best introduced, and articulated, by describing some representative major environmental episodes. In the course of selecting the three case study episodes reported on below, we reviewed many other candidate episodes, and we have drawn on those, too, in what follows.

The Three Mile Island Accident

The April 1979 accident in Unit 2 of the Three Mile Island nuclear power plant operated by General Public Utilities occurred at a time when the nuclear power issue was high on the nation's policy agenda. As the first major accident in an American commercial nuclear power plant, it dominated the national news for almost one week. In its wake came at least three major postmortems, including the report of a Presidential Commission specifically created for the purpose of investigating, and interpreting, the accident.

Thus we know about as much about the Three Mile Island accident as we are likely to know about any future episode. And two salient facts arise

from the postmortems. The first is that there were no damages in the strict sense: there were no radiation releases and therefore no health damages imposed upon the local community. The second is that, superficially, the accident was eminently avoidable, for at least an hour after the trouble began. Each of these facts has something to tell us about the analysis of major episodes.

Suppose, then, for argument's sake, that we accept those conclusions of "no physiological or physical damages" at face value. Does it follow that there were no economically relevant damages? Even a convinced materialist will, I think, answer that question in the negative. And, as we shall see, the federal courts, which are about as far from being representatives of philosophical mentalism as an august body of representative individuals can be, have already answered it in the negative.

The issue at hand is, of course, the issue of stress, or anxiety, or of psychic cost, names which we will take to mean about the same kinds of things. The principal injury visited upon the local community by the Three Mile Island accident was the anxiety associated with days of uncertainty about a possible catastrophe at the plant, and subsequent months and years of uncertainty about the future performance of the plant. Even in economics, mental states do matter: though Doctor Johnson maintained that the prospect of imminent hanging concentrated the mind wonderfully, decisions taken on the basis of an imminent execution which never occurs will generally be wasteful. To make the rather obvious point as quickly as possible, investments in human or physical capital are postponed or cancelled, and investments in individual and capital mobility are substituted. Both kinds of investments will be wrong if the motivating anxieties are in fact unwarranted.

The fault then lies in the motivating anxieties, but how are people to do better? To ask them to do better is to ask, in effect, that they be able to form accurate probability judgments about questions from which their experience is quite remote, such as questions about the inherent safety of a pressurized water nuclear power reactor. In other, similar contexts, where there are important decisions to be made about which an individual is uncertain, he or she typically hires a supposedly better informed expert. Society tries to arrange things so that the expert uses his or her specialized knowledge in the client's interest. Legal and medical services are the classic examples of this set of arrangements, which are still, as they were in Roman law, called principal-agent relationships.

But in the case of nuclear power plant safety, those who would play the roles of agent are bitterly divided along many lines. It is far from clear who, if anyone, represents the disinterested agent of the public interest. In the Three Mile Island accident, for example, the public was led to believe that there was, for several days during the accident, a serious danger that a hydrogen bubble in the reactor vessel would explode, showering the area with the intensely radioactive contents of the primary cooling loop. This was the position of the technical staff of the Nuclear Regulatory Commission. But the representatives of Babcox and Wilcox, the vendor firm that built the actual reactor core and pressure vessel, maintained that there was no such possibility, and that the public was being unduly alarmed. With hindsight,

it seems that the vendor's technical staff was correct. But even with that hindsight, who is to be believed the next time around? And how is this pattern of serious disagreement between credentialed experts to be reflected in a layperson's probability assessment of a prospective incident?

The point is strengthened by the second. feature of the accident that we have identified as salient: it was eminently avoidable. For at least an hour after the operating staff of the reactor recognized that there was a problem, relatively straightforward actions could have halted the development of the accident sequence. But those actions were not taken. Worse, the operators did take actions to override the automatic plant safety features that might have halted the accident sequence.

And therein lies a lesson for, and a guide to, thinking about major episodes. To draw that lesson somewhat more sharply, consider what actually happened during the accident sequence. In brief, highly radioactive primary coolant water was leaking from the overpressure relief valve, threatening to uncover the reactor core: an uncovered core will melt down in about thirty minutes. The operators received, through the gauges and signals of the operating room, four or five clear indicators of trouble: one indicator susceptible of the interpretation that there was a danger of overpressure in the primary coolant circuit, and the other three or four unambiguous indicators of an underpressure situation--a primary coolant circuit leak. Drawing upon a rule of thumb pounded into all reactor operators during their training--avoid an overpressure situation, almost at all costs--they chose to acknowledge the first signal, and effectively to ignore the others. It is hard to exaggerate the strength of their determination to "stay the course": when the automatic high pressure core injection spray was activated to cool the uncovered core, the operators elected to override that system manually. By the time that they reconsidered their analysis of the situation, a reconsideration largely prompted by discussions with a representative of the reactor vendor, it was too late.

But remember that, with the possible exception of aviation, no technology has been subject to such intensive and persistent scrutiny as civilian nuclear power. In what was the capstone of that effort, the monumental Reactor Safety Study (or "Rasmussen Report"), there is no trace of the Three Mile Island accident sequence. While it is logically possible that Three Mile Island was a realization of one particularly rare event lying much lower in probability than the dominant accident sequences identified and quantified by the study, that logical possibility seems implausible. In fact, many reactor safety experts now believe that, taken together, the class of sequences initiated by small leaks may dominate the class sequences initiated by large leaks. But even more important is another implication of the accident: that a credible estimate of the accident probability could only be based upon a plausible cognitive model of reactor operator error. While cognitive science generally, and the branch of applied cognitive science called human factors research specifically, have grown remarkably, few believe that we are anywhere near being able to design and validate such models. For that reason alone, release probability estimates for major episodes will, in the foreseeable future, be conjectural. We will argue below that there is nothing wrong with using such probabilities as summaries

of belief, but that there is a danger, requiring almost eternal vigilance, of those summaries being promoted into something more.

Kepone Contamination of the James River

During the 1960s, the Allied Chemical Corporation transferred production of the pesticide Kepone to an independent company, Life Sciences of Hopewell, Virginia. In 1974, the Hopewell operation was abruptly halted by the Health Department of the state of Virginia. The state moved only after reports of acute exposure symptoms among the relatively few employees of the Life Sciences plant: those employees had served, inadvertently, in the way that miners once used canaries in small cages to detect life-threatening buildups of methane.

Put on notice, the state rapidly expanded the investigation. The state investigation found serious Kepone contamination of the James River estuary, a major estuarine tributary of the Chesapeake Bay and an important provider of recreational and commercial services in its own right. Reconstructing what happened, it seems likely that there were intermittent and illegal discharges of Kepone into the James for a period of about ten years. Because Kepone is relatively insoluble and nondegradable, much of that discharged Kepone is now resident in James bottom sediment. The portion that is in the water column is slowly being flushed out into the Bay in one way or another, so that there is persistent low-level contamination of the biota of both the James estuary and some portion of the Bay. But there is also an improbable catastrophe that should not be ignored: the sudden flushing into the Bay of substantial amounts of the Kepone currently resident in James sediment. Since at current normal flushing rate half of that Kepone will remain in James sediment in fifty years' time, that improbable catastrophe will be our legacy to at least the next two generations.

Unless, of course, something is done to remove the Kepone. The possibilities for remedial action were in fact carefully enumerated and examined by a major environmental postmortem of the accident. The saddest aspect of the conclusions of that study is their uniformly high cost and uncertain benefit. The costs of ex post mitigation of the incident run into the billions of dollars. And it is not clear that even incurring those costs will improve the situation, as opposed to moving the Kepone to another place and environmental medium. For example, one possibility considered was dredging the Kepone-contaminated portions of the James estuary bottom. But dredging operations will inevitably disperse much of that bottom-sediment Kepone into the James. And if the dredged sediment is placed in a landfill, there is the problem of groundwater contamination to be reckoned with.

Thus, ex post, Kepone contamination of the James seems to have been a very bad bargain. If this generation chooses to undertake the uncertain and costly remediation effort, the ex post accounting requires subtracting a figure running into the billions of dollars from what, in principle, would have been the opportunity costs of the resources required by the Life Sciences plant to avoid the incident.

About those latter costs there will always be a penumbra of uncertainty. In the criminal prosecutions that followed the incident, internal Life Science memoranda came to light in which corporate executive officers mentioned figures in the hundreds of thousands of dollars as the cost of certain kinds of abatement equipment. In a more perfectly competitive economy, such estimates could be taken almost at face value. Even if they are, the Kepone incident transaction seems, ex post, to have been a spectacularly bad one for this generation. If this generation chooses, for whatever reasons, not to mitigate the situation in the James, we leave to at least the next two generations a seriously damaged version of the estuarine system our own generation inherited.

And, of course, we leave to those two generations the remote possibility of a catastrophe affecting the Chesapeake Bay. For the Bay is a unique natural asset, a productive fin and shellfishery that may contribute to the productivity of fisheries as far north as the Georges Bank fishery off New England. At some remove, the argument that the James estuary has close substitutes and can rationally be depreciated and scrapped may appeal to some; the same argument is unlikely to be made, in any serious way, about the Bay.

Contamination of Groundwater by Price's Pit

During the 1960s, Charles Price, the operator of a landfill located about one mile east of the water supply well field for Atlantic City, New Jersey, illegally accepted for disposal tons of liquid waste containing hundreds of toxic chemicals and hazardous metals. The location of the site and the geohydrology were about as bad as could be: only ten feet of porous soil separated the bottom of the pit from the Upper Cohansey Aquifer, which flows roughly eastward from the site toward the Atlantic City wellfield.

The Price's Pit site now routinely appears on lists of the country's worst hazardous waste site problems and has been extensively studied by hydrologists and public health officials. Thus, it helps us understand more clearly what kind of a problem we have in this, and unfortunately in many other, sites.

For even though groundwater moves relatively slowly and can, at some cost, be monitored for contaminants, it is difficult to establish just what chemicals, and in what concentrations, the population of Atlantic City has been exposed to. This is because there are several hundred chemicals involved, so that the concentrations of only about 120 could be measured. And because drilling wells and running field chemical analyses accurate to the parts per billion level are expensive and uncertain undertakings, the information that we actually have is far from what would be required to make exposure estimates with some confidence. We might call this the large numbers problem: there are what may be an unmanageably large number of chemicals running around.

And even the interpretation of the data that we do have on particular chemicals is complicated by what might be called the geohydrological facts of life. A real world aquifer often has irregular boundaries and is highly

heterogeneous. Real world hazardous wastes often infiltrate such an aquifer in lumps, or "slugs," and not in the steady flows that are easiest to model. Thus, what is really happening in the Cohansey Aquifer below Price's Pit may be very complicated: numerous slugs of varying concentrations may be moving through the aquifer, to be detected only when they pass the screens of observation wells or--sometimes years too late--residential or public water supply wells.

Something like that happened at Price's Pit. About thirty residential supply wells were located close to the landfill, standing between the landfill and the public supply wells. Contaminants from the landfill almost certainly reached those wells years ago. Again, the analogy with the nineteenth century miners' canaries seems apt. Had those wells not been there, recognition of the incident might have followed serious exposures of a population of 40,000 individuals to many toxic chemicals and heavy metals.

But a careful analysis of what we stood to gain from avoiding the Price's Pit incident must, in some way, consider all these possibilities: of recognition, of transport, of exposure.

The Fire at Chemical Control

In the early morning hours of April 21, 1980, the Chemical Control site in Elizabeth, New Jersey, caught fire. The site had already been closed by the State Department of Environmental Conservation for violations of state laws. Substantial amounts of hazardous materials of various kinds had been removed from the site between the date of closure and the morning of the fire.

Nevertheless, many thousands of drums of hazardous chemicals remained on the site, and were there on the morning the fire began. That fire lasted for ten hours and produced pictures of drums lifted tens of meters into the air after exploding. Most seriously, perhaps, the fire sent a plume of toxic smoke into the air. That plume could have been transported over the nearby and heavily populated areas of the borough of Staten Island in New York.

But transport was in fact restricted by favorable meteorology: the gods, or at least the winds, were with us. During the course of the fire, wind speeds and directions were such that toxic fumes were not carried over heavily populated surrounding areas. The actual, or ex post, health damages were probably small. But the ex ante damages may be significantly larger: had the fire occurred while wind speeds were higher and wind directions pointed toward densely populated areas, things could have been worse. Since it is ex ante damages that are relevant in thinking about what we are willing to do to avoid a Chemical Control-type incident, we are forced to think about how to construct measures of ex ante damage. In some way we have to take account of the full range of possible meteorological conditions, and be on our guard against averaging out extreme conditions too early in our accounting. For by excluding even less probable extreme events, we may very well be throwing away large associated damages, and thereby biasing our estimates.

This is not just a theoretical issue: it was just this undue haste to average over meteorological conditions that was one of the principal targets of critics of the Reactor Safety Study. For that study presented a computation of the damages associated with a reactor accident at a representative site. That representative site was constructed early in the game by averaging over meteorological conditions at the approximately sixty American commercial nuclear power reactor sites. But the actual risk to populations from the nation's power reactors is dominated by the risk at two particular sites, Indian Point in New York State and Zion in Illinois. Both sit just north of large cities: Indian Point north of New York City, and Zion north of Chicago. The coincidence of a major release of radioactivity and southerly winds would, at either site, lead to huge losses of life and significant damages to property. By averaging too early, those losses were inadvertently eliminated by the Reactor Safety Study analysts.

MAJOR EPISODES: DILEMMAS OF EVALUATION AND POLICY

The four episodes described above could be supplemented by many others, but those four raise many of the dilemmas such episodes pose for society and for efforts to mitigate or avoid such episodes--and efforts to evaluate proposals toward those ends. The Three Mile Island episode reminds us of the reality of the question of psychic cost or of the cost of bearing what we have called anxiety. That same incident at the same time tells us that it can be difficult, both conceptually and practically, to talk in terms at once sensible and probabilistic about the events triggering such episodes.

The Kepone episode and the Chemical Control episode both force on our attention the necessity of thinking about extreme events which, if not unthinkable, are at least relatively improbable. In the Kepone incident, the worrisome extreme event is sudden transport of significant amounts of bottom-sediment Kepone into the Chesapeake Bay, with unpredictable consequences for that unique and imperfectly-understood natural asset. Similarly, in the Chemical Control incident, we suspect that we will have to examine the full range of meteorological conditions that might have accompanied the fire, and that the early averaging of extreme meteorologies may be a serious mistake.

From the Kepone episode, we may also draw the moral of what we might call ex post cost asymmetry: the ex post costs of prevention seem to have been small relative to the ex post costs of remediation, and perhaps relative to the ex post damages, even if no remediation is undertaken. And finally, from the Price's Pit incident, we abstract a feature we might call the many substances problem: there are many chemical contaminants in the Upper Cohansey aquifer under Price's Pit, and keeping track of those many potential "bad actors" is itself difficult and costly. The number of those actors is, of course, another one of the legacies of our period of chemical inventiveness. How do we keep track of, and govern, those many chemicals once they have contaminated an aquifer or, for that matter, even before such contamination?

As we contemplate changes in our policies for managing hazardous substances, those issues will receive serious attention. They are

sufficiently novel, and difficult, to be labeled characteristics of what we have called major episodes. We attach no enormous importance to that choice of terminology; some or all of these features are shared by what we have become accustomed to thinking of as routine, or ordinary, pollution. But we believe that choosing and using the term major episodes is useful in the way that caricature can be: it focuses attention on those special features.

Further, in the forthcoming discussions of hazardous substance policy, it is far from clear that one disciplinary perspective should rule the roost. But our own perspective is that of economics leavened by a bit of cognitive psychology and natural systems modeling. And because the decisions to come in hazardous substance policy, possibly including the proverbial decision to do nothing, will have important consequences for the way we allocate our scarce human capital and environmental resources, that perspective must be present.

THE ECONOMICS OF MAJOR EPISODES: A FOCUS

Costs and Benefits

Perhaps the first question a citizen might ask about the episodes we have described, or about any of the many similar episodes, is: are those episodes necessary? And perhaps the simplest rendering of that question into economics, which abhors the notion of absolute necessity and substitutes the notion of cost, is: what would it cost to avoid such episodes, or to mitigate their consequences after they occur?

Surprisingly, answering such questions probably poses no serious conceptual problems, even though the practical problems may be severe. Marching through the illustrative incidents we have chosen may be the best way to make the point. For the Three Mile Island accident, an upper bound on that cost estimate is given by the cost differential, over and above nuclear electricity, for base load power from a plant not susceptible to a core meltdown. And there are representatives of that latter class: clean (but expensive) natural gas generation, or perhaps even nuclear electricity from a gas-cooled reactor. For the Kepone incident, the relevant cost is the cost of avoiding occupational exposures and releases to the environment: but such costs are routinely generated by cost-estimation specialists working for engineering firms. For the Price's Pit episode, what matters is the differential between the costs of "safe disposal of hazardous waste"--by incineration or disposal in securely-lined landfills--and the operating and capital costs of an operation like Price's Pit. And for the Chemical Control episode, we might look at the difference between the costs of transportation to, and disposal at, a site far from any populated area, and the costs of disposal at the actual Chemical Control site.

There is no great difficulty in producing such cost estimates, and many have been produced recently: more estimates are sure to be forthcoming as we learn more about safe land disposal technologies, about incineration, and about process changes and abatement technologies. But if we had all those present and future cost estimates in hand in some ultimate cost-engineering book of blueprints and costs, we would still have to ask ourselves an

important question: what are we willing to pay to avoid those episodes? For it seems likely that, here as elsewhere, we can get what we pay for. If by incurring greater costs we can further reduce the likelihood and severity of major episodes, how much do we want to spend out of scarce resources with alternative uses? Put in other words, what are the benefits of reducing the likelihood, and the severity, of prospective major environmental episodes?

The Cost-Benefit Calculus

Under certain highly idealized conditions, competitive market prices can guide consumers and producers in allocating the scarce resources under their control, and guide them so that everyone benefits. That this can happen was of course recognized early, and publicized by Adam Smith. Much of the development of theoretical economics since has been devoted to demarcating the situations in which those idealized conditions do not hold. When they do not, there is sometimes a presumption that the government must or should play a role in correcting that so-called market failure.

But, granted the presumption that the government should do something, the question remains: what should it do? The subdiscipline of economics called public economics is, at least on the expenditure side, largely concerned with formulating rules for government action. Invariably, those rules are aimed at "removing the market failure" by generating the information necessary for the government to act in a "market-like" way.

Cost-benefit analysis is one kind of public economics calculus for guiding government actions in cases of market failure. It came of age as economists wrestled with water resource project analysis issues. In principle, little could be simpler than the rule for project evaluation. If a project has been proposed which provides, say, both market and nonmarket benefits and costs, then try to do what the market would. Estimate both nonmarket costs and benefits, and then apply the decision rule: proceed with the project if total benefits exceed total costs. If they do not, cancel the project.

This is about as simple as the Golden Rule, and may have lead to about as much confusion in practice. It has certainly generated an enormous body of academic disputation that may be only loosely related to the implications of cost-benefit analysis in practice. Some have conjured up Burke's image of a group of heartless "sophistikers, economists, and calculators," siting plants or damming free-flowing rivers on the turn of a decimal point. Charles Schultze, for example, believes that the main impact of cost-benefit analysis in water resource project evaluation has been to help kill projects that were unambiguous and dramatic losers. A similar modest success in the hazardous materials policy area is a worthy ambition. But in any event, the comparison of costs and benefits in that area is inevitable. The work of cost-benefit theorists and practitioners over the past decades can help guide those comparisons. And more recent work by scholars in several areas can, as we shall see, help shape those methods to the particular and intriguing requirements of the major episodes.

Cost-Benefit Analysis and Major Episodes: First Approach

Imagine, then, that we are charged with writing a manual for cost-benefit analysts of prospective facilities, or siting arrangements, which may conceivably lead to major episode. We may, for example, be siting a hazardous waste facility, or we may be locating a chemical plant on a major estuary. Imagine further that the manual must be written quickly, using well-developed and well-understood methods, and that there is no time for serious reflection particularly about the special characteristics of major episodes that we have listed. What do we do?

H. L. Mencken once said that for every question, economists have an answer: one that is "simple, neat, and wrong." We prefer the formulation "simple, neat, and conventional." In developing that answer, begin with the project analogy. A policy aimed at preventing, or mitigating, a major environmental episode is considered a project, and therefore a fitting subject for cost-benefit analysis. The alternatives are eliminate, or do not eliminate, the possibility of an episode. The benefit associated with eliminating the possibility of an episode is precisely the eliminated cost of bearing the risk associated with that episode.

Thus, we can apply standard cost-benefit methods to the major episode case if we have a way of estimating the cost of bearing the risks associated with an episode. And we do have a more or less conventional way to make that estimate. We simply apply the standard theory of the cost of risk-bearing that has evolved over the last few decades. Let us suppose, for illustrative purposes, that some prospective episode would impose damages largely in the form of damages to human health, and proceed from that supposition to a method for estimating the cost of bearing the episode risk. There are two steps in that exercise. The first consists of what should be called risk assessment, or probability distribution estimation. The second consists of an exercise in valuing the range of consequences associated with a major episode.

The two steps are logically independent, and in fact require very different kinds of skills. The first step, risk assessment, might be undertaken by reliability engineers and natural systems modelers. If, for example, we are considering episodes associated with groundwater contamination from landfills, that risk assessment group would be charged, in principle, with producing a probability distribution on population exposures to various levels of chemical contaminants in the years after the landfill operations begin. In principle, there is nothing remarkable about such a distribution. Though we have argued, from the Three Mile Island incident, that release probabilities may in practice be very difficult to construct, the notion of a probability distribution over exposure time sequences is, in itself, straightforward.

The same can be said of at least one approach to the second, or valuation step. We will shortly see that the approach itself is deserving of scrutiny, but to see why we must think the logic behind that approach. We begin with a few radical, and radically simplifying, assumptions. We assume that individuals know, and act upon, the correct probability distribution

that our risk assessment team will be charged with estimating. We assume that individuals make rational choices under uncertainty, and define rationality with a set of axioms similar to those that are the basis for the modern theory of consumer choice under certainty. We further assume that individuals treat certain risk categories as homogeneous commodities, and in particular that risks to human health encountered on the job are treated as homogeneous with risks to health from exposures to environmental pollutants in the ambient environmental media. And we further assume that individuals value collective risk much as they value individual risk.

Bringing those assumptions together allows us a simple, neat set of instructions to our valuation team. To wit: take estimated risks to life, and then value those risks at the prices revealed by individuals as they choose among risky jobs in the labor market. We will, in fact, carry out this program for the cases we shall study. But the radical simplifying assumptions underlying such estimates call for some closer examination, particularly when we recall the distinguishing characteristics of our major episodes. If risk assessment is a difficult task for a group of highly-trained reliability engineers, and if natural systems modeling under extreme conditions is the purview of scientists working with state-of-the-art supercomputers, how are individuals supposed to derive good estimates of those probability distributions? Were we to take the trouble to test, in the laboratory, axioms defining rational behavior under uncertainty, will people behave that way?

Do individuals treat all risks, or even all risks to life, as amounts of some homogeneous quantity or do they value specific kinds of risks to life in specific ways? And, finally, do individuals value individual and collective risks differently?

Cost-Benefit Analysis and Major Episodes: The Dilemma

Those questions have been raised frequently, and forcefully, in recent years, sometimes implicitly by noneconomists and, less frequently, by economists with nagging doubts about the application of standard methods to what seem to be some very nonstandard cases. If things run true to form in economics, where good and bad ideas alike can survive over the centuries and real refutation is rare, those questions will be with us for some time. But at least one seems to have special claims on our attention and effort here. We restate the question: is it plausible, in thinking about major episodes, to imagine that individuals value those episodes on the basis of true probabilities, and in the same ways that they value other more familiar risks?

To pose the question is almost to answer it: it is not. But that answer alone is insufficient, for Stigler's Law says that even a bad (economic) theory is better than nothing, or seems better. In working our way towards an alternative to a bad theory, it may help to recall the bases for the familiar, but implausible, theory.

The modern economic theory of choice under uncertainty is largely the result of the efforts of economists working in the 1960's and 1970's, but it

had its origins in the eighteenth century. The critical stimuli, as so often in the history of science, were certain observed and seemingly paradoxical facts of life. The same individual could be observed gambling, or courting risk, and buying insurance, or avoiding risk: why? And most individuals will not pay an arbitrarily large sum to play a game in which they win 2^n dollars if the first head appears on the n^{th} toss of a fair coin, even though the expected value of that game is infinitely large: this is the famous St. Petersburg paradox.

The decisive resolution of these paradoxes was given by Daniel Bernoulli, who for the first time noticed, and formalized, the phenomenon of risk aversion: individuals are averse to bearing risks, and will pay not to bear particular risks. What they are willing to pay is in effect the cost, to them, of bearing the associated risk.

But Bernoulli simply assumed that individuals are risk averse. Such assumptions are incompatible with the temper of modern economic theory, which has taken the axiomatic method of pure mathematics as a model for the derivation of behavioral rules. And in many real-life decisions under uncertainty things are far less tidy than, say, in gambling. For the odds of a certain roll of a fair die are known, but in real-life choices under uncertainty there is often no actuarial or a priori basis for probability estimates.

It is a measure of the achievement of Leonard Savage that his book The Foundations of Statistics filled, in the 1950's, the gaps left by Bernoulli in the eighteenth century. Savage begins with axioms defining rational individual choice under uncertainty. From those axioms he derives a representation of a rational individual's decision rule, and in that representation there appear weighting factors having all the mathematical properties associated with a probability distribution. The terms those factors weight are exactly Bernoulli's utility terms. The theory of choice under uncertainty has been closed; probabilities appear as convenient summaries of individual beliefs about uncertain events, and they weight individual utilities.

Savage's achievement is powerful and elegant. It can fairly be said that his version of expected utility theory is the dominant influence in the modern economic theory of choice under uncertainty, and thus that his influence permeates our understanding of insurance and labor markets, of corporate finance and the stock market. In fact it is everywhere that modern economists reason about the choices rational agents make when facing uncertainty. And since the normative appeal of Savage's axioms is enormous, this is as it should be.

But there is a problem with the theory: it lies on the positive, or descriptive side. In the simplest terms, individuals don't seem to act the way Savage's theory says they should. This was pointed out soon after the publication of Savage's book by the French economist Maurice Allais, and the results of his experiments go by the name of the Allais paradox. It has been repeatedly confirmed by laboratory experiments with human subjects over the past three decades, and the evidence is impressive. Individuals violate the

Savage axioms, and they violate them in ways that are both routine and systematic. The issue is not one of random deviations from a correct behavioral theory: the behavioral theory is wrong.

We are a long way from understanding what to put in its place. Daniel Kahneman and Amos Tversky, who have studied these issues for several decades, have come as far as a proposal that individuals use several rules of thumb in making choices under uncertainty, rules with some basis in previous experience that are then applied in making decisions. Some of those rules are compatible with the notion that some meaning can be attached to individual risk valuation, but others are not. And attempts, typically based upon models of cognitive processes proposed and developed by psychologists, to put those rules of thumb on a firmer theoretical foundation are in their infancy.

The dilemma this situation poses for practical exercises in the cost and benefit analyses of major episodes can now be seen. If an analyst decides to accept the Savage theory as a positive theory, and decides to assume that individuals act as the theory suggests with more or less correct probabilities, he or she is on thin ice. For the episodic situations are exactly those for which experience does not provide a good basis for the formation of probabilistic judgements.

If, on the other hand, that analyst rejects any application of the Savage theory as a positive theory, and in fact embraces the list of rules of thumb as descriptively, then he or she is left with no basis for valuation whatsoever. In fact the conclusion is that there is no basis for application of the cost-benefit calculus to major episodes, because there is no such thing as a consistent individual valuation of the cost of bearing the risk of such episodes.

Cost-Benefit Analysis and Major Episodes: Some Compromises

We regard this dilemma as real and inescapable; it should be the focus, over the next ten years or so, of major cooperative research efforts involving cognitive psychologists and economic theorists. But the resolution of many of our outstanding issues in hazardous substance policy cannot and will not wait for the success of that enterprise. And we are reluctant to render a judgment that, in the interim, the cost-benefit calculus has nothing to contribute.

Thus the following guidelines seem attractive. For starters, we do perform our analyses of the costs of bearing the risks associated with major episodes under the unpersuasive and highly idealized assumptions most naturally compatible with expected utility theory. Those assumptions do support a consistent notion of individual valuation of the cost of risk bearing, and moreover they almost certainly give conservative lower bounds on that cost. Good lower bounds can convey useful information.

Eventually, we hope to go beyond that framework and perform some auxiliary estimates of the cost of risk bearing. Such estimates can be based upon a few assumptions about the way in which individuals react to, and act

after, a major episode. Specifically, we retain the Savage expected utility formalism with its normative implications. But we enrich its descriptive power by assuming that individuals react to a major episode by increasing their estimates of the probability of such an episode. Over time, those probability estimates can be revised gradually downward, if further major episodic events do not occur. That revision may be slow or rapid, depending upon the time period the individual chooses for revision. But the following argument suggests that, in practice, it may be quite protracted. Consider the situation of an individual living near the Three Mile Island site. He or she may or may not have heard of the Reactor Safety Study bottomline estimate of one major reactor reactor accident in one hundred thousand reactor-years. But like the billion dollar unit in which deficits are measured, that number is so small that it is not internalized: it is cognitively negligible.

Suddenly that impossible rare event occurs. The event, or episode, is now cognitively plausible, and is effectively awarded a subjective probability comparable to that awarded other, more familiar hazards. Almost everyone in the United States is aware that driving is potentially hazardous to their health, for example. Newly-recognized dangers may then be assumed about as dangerous as familiar ones, at least until additional experience forces revision of that judgement.

But for rare events, the relevant measure of experience may be long in coming. The residents of the Harrisburg, Pennsylvania area around Three Mile Island are, for example, bringing suit in the Supreme Court of the United States. That suit aims at forcing the Nuclear Regulatory Commission to prepare an Environmental Impact Statement for the startup of Three Mile Island Unit 2, the reactor and plant that were not involved in the accident. Suppose that in fact there is no reason to believe that Unit 2 is any riskier than any other operating American nuclear power reactor. Are the plaintiffs, justified, or rational, in clouding their judgements about Unit 2 with their experience gained during the Unit 1 accident? Alternatively, suppose that they have, after the Unit 1 accident, revised their accident probability estimates sharply upward. Are they simply applying those estimates to Unit 2, and is that application justified? And if so, should the Court acknowledge that those beliefs, or anxieties, deserve administrative and regulatory recognition, and require the Nuclear Regulatory Commission to take them into account?

We have no simple answers to those troublesome questions. But the scheme we propose, in appendix E, is both testable and estimable, in principle. And we believe that, since this issue is so critical to our social and political judgements about policies aimed at making major episodes less likely, it should be explored.

PLAN OF THE BOOK

Now let us look forward over the next eight chapters. Chapter 2 is devoted to the problem of valuing major environmental episodes. It begins by taking up the conceptual problem of defining a willingness to pay to avoid bearing a particular risk. Because there is a simple and more or less standard approach to that problem, we begin with an exposition of that

approach. But because at least some of the assumptions underlying that approach are questionable in the major episode case, a good deal of care is taken with those assumptions. The reader is warned that there are stretches of chapter 2 that are somewhat inconclusive; that, unfortunately, is true of the state of parts of this art.

The next six chapters, chapters 3 through 8, are devoted to the three case studies. Our description and analysis of each case study extends over two chapters. The first describes the episode, and provides an overview of the problems of modeling environmental transport for that particular episode. The second chapter then carries through the actual analysis and computation of the costs of bearing the risk associated with the particular episode.

In a ninth, and final, chapter, we bring together some conclusions drawn from the three case studies taken together, and provide the conventional statement of further research needs. Moreover, we offer a tentative and very preliminary set of reflections on what we call hazardous substance policy. The reader is warned that those suggestions cannot pretend to have any firm statistical basis in the three cases we have examined. Three is, by almost any measure, far too small a sample from what seems to be a rapidly expanding universe of major environmental episodes associated with toxic substances. Nevertheless, on the grounds that there has probably been a good deal of what the cognitive psychologists call "general learning" during this project, those ruminations on policy are offered, clearly marked for what they are.

Finally, a word about the appendices. This project has embroiled the author in more kinds of data than he cares to, or can, remember. Not all of those kinds of data were ultimately used. On the chance that even those data may be of some use to other scholars working on these problems, even they have been described, in three data appendices, one for each case study. And in addition to the data appendices, there are two others. One of those lists the incidents we considered, but ultimately rejected, as case studies. The other, parts of which are described in the research needs sections of chapter 9, deals with the valuation of what we call anxiety effects.

CHAPTER 2

ASSESSING AND VALUING RISKS POSED BY MAJOR ENVIRONMENTAL EPISODES: A METHODOLOGICAL OVERVIEW

INTRODUCTION

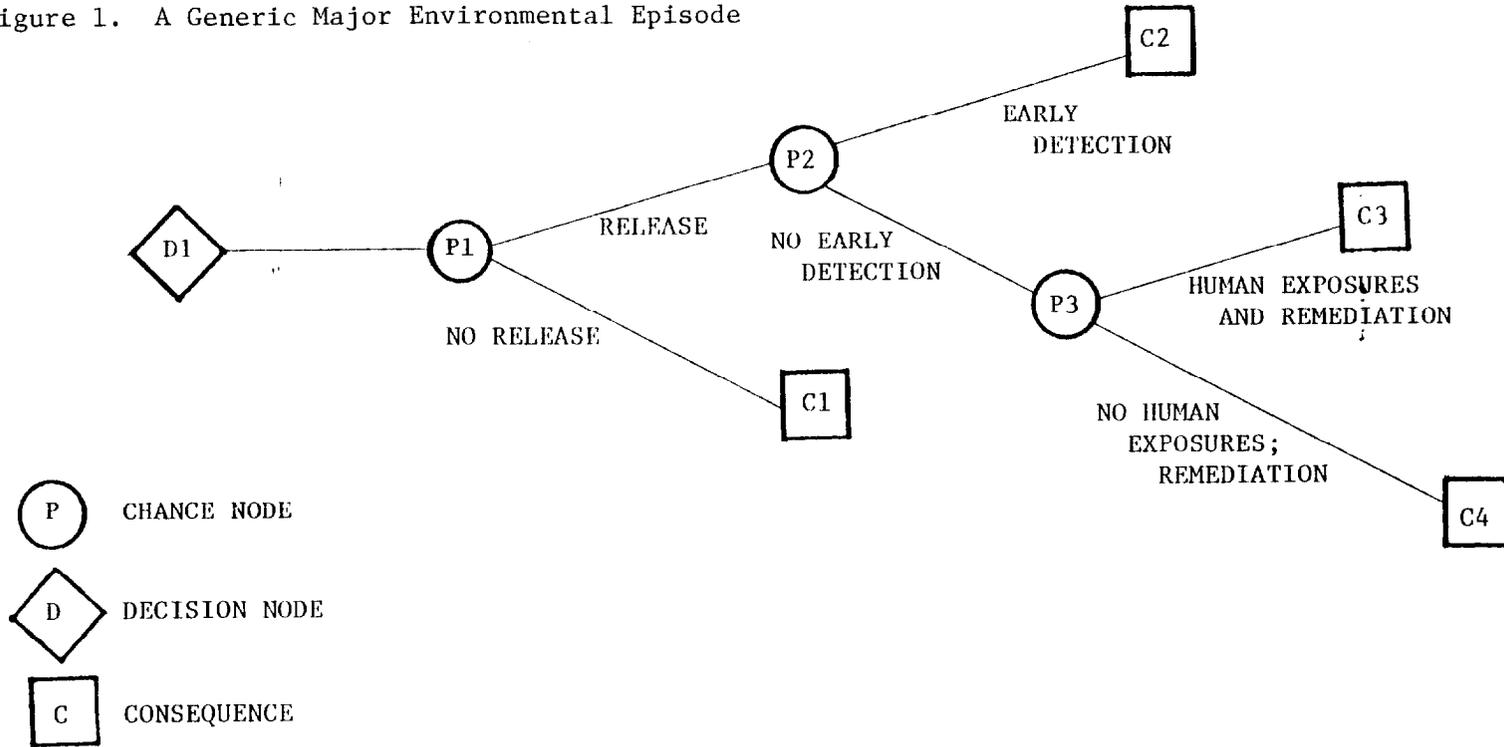
Let us begin by setting out what we will call the standard method for valuing major environmental episodes. Imagine that a community is debating whether to accept the siting of a facility which may be the cause of what we have called a major episode. That facility may be a chemical plant located upstream of the community's water intakes; it may be a toxic waste disposal site located above the community's drinking water or downwind from the community under prevailing wind conditions. What these examples share is easy to describe: the possibility of releases from the facility being transported by some ambient environmental medium, and resulting in environmental damages or human exposures.

Figure 1 below is a generic representation of a figure we shall see for each of our case study episodes. The initial decision node D1 represents an explicit decision to site the facility, with the decision to site "voluntary" in the sense that a decision not to site could, in principle, have been made. Acceptance of that lottery implies acceptance of what may happen as a result. The chance node P1 corresponds to the possibility of a release of a hazardous substance from the plant. If no such release occurs, the innocuous consequence C1 stands for "no losses." But the other branch of that subplottery corresponds to a worse outcome. There may be human exposures, and human health effects, from the initial release: see consequence C2.

Presumably a large, unscheduled release could not escape detection; a small, steady release might, on the other hand, go undetected for many years. As we shall see, the relationship between releases and detection, and the translation of that relationship into this or any other calculus of risk is subtle. For now, suppose that small, steady releases go unnoticed for many years. If the released materials do not persist in the environment, then a relatively innocuous consequence C3 may be the result. But if those materials are persistent, there is the possibility of steady, low-level human exposures, and the possibility of secondary releases from the stock of pollutant, resulting in "delayed" human exposures--and damages to the service flows provided by the environmental medium.

Now let us suppose that the community making the facility siting decision is a community of one, the economist's traditional device for separating the problem of individual valuation from the aggregation problem.

Figure 1. A Generic Major Environmental Episode



D1: Lottery accepted; landfill sited

P1: Release subplottery

C1: No release from landfill

P2: Early detection subplottery

C2: Early remediation (or substitution of alternative water supplies)

P3: Human exposures subplottery

C3: Human exposures and remediation

C4: No human exposures

Further suppose that the community accepts the approach to decisions under uncertainty associated with the classic work of Leonard Savage, an approach that is the basis for almost all of the "economics of uncertainty." What is the cost to that individual of bearing the risk associated with the facility? Put in other words, were the individual asked by some higher facility-siting authority to submit an honest estimate of how much the individual must be compensated for bearing the risk of the facility, what figure would the individual submit?

Here is the standard method for answering that question. Suppose that our community understands, from figure 1 or something similar, that there are three kinds of uncertainties involved: those surrounding an initial release, those bearing on subsequent environmental transport, and those determining consequent environmental and health effects. In the state-of-nature approach to uncertainty, the state of nature s is a triple (s_R, s_T, s_E) , with components corresponding to release, transport, and exposures.

Then what is required is a combination of "risk assessment" and "valuation" exercises. What has come to be called a risk assessment is nothing more than the distillation of vague beliefs into a probability distribution (s_R, s_T, s_E) on the states of nature. That task is made easier because the problem naturally decomposes into a set of somewhat independent exercises requiring very different kinds of skills. It is conventional to suppose that the random variables governing releases, transport, and the consequences of exposure are independent, so that the distribution factors:

$$\pi(s_R, s_T, s_E) = \pi_R(s_R)\pi_T(s_T)\pi_E(s_E) \quad (2.1)$$

In this case, safety experts can be assigned the job of estimating $\pi_R(s_R)$, natural systems modelers of the task of estimating $\pi_T(s_T)$ and environmental health specialists can work on $\pi_E(s_E)$.

Given a risk assessment of this kind, the valuation exercise is, in principle, routine. Our community knows its own risk preferences, summarized by initial wealth parameter W_0 and a utility function u . If the monetary loss suffered by the community when SE occurs is $L(s_E)$, then the community should be willing to pay V to avoid bearing the lottery associated with the facility, where V is implicitly defined by:

$$u(W_0 - V) = \sum_{s_R, s_T, s_E} \pi(s_R, s_T, s_E)u(W_0 - L(s_E)) \quad (2.2)$$

Equation (2.2) in effect organizes the rest of this chapter. The next section, Estimating Source (or Release) Probabilities, treats the release, or source, probability estimation problem. In the following section, Environmental Transport Uncertainty, we turn to the contribution to the overall episode lottery from uncertainties in the environmental transport media at the time the episode occurs and thereafter. The general idea will be familiar to anyone who has ever gambled on the weather in planning an outing or a trip. The returns to the enterprise will be very different in rain and shine, and in deciding whether to make the outing or trip, we must

form some judgment of the relative likelihood of rain or shine. In the section Exposure Estimates, we take up the next of the major sources of uncertainty: the linkage between ambient environmental concentrations of pollutants and human exposures to those pollutants. That linkage is critical to our estimates of one of the most important damage categories, damages to human health. The section Dose-Response Estimates takes up the uncertain relationship between exposures and health effects, sometimes called the dose-response relationship.

Thus far, all of those relationships are the proper subject matter of the natural and engineering sciences: reliability engineers have much to tell us about release probabilities, natural systems modelers about translating source terms into ambient concentrations, and physiologists and epidemiologists about the relationship between exposures and health effects. But in the end, we are interested in dollar valuations of those damages, so that they can be compared however roughly, against estimates of the costs of avoiding major episodes. That raises the valuation issue: how do we translate our damage categories, such as environmental degradation and human health effects, into dollar losses? That issue is one on which reasonable persons can and do differ, sometimes vehemently. While some of those differences are philosophical, they have a firm basis in the serious conceptual problems still plaguing the so-called value of life literature. Those problems, and the interpretation we have adopted of them for present purposes, are discussed in the section The Valuation Problem.

For the most part, the discussion in this chapter is quite general, leavened by references to the case studies. The case study chapters are the place to look to see what was actually done in each case. But there is some purpose in seeing these episodes in a more general light, and in some comparative perspective.

ESTIMATING SOURCE (OR RELEASE) PROBABILITIES

The initiating event for our generic case study episode is an extraordinary release of some pollutants into the environment. It is obvious that no release implies no major episode, and equally obvious that estimates of the value of policies aimed at avoiding such episodes will be sensitive to estimates of release probabilities. And it is almost as obvious that estimating those probabilities is difficult, precisely because the rarity of the events in question means that there is little or no actuarial base on which to build.

In the end then, we are thrown back on either treating release probabilities as parametric, or on strong assumptions about the underlying probability model for our purposes of risk assessment. The first strategy seems easy; the second clearly requires some disciplined presentation and judgment. In that presentation, we follow, at least for a time, Pratt and Zeckhauser (1982).

Should Release Probabilities Be Taken as Parametric?

In chapter 1, we argued from the Three Mile Island accident record that good estimates of failure probabilities for rare events are extremely difficult to produce. Returning to the lottery depicted in figure 2.1, how will that difficulty be reflected in discussion of public policies aimed at preventing, or mitigating, major episodes? One possibility is treatment of release probabilities as the undetermined residual. The kind of question suggested by the apparatus of figure 1 would then be: how high would the probability of a major episode at a prospective site have to be to justify the incurrence of a specified level of avoidance or mitigation costs? The work product of the kind of exercise suggested by figure 1 would be a function giving the dependence of those critical probabilities on avoidance, or mitigation, costs.

This is one kind of practical argument for treating release probabilities as parametric. But there is another, perhaps more compelling, argument. It is far from clear that the probabilities relevant to individual risk perceptions are the true release probabilities: to the contrary, there are powerful reasons for believing that some cognitively-determined probabilities, and not the ones that a competent risk-assessment team would produce as estimates, are what matters.

If that is the case, then our analyses should be done in such a way that the sensitivities of all results to release probabilities are easy to explore, and easy to understand.

The Event Space for Release Probabilities

Return to the innocuous-looking expression $\pi_R(s_R)$ that occurs as a factor in equation (2.1). Since by definition knowledge of the state of nature (for releases) s_R removes all uncertainty with respect to those releases, the literal interpretation of that expression is: our subjective probability distribution over releases. In principle, those individuals bearing the risk of an episode should, according to the normative theory of choice under uncertainty, summarize their beliefs about plant hazards by constructing such a distribution. In a world of perfect information and more or less trustworthy experts, a risk-assessment team might prepare estimates of that distribution, and those estimates would find general acceptance among the population at risk.

But in the world we actually live in, things are not so simple. The strictly technical task of release-probability assessment is challenging enough: for a large and complicated facility like a civilian nuclear power plant, it can push our current computational abilities to their limits. But even worse, the relevant release probabilities depend, in some intricate ways, upon the relationship between the incentives facing the operators of these technologies and upon the technologies themselves. For that reason alone, plausible subjective release probability estimates will be hard to come by.

We have already illustrated this argument in our discussion, in chapter 1, of the Three Mile Island accident. There we suggested that a good failure model of the plant, the key constituent of a good release-probability model, would necessarily build upon a good cognitive model of operator behavior under accident conditions. Take the argument one step further: upon what principles should that cognitive model be built? Certainly, in part, it should be built upon psychological principles drawn from the human-factors studies of the last decade. But such general principles alone are unlikely to be sufficient. There is good reason to believe, for example, that operator performance in more or less identical nuclear plants operating in different countries varies across those countries. The explanation almost certainly lies in the different recruitment and incentive systems deployed by the different countries. Thus, the incentive systems cannot be ignored in modeling release probabilities. But it is far from clear how to take them into account.

The situation is broadly similar for the case study incidents that we examine below. In each case study setting, it seems plausible that decent subjective release probability estimates must build upon substantially more knowledge of the governing incentive systems, and their relationship to the technologies involved, than we have now. In the Kepone incident, there seems to be little doubt that the Life Sciences management was knowledgeable about the releases of Kepone into the James. For the sake of the argument, grant this premise. For the same reason, grant the further premise that management made its decision not to halt operation, not to install abatement equipment, and not to make process and configuration changes to limit occupational exposures on "rational" bases. Here "rational" means rational in the sense of Gary Becker's criminals, who maximize expected utility. They do so by committing crimes which pass the ex ante test: given the probabilities of detection, apprehension, conviction, and payoff, the "commit" lottery increases expected utility. No moral judgment beclouds their calculations.

Disturbing as it is in some respects, Becker's theory has substantial plausibility for many run-of-the-mill crimes in the United States in the 1980s, and seems to have substantial explanatory power. And for crimes like auto theft, the crime is so frequent, and the data on detection, apprehension, conviction, and payoff so abundant and accessible, that almost anyone can do the required calculations.

But if we seek to apply the Becker model to the Life Sciences management, again for the sake of argument, what do we learn? Even if we accept its general applicability, its explanatory power is sharply limited by the very different kinds of probability estimates involved. For now the relevant probability estimates will be the managements' estimates of the detection, conviction, and penalty odds they face. Beyond a few generalities about the relative ease of white collar, as opposed to blue collar crime, we know very little about these things. In any event, what will really matter are probability estimates at one further remove: the estimates of those bearing the risk of the lottery of the probability estimates held by those contemplating imposing the risk of the episode lottery on the community.

Turn now to the Price's Pit episode. There seems to be little reason for euphemism here: Price rather systematically and willfully violated the conditions of his permit, and apparently almost every cannon of safe landfill operating practice. That he evaded the consequences of his behavior for so long must be interpreted by many as evidence that Price's own estimate of the probability of detection and conviction was close to zero. By perhaps imputing more rationality to Price than is justified, we can see where those calculations will lead a landfill operator who is "rational" in the sense that Becker's criminals are rational. Because groundwater moves so slowly, a landfill operator whose site is only a mile from the nearest drinking water wells will have more than ten years before contamination from his landfill is detected in drinking water. That operator, if unsuspected, can accept nonpermitted wastes now and virtually drop them into the aquifer. In that way, he collects the fees and avoids the cost of liners or other retentive devices. If he abandons the site before contamination is detected at any water supply well, he has in effect privately appropriated, and depreciated, a common proerty resource, the aquifer.

There is abundant evidence that something like this goes on, so that the probabilities of detection and conviction have, in the past, been judged as unimpressive. But that is only part of our problem of constructing release probabilities for landfill sources. The part is knowledge of what was accepted for disposal at the landfill. It seems plausible to argue, at least in the Price case, that what was accepted was limited only by supply: anything brought to the landfill would have been accepted by Price. But we don't know very much about what will be offered under those terms.

Finally, consider the Chemical Control incident. Here the parallels with the Price case are extensive: the major difficulties in constructing the probability distribution over releases arise from our poor knowledge of what was actually on the site. In the Chemical Control case, because the state was in the process of clearing the site when the fire occurred, we do know something, and we can use that information as a guide in constructing our distribution.

Inferring Extreme Event Probabilities from Occurrences: The Zeckhauser-Pratt Argument

The events that we have called major episodes all are extreme events in the sense of being low-probability events. Our perceptions of this may change as time goes on and more evidence accumulates. But from the vantage point of the present, each of our case study episodes seems extraordinarily bad, the result of impressive negligence or worse, abetted by breakdown of governance. And in some ways, each of those episodes contributed to public recognition of what was seen as a new class of hazard. This general situation, with variations, is familiar elsewhere: though floods of various magnitudes are far from infrequent, catastrophic flood episodes significantly increase public awareness of flood dangers.

Thus there may be a tendency in perceptions, and a temptation in analysis, to confound the first observed event of a particular class, or the most recent event of a particular class, with some recognition threshold.

The next step, too, is tempting, a reconstruction of the probability distribution of events from those perceived recognition thresholds.

Recently, John Pratt and Richard Zeckhauser (1982) have examined the bases for such an identification. Those authors are interested in demonstrating that when an extreme event occurs and draws attention to a class of events, identifying the observed event with the recognition threshold for the class leads to biased conclusions. Toward that end they formulate the following probability model for the process generating the class of events. Introduce notation as follows:

d	Damages from an event, or "episode"
d'	Magnitude of the first observed episode
t'	Time of observation of first episode
d_A	Recognition threshold for an episode
h	Poisson-process parameter for event occurrence, assumed uniform over time
$f(d)$	Probability density of event magnitudes
$F_{\geq}(d)$	Probability of an event of magnitude $\geq d$, conditional on occurrence

The model works as follows. There is some probability per "unit time" that an event in the class occurs, with the time unit a natural measure of the amount of "exposure to hazard." Conditional on the occurrence of an event in the class, there is some distribution of event magnitudes, with probability density $f(d)$ and right-tail distribution of mass $F_{\geq}(d)$. Finally, there is some recognition threshold d_A : events smaller than d_A go unnoticed, or unreported. The first noticed event arrives with magnitude d' , so that d' necessarily satisfies $d' > d_A$. The natural question then arises: what are we entitled to infer from that occurrence about the class of events?

Suppose that both h , the occurrence rate, and d_A , the recognition threshold, are unknown, but that the event magnitude distribution $f(d)$ (and thus also $F_{\geq}(d)$) is known. Then a standard (Bayesian) answer to the question is as follows. Begin with a prior distribution $f_0(h, d_A)$. Next construct the likelihood function for the "data," in this case t' and d' . Here that likelihood function is zero for d' less than d_A . For d' greater than or equal to d_A , it is the product of the probability of an episode greater than the threshold (per unit time around time t') times the probability that no event occurs before t' , or:

$$l(t', d' | h, d_A) = \begin{cases} hf(d')e^{-t'hF_{\geq}(d_A)} & \text{for } d' \geq d_A \\ 0 & \text{otherwise} \end{cases} \quad (2.3)$$

Then by Bayes' Rule, the posterior distribution $f_1(h, d_A)$ is given, up to a multiplicative constant, by

$$f_1(h, d_A) \sim \ell(t', d' | h, d_A) f_0(h, d_A) \quad (2.4)$$

From that posterior joint distribution the marginal posterior distributions of h and d_A can be computed. But those marginal distributions are just what we want: they tell us how frequent episodes are (the marginal distribution on h) and how high or low the recognition threshold is (the marginal distribution on d).

What has been gained by this formalization? Pratt and Zeckhauser are concerned to show that mistaken identification of d' and d_A imparts an upward bias to estimates of the hazard rate. As they point out, that result is obvious from the form of the likelihood function for their probability model: since the factors multiplying t' in the exponential are reduced, the likelihood function overweights high values of h , and thus so does the marginal distribution.

But the Pratt-Zeckhauser construction assumes that the distribution f is known. In fact it typically is not, so that it, too, must be treated as "parametric" in the inference problem for h and d_A . That leaves us with the question of whether there is any real gain over the mistake against which these authors warn. The answer is the usual one: the gain is one of explicitness. In the above scheme, the sensitivity of inferences to the underlying extreme-event distribution can be tested: in the case they warn against, this cannot be done.

Finally, note that in an intermediate case, where there are several observed extreme events, the distribution f can be chosen from some parametric family, and those parameters estimated, along with h and d_A .

ENVIRONMENTAL, TRANSPORT UNCERTAINTY

Two Kinds of Uncertainties: A Formal Similarity

Between releases of pollutants to the environment and damages to the environment and human health stand the transport processes of the ambient environmental media. We begin with an essential distinction between two kinds of uncertainty associated with transport and then show that, in some important sense, that difference makes very little difference in the method for making benefit estimates.

One of the ambient media we will consider is groundwater. Groundwater aquifers are extremely sluggish natural systems, with very low flow velocities and very long turnover times. In a sense those systems are, or can be thought of as, entirely deterministic. But that does not mean that we can predict, with great confidence, how contaminants injected into an aquifer will spread. The source of the predictive uncertainty is our imprecise knowledge of the parameters of the transport system--the geometry of the permeable layer of the aquifer, the transmissivity of the permeable layer,

and so on. Because information on those parameters is costly, there generally will be substantial residual uncertainty even after extensive site surveys. But we will see that this formal similarity, while intriguing, is not enough to save us from some very tough problems in modeling environmental transport.

Practical Environmental Transport Uncertainty: The Deterministic Case, "With Ignorance"

Here the relevant example is groundwater, and the problems are correspondingly easy to visualize. Most of the water in the freshwater cycle at any time is in fact in the form of groundwater aquifers, more or less confined within well-defined boundaries by layers of earth and rock differing in permeability to water. Because most of the volume of an aquifer is composed of solid earth or rock particles, water flows very slowly in those media: the flow volume is restricted, and much work must be done against the large surface area of solid with which the fluid is in contact. Some numbers may help: if many rivers flow at speeds of about one foot per second, the convective flow velocity in many aquifers is often as low as one foot per day. Remembering that there are 86,400 seconds in a day, that is a factor of 86,400 more slowly. Aquifers do, of course, turn over: they lose water naturally from outflows and by evaporation upwards through the water table, and "unnaturally" through pumping withdrawals. And they are recharged by inflows in particular sites and areas. But both those processes are slow, and many important aquifers turn over their contents only over periods as long as three to seven or eight centuries.

Consequently, the stochastic features of the more easily accessible environmental media, such as surface water flows and rainfall, are simply not present in anything like the same way in the groundwater cycle. Short-term fluctuations in rainfall and other sources of recharge will be "averaged out" by the long aquifer turnover times, so that only mean rainfall over many years matters. For all practical purposes, we can treat aquifers as deterministic media.

But what we gain by banishing determinism, we lose in accessibility: aquifers flow underground, sometimes tens of hundreds of feet below the surface, and measurements of their boundaries and flow characteristics are necessarily expensive and imprecise. In some extremely homogeneous and regularly-configured aquifers, things are simple. Particularly when an aquifer consists of a homogeneous permeable layer, say of sand, confined between two regular impermeable layers, a few well borings may be sufficient to establish the geometry. But in many important practical cases, things are far from simple. The aquifer boundaries may be irregular, and the permeable medium of the aquifer itself may differ substantially in its properties over space and even in different directions at each point.

In those heterogeneous aquifer cases, we will always be somewhat uncertain about both the boundary of the aquifer and the spatial distribution of flow characteristics within the aquifer. At some cost, we can drill more and learn more, but it will always be too expensive, and irrational, to remove all our ignorance of the aquifer's geohydrology. The ignorance that

remains we summarize in the probability distributions $\pi_T(s_T)$, with the state of nature variable s_T referring to that ignorance.

Practical Environmental Transport Uncertainty: The Stochastic Case

For the remaining two cases, the Kepone episode and the Chemical Control episode, the kind of uncertainty that arises from the environmental transport medium is more familiar, particularly to water resource specialists. In both cases, the properties of the environmental transport medium are inherently stochastic, and must be treated as such.

In the Kepone episode, there are two relevant sources of stochastic behavior for the environmental transport medium, the waters of the James estuary. The first is the more or less "normal" stochastic variation in freshwater inflows into the James. That variation has been abundantly measured, and models incorporating that variation extensively developed, over the past few decades. The subdiscipline of stochastic hydrology is in fact concerned with the problem of generating the best possible estimate of the "true" underlying flow distribution from existing historical flow data.

But there is a second source of stochastic behavior in James flow: there can be periods of very high flow associated with extraordinary events, such as hurricanes or coastal floodings. Unlike the "normal" sources of flow variation, those extraordinary events can not easily be modeled as manifestations of an underlying probability distribution. Nevertheless, there have been attempts to estimate such a distribution.

And conceptually, the two kinds of probability distributions merge into one master distribution, the distribution of flows in the James estuary. The latter distribution is essentially the one we need, for our principal concern will be the possibility of transporting substantial amounts of bottom sediment into the Chesapeake Bay. Since the probability distribution of Kepone transport can be constructed from the probability distribution of James estuary flows, we can identify the state of nature variable s_T for the Kepone episode with James flow. The relevant distribution $\pi_T(s_T)$ can be constructed by bringing together, in the appropriate way, the normal and extraordinary parts of the James flow distribution.

In carrying out that exercise, we will in fact use a rather primitive kind of stochastic hydrology: in constructing the normal part of the James River flow distribution, we will expand upon the evidence in the historical trace of James flows by generating additional flow sequences representing draws on the same underlying distribution.

EXPOSURE ESTIMATES

For damages to human health resulting from exposures to environmental pollutants, a particular factorization of the problem has become conventional. We separate the process into exposure and dose-response components. Conceptually, exposure is simply the time profile of contaminants to which an individual or a population is subject: there are as many observations as the product of the number of time periods and the number

of contaminants. And, again conceptually, there is a time profile of health status variables with some functional dependence on the exposure time series. That functional dependence is called the dose-response function.

In our equation (2.1), both of these components, exposure and dose-response, have been summarized by the state-of-nature variable s_E and by the distribution $\pi_E(s_E)$ defined over that variable. The ambiguous subscript E used in that summary variable allows, in principle, for a two-component state-of-nature variable summarizing what we know about both exposure and dose-response.

But what we know about exposure is relatively easy to summarize. For exposures through the drinking water route we know, more or less, that individuals tend to drink a liter or two of water a day. Never mind that many people by now have switched to bottled water, for reasons of both taste and concern about health. In principle, this latter complication could be accommodated by estimating the demand for bottled (and presumably contaminant-free) water as a function of price, beliefs about the health effects of contaminants, and other variables. This we have not done; all our drinking water route exposure estimates are based upon the two liters per day assumption.

But for one of our case studies, the Chemical Control case, the exposure route is air: individuals breathe toxic contaminants put into the air by the fire. Our computation of contaminant concentrations give concentrations in the ambient air. Translation of those concentrations into exposure must recognize the rather obvious fact that people spend most of their time indoors. It has become conventional to adjust for that fact by multiplying the "raw" exposures by a factor of 0.1, and that is the adjustment procedure we have adopted.

DOSE-RESPONSE ESTIMATES

We have already suggested that the really difficult component of the state of nature variable s_E summarizing what we know of exposure and dose-response relationships is the component summarizing the latter. For in principle we know how to measure exposures: measure ambient concentrations of the relevant contaminants at several points in time. However expensive and subject to error the actual procedures are, the principle is at least clear.

But the same is not the case for what have come to be called dose-response relationships. At the simplest level, even if the notion of a functional relationship between dose and response was beyond question, it would be far from clear which functions to estimate from which data. For exposure to a particular chemical contaminant, for example, what is the relevant, and presumably related, response: all cancers, some cancers, or some other disease categories?

That is the problem under attack, at present, by both laboratory toxicologists and epidemiologists. The former actually subject organisms, ranging from bacteria to rodents, to the chemical in question, and from the

observed response attempt to infer, or estimate the parameters of, a dose-response relationship. They therefore have the advantage of working with a presumably controlled population under presumably controlled conditions of exposure.

But they are, in another sense, seriously disadvantaged: they do not work with human populations which are, after all, the relevant ones for human health effect dose-response estimates. Epidemiologists, who attempt to infer dose-response relationships in humans from exposure and health data on human populations, in a sense make the opposite tradeoff. While they do work with data on human populations, they necessarily surrender the advantages inherent in data from controlled experiments.

The state of the art in dose-response estimation is such that neither of these methods is clearly superior to the other. Where evidence from one or the other method is available, that will have to do; where evidence from both methods is available, judgment must be applied in synthesizing those data.

For our three case studies, the situation varies considerably. Perhaps the most straightforward is the Kepone case. Kepone is a single, well-defined chemical substance that has been tested in animal bioassays for carcinogenicity. Assume for the moment that cancer induction is the only health effect of Kepone ingestion. Then the most important remaining question is what is sometimes referred to as the question of extrapolation: are chemicals that test positive as animal carcinogens necessarily human carcinogens? And if so, how are the carcinogenic potencies in animals and humans related? Because there are no settled answers to those questions, we have, in the Kepone case study, examined the sensitivity of damage estimates to carcinogenic potency.

In the Price's Pit case study, things are better in some respects but, unfortunately, worse in others. They are better, or at least better defined, in the precise sense that the Environmental Protection Agency has, for each of one hundred twenty-nine so-called priority (water) pollutants, promulgated lower and upper limits on toxicity. They are worse in the sense that those are single-chemical toxicity estimates, and restricted to the chemicals on the priority pollutant list. It is far from clear what the human health effects of those pollutants, taken together, are. And it is far from clear that the chemicals present in Price's Pit but not listed as priority pollutants have negligible effects. Nevertheless, we have, for the Price's Pit case study, simply added the risks from the priority pollutant chemicals. In doing so, the lower and upper bounds of the individual-chemical add up to overall lower and upper bounds.

The Chemical Control case study involves human health damages from inhalation of toxic smoke. We will see that it is difficult to identify the chemical constituents of toxic smoke from burning of even such prosaic materials as wood. For that reason we will simply use a range of chronic effect toxicities based upon linear extrapolation of the acute effect toxicity of an important product of any fire: carbon monoxide. That this approach is unsatisfactory hardly need be emphasized. But until we have a

better theory of combustion it, or something that is not much better, will probably have to do.

Finally, let us return to a point mentioned above: the presupposition that there is a stable dose-response relationship. Similar presuppositions are common in all empirical science, and the phenomenological laws derived by combining such presuppositions with some data are the basis of much engineering. But there is something troubling about carrying this method over into the biological sciences, particularly when the relevant exposures are at the parts per billion level. Perhaps an analogy will suggest what seems questionable: what relationship can capture the response, in degraded performance, attributed to various doses of error in system programming? Both "dose" and "response" seem too broad to capture much of what is really happening.

THE VALUATION PROBLEM

Return now to equation (2.2). Thus far we have been worrying about producing estimates of the joint probability distribution $\pi(s_R, s_T, s_E)$. But we are ultimately interested in V : the cost, to the community at risk, of bearing the risk associated with the hypothetical episode. In principle, that requires that we know something about u , the community's utility function, and W_0 , an "initial wealth" parameter. Taken together, the function u and the number W_0 summarize the community's attitude toward episode risks.

But who is "the community"? Individuals certainly will differ in their attitudes toward all kinds of risks, including the risks associated with major episodes. And how are we to measure, or even learn about, constructs like u and W_0 , when most of our observations are necessarily of individual behavior? These questions are active research topics, and will be for some time. For present purposes, what we need are plausible answers that seem appropriate to the major episode cases we will examine. We develop those answers in two stages. First, we sharpen the distinction between individual and collective risk, a distinction central to the valuation of major episodes. Then we review the evidence on individual risk valuation, and close with some warnings about its relevance here and uncritical adoption elsewhere.

The Distinction Between Individual and Collective Risk

Here we want a clear statement of the distinction between individual and collective risk, and of the relevance of that distinction for major episode risk valuation. Individual risk is the easy case: these are the familiar, insurable risks. Roughly, a risk is individual if it is independent across individuals: the event a New Yorker will have an auto accident imposing a \$10,000 loss is independent of the event a San Franciscan will have an auto accident imposing a \$10,000 loss. Here we use the term independence "naively," but the meaning should be clear: the occurrence (or nonoccurrence) of the New York event will be of no use in predicting the San Francisco event.

However unpleasant the losses associated with individual risk are to the individuals involved, their very independence allows society to "adjust" for them--to spread those losses--in a particularly convenient and appealing way. The mechanism, of course, is insurance. If there are many individuals subject to the same kind of individual risk, an insurance company can sell insurance against the corresponding losses at a price equal to the expected, or "actuarial," loss faced by each individual. And if those individuals are risk-averse, so that they prefer a certain, but slightly smaller, income to an uncertain higher income (with the possibility of a large loss), they will buy that insurance.

But we all face risks of a very different kind: collective risks such as the risk of recession, the risk of war, and the risk of exposures of toxic pollutants released in a major environmental episode. The distinguishing feature of those risks is their collectiveness: when the unemployment rate is greater than 10%, say, earnings losses from furloughs and layoffs are not independent across individuals. Put another way, the unemployment rate in New York is a good predictor of the unemployment rate in San Francisco. And an insurance company cannot sell unemployment insurance at reasonable rates: any company that did so would expose itself to intolerable danger of bankruptcy during the slump.

The same kind of argument goes through, with modifications and some qualifications, for major environmental episodes. Imagine that all asbestos workers during World War II had been sold health insurance and life insurance policies specifically tied to subsequent incidence of asbestos-related cancer. Because the true state of nature has turned out to be "asbestos causes mesothelioma," and because the risks facing individual asbestos workers of subsequent mesothelioma are highly correlated, any company selling such insurance would have long since been bankrupt. Note that all asbestos workers have not contracted mesothelioma, probably because of significant individual variations in exposure and susceptibility. But even in the case in which exposures were identical across individuals but susceptibilities differ, so that not all individuals contract mesothelioma, there is a collectiveness to the risk. Since individuals do not know their own susceptibilities, each individual has, upon exposure, been "endowed" with the lottery on future mesothelioma. The "endowment" is a collective good, and the future risk a collective risk.

Some formal rendering of these arguments may help here. We employ the state of nature (SON) approach to uncertainty. For n individuals in the SON approach, take

$$s = \sum_{i=1}^n s^{(i)}$$

$$s^{(i)} = \sum_{k=1}^{\ell} s_k^{(i)}$$
(2.5)

An element of $S^{(i)}$ is a k-tuple

$$(s_1^{(i)}, \dots, s_k^{(i)})$$

Component 1 describes an individual risk if the marginal distributions p_1 , on S and p_{11} on $S_1^{(i)}$ satisfy

$$P_1(s_1^{(1)}, \dots, s_1^{(n)}) = \prod_{i=1}^n p_{11}(s_1^{(i)}) \quad (2.6)$$

Component 1 describes a collective risk if, necessarily, the realized $s_1^{(i)}$, $i = 1, \dots, n$, satisfy

$$s_1^{(1)} = s_1^{(2)} = \dots = s_1^{(n)} = s_1 \in S^{(1)} \quad (2.7)$$

and the $S_1^{(i)}$, $i = 1, \dots, n$, are copies of one another. In the individual risk case, were we to set out to value a policy which might reduce the exposure of individuals to that risk, we would have to remember that insurance arrangements will arise to reduce that exposure. That is made possible by the statistical independence of the random variables in (2.6). But for a collective risk, things are both tougher for society, and easier for the analyst: insurance arrangements will not arise, and need not be considered, in valuing policies aimed at reducing our collective exposure to the particular collective risk in question.

Thus we have an answer to the aggregation problem: for n identical individuals, the utility function u in equation (2.2) is n times each individual's utility function with respect to the collective risk.

But where are we to find evidence on individual attitudes toward collective risks? Much of the evidence we have--from insurance markets, from the stock market, from participation rates in dangerous sports--is obviously about attitudes toward individual risk. Attitudes toward collective risk show up, but in ways that makes the evidence harder to interpret: in attitudes toward defense spending, flood control programs, price support programs, and the like. That evidence is only now getting its due share of attention. For the present, we have to make what use we can of evidence on individual attitudes toward individual risks.

Individual Valuations of Individual Risks: The Ambiguous Evidence

For almost thirty years, Leonard Savage's (1954) version of expected utility theory has dominated the thinking of economists on decision making under uncertainty. Savage's work completed a development that began with, or before, the efforts of Bernoulli in the eighteenth century. And his exposition in his seminal The Foundations of Statistics is memorable. These

facts account, in part, for the dominant position of Savage's expected utility theory.

It is easy to go from Savage's expected utility theory to two kinds of methods for incorporating risk and uncertainty into cost-benefit analysis. In the first, one "estimates" utility functions, then uses those functions to compute monetary equivalents (either equivalent or compensating variations) of project-related uncertain consequences. In the second, look for a setting in which individuals reveal, through their trades in "implicit" markets, their marginal valuations of incremental mortality risk. Then value policies which increase, or decrease, mortality risk incrementally at that revealed value.

We will follow a variant of that second method in some of our valuation exercises, but only with some misgivings. Those misgivings arise in two ways: from comparisons of existing attempts to derive such implicit market valuations, and from laboratory experiments on human decision making under uncertainty. We briefly review each of these sources of evidence, and then turn to the implications of both, taken together, for our work.

Table 1 below, compiled by Mordechai Shechter, summarizes the findings of some recent empirical work on the valuation of mortality risk. Perhaps the salient feature of that summary is the wide range of estimates of the "value of life" that arise: for example, the Jones-Lee (1976) and Acton (1973) values differ by two orders of magnitude. One of the best known studies, the Thaler and Rosen (1975) study, produces a value very near the lower end of that range.

There are at least two obvious implications of table 1 for cost-benefit analysis. In the case, fortunate for cost-benefit analysts, where policy conclusions are insensitive to variation over the whole range of "life values" between \$100 and \$10,000, there is no problem. But where such conclusions are sensitive to such variation, there must be some effort to understand, in a systematic way, the determinants of the values reported in table 1.

Now take note of the second line of work requiring that we push beyond the standard treatment of uncertainty in cost-benefit analysis. Laboratory experiments over the past twenty-five years aimed at testing the hypothesis that individuals act in a manner consistent with the Savage axioms have converged on the conclusion that they do not. Some of those axioms are in fact consistently violated, and the pattern of violation is sufficiently stable to require renewed attempts at explanation. This area is in its infancy, but there are already some plausible alternative approaches. In the next subsection, we take up our own.

The Assumptions of the Standard Method Revisited

The variation in the values of table 1 is so wide that it is natural to return to the assumptions underlying equation (2.2), so that we can reconsider both the plausibility and their applicability to the major episodes case. Those assumptions include the bases for the standard

Table 1. Values-of-Life Estimates

Source of Evidence	Authors	Value of Life (1980 U.S. \$, thousands) ^a	Associated Incremental Lifetime Mortality Risk
<u>Implicit Values From Labor Market Activity:</u>			
Blue collar workers in manufacturing & construction	Dillingham (1979)	\$ 378	10 ⁻⁴
Workers in risky occupations	Thaler & Rosen (1975)	494	10 ⁻³
Males in manufac- turing industries	Smith (1976)	2,785	10 ⁻⁴
Blue collar workers	Viscusi (1978)	2,820	10 ⁻⁴
<u>Implicit Values From Consumption Activity:</u>			
Residential housing market	Portney (1981)	180	10 ⁻⁴
Residential smoke alarms	Dardis (1980)	351	10 ⁻⁵
Highway speed	Ghosh, Lees, & Seal (1975)	419	10 ⁻⁴
Auto seat belt use	Blomquist (1979)	466	10 ⁻⁴
<u>Contingent Values:</u>			
Air travel	Frankel (1979)	57	10 ⁻³
		3,372	10 ⁻⁶
	Jones-Lee (1976)	10,120	10 ⁻⁶

Table 1. Continued

Source of Evidence	Authors	Value of Life (1980 U.S. \$, thousands) ^a	Associated Incremental Lifetime Mortality Risk
Heart attack prevention	Acton (1973)	59 91	2×10^{-6} 1×10^{-3}
Nuclear power	Mulligan (1977)	428 3,576	10^{-4} 10^{-5}
Reducing cancer mortality	Landefeld (1979)	1,632	10^{-4}
<u>Modified Human Capital:</u>			
U.S. population by sex & age group	Landefeld & Seskin (1981)	898 ^c	
Neoplasm (cancer)	Arthur (1981)	185 ^d	4×10^{-2e}
<u>Implied Policy Values:</u>			
Trihalomethanes in water	EPA (1979)	227	(status quo to 100 mg/l)
Arsenic	CWPS ^b (1976)	6,800	(approx. 10^{-3} to 10^{-5})
Vinyl Chloride	Perry & Outlaw (1978)	9,450	(approx. 2×10^{-5} to 10^{-6})

^aAll values are converted to 1980 dollars using the B.L.S. Consumer Price Index.

^bCouncil on Wage and Price Stability.

^cFor males aged 40 to 44.

Table 1. Continued

^dThe value depends on the (assumed) value of the elasticity of consumption, ϵ , assumed constant across age, where $\epsilon = cu'(C)/u(C)$. $\epsilon = 1$ implies u is linear in consumption. An intermediate value, $\epsilon = 0.6$ was postulated here.

^eRisk reduction is age-specific and ranges from 3×10^{-3} (age 0) to almost 3×10^{-1} (age 80). The listed risk is associated with an age of 40.

Sources: Blomquist (1981), Graham and Vaupel (1981), Landefeld and Seskin (1981), Kimm, et al. (1981), Arthur (1981).

normative theory of individual choice under uncertainty, the interpretation we have given them for the episodes cases, and the use of inferences about risk attitudes drawn from implicit-market studies.

The first set of assumptions has to do with the probabilities appearing in equation (2.2): those are taken to be objective probabilities. By "objective," here we mean corresponding to the best efforts of a team of expert risk assessors: the more usual interpretation of objective as associated with repetitive identical events obviously has no relevance to the major episode case, since major episodes are definitionally infrequent.

The second set of assumptions has to do with the utility function u and the parameter W_0 of equation (2.2). Taken together, they determine the community's degree of aversion to the risk associated with a major episode, and therefore the amount the community is willing to pay not to have to bear the risk of a major episode. If we simply apply the evidence generated by individual behavior in implicit markets to our problem of "estimating" u and W_0 , then we have made two assumptions. The first is that individual attitudes toward collective risk are identical with individual attitudes toward individual risk. And the second is that individual attitudes across classes of similar mortality-related risks are identical: for example, that individuals treat incremental mortality risks of various kinds as representatives of some homogeneous commodity.

It is not hard to concoct arguments for the implausibility of most of these assumptions, and even easier to indicate why they are particularly inapplicable to the major episodes case. But we suspect that the most important problems arise from the assumptions on individual knowledge of, and treatment of, the probability distribution. For that reason we focus, in the next subsection, on that class of objections to the literal interpretation of equation (2.2).

THE NOTION OF ANXIETY

In chapter 1 we argued, in our description of the Three Mile Island accident, that the anxieties associated with the incident were an important component of the costs of bearing the risk of that, and future, incidents. We used the term "anxiety" naively, as if its meaning were clear, and most of us have some more or less clear general idea of what it means to be "anxious."

But if we are to make estimates of the cost of such anxieties, and if those estimates are to have any pretensions to precision, we have to sharpen our definition of "anxiety," and bring that sharpened notion into the domain of utility theory. When we try to do that, we will see that there are several plausible candidate notions of anxiety, and that each has something to recommend it.

Perhaps the easiest way to being to set out the problem is with a formulation of Thomas Schelling's (Zeckhauser, 1974). In commenting on one proposal for incorporation anxiety into utility theory, Schelling argues that

the notion of rationality and the notion of anxiety may be incompatible. For economists, "rationality" of course has a very precise meaning. The axioms of consumer preference theory, for example, define what is meant by a rational consumer; the usual use of those axioms, the derivation of a utility function representing those preferences, simply provides a convenient summary representation of those axioms.

But, asks Schelling, does that mean that a "rational" individual should not be subject to anxiety? For if anxiety is uncomfortable or disabling, the "rational" individual must worry about it, just as he or she would worry about the utility effects of a crippling disease. But then the question raised is the familiar one about the "real me": is a rational individual the one living in a nervous body, or is that individual irrational because his or her anxieties spoil much else that might be pleasurable, thereby diminishing utility?

Schelling's conundrums can only be straightened out if we pick apart the relatively broad and vague notion of anxiety, and parse it into better defined, and more manageable, notions.

At first glance there are many. We list only the most obvious. There is:

- (PR) Anxiety as probability revision: the underlying idea is that individuals are skeptical of the independence of even "genuinely" independent events. When rare, damaging events occur, they revise subjective probabilities upwards; subsequently, barring repetitions, those same probabilities are reviewed downwards by the same moving-average model. Since subjective probabilities are the relevant ones for risk valuation, those revised probabilities are what matter for estimates of the cost of risk bearing.
- (SD) Anxiety as state dependence: choosing some gambles, or being subjected to others, effectively changes the individual's utility function, so that subsequent lotteries are evaluated differently.
- (SEQ) Anxiety as sequentiality: the valuation of intertemporal lotteries can depend on the timing of the resolution of uncertainty. This can happen in two distinct ways, either or both of which may be present in particular cases. The first is a pure "rebudgeting," or intertemporal reallocation effect: earlier knowledge is valuable, because it makes better intertemporal allocations attainable. The second we call a "pure knowledge" effect: utility depends upon knowing something will happen in the future, independently of being able to do anything about it. This latter case is indistinguishable from what we have called "state dependence."
- (BPI) Anxiety as belief-preference interdependence: here utility depends directly upon probabilities; the characteristic independence results of expected utility theory break down.

(CD) Anxiety and cognitive dissonance: anxieties impose a real cost of fear. Individuals reduce that cost by "rationale" deluding themselves about the risks producing those fears. Those delusions are rational in the following sense: they are pushed just far enough to balance the reduced costs of fear and the increased costs of accidents associated with the "mistaken" probability judgments.

CHAPTER 3

KEPONE CONTAMINATION OF THE JAMES RIVER

INTRODUCTION

Kepone, an environmentally persistent and highly toxic pesticide, was developed by Allied Chemical Corporation in 1949 at its plant in Hopewell, Virginia. Patented by Allied in 1952, it was registered as required by the federal government for commercial sale as a pesticide. Kepone has since been used to control ants, roaches, potato beetles, and banana borers. During the sixteen years of its production, Kepone never exceeded 0.1 percent of America's total pesticide production. Annual sales during this period were less than \$200,000 (Goldfarb, 1978).

In the 1960s, the Food and Drug Administration banned its domestic use on food crops (Taylor, 1977). Thereafter, Kepone was produced primarily for export markets.¹ Allied produced Kepone at its Hopewell, Virginia, plant intermittently between 1966 to 1974. In 1974, the corporation contracted with Life Sciences Products, Inc., a company formed by two former employees, to produce Kepone. Under the contract, Life Sciences was to produce the pesticide for a fee, while Allied would retain ownership of both raw materials and final product. Life Sciences produced Kepone at Hopewell for sixteen months. Those production operations were halted in July 1975, when the Department of Health of the Commonwealth of Virginia ordered the plant closed. The reason: symptoms of acute Kepone exposure among Life Science production workers.

What had gone wrong? It does not seem that information on the acute toxicity of Kepone was unavailable. Allied Chemical had conducted toxicity tests in the 1950s to obtain registration under the Federal Insecticide, Fungicide and Rodenticide Act. And Allied subsequently contracted with the Medical College of Virginia to study the acute, subchronic, and chronic toxicity of Kepone (Sterrett and Boss, 1977). In those tests, Kepone was found carcinogenic and highly toxic in all species tested, but those results remained confidential (Jaeger, 1976). For that reason, Allied gave the Food and Drug Administration only "limited toxicological data" (Johnson, 1976) when that agency registered Kepone for marketing.

Moreover, it seems clear that Kepone production operations can be conducted without undue risk to production workers. There were no reports of acute toxic effects during the period in which Allied Chemical directly controlled production operations. Kepone spills were closely controlled, and workers wore safety glasses, rubber boots, and gloves (Goldfarb, 1978). All

this changed radically when Life Sciences assumed control of Kepone production. In April 1974, within three weeks of the start of Life Sciences plant operation, symptoms of acute Kepone exposure among Life Sciences production workers were observed. Several of those workers became sick with tremors, dizziness, and nervousness, symptoms that were called the "Kepone shakes." In September 1974, a former Life Sciences employee filed a complaint with the Federal Occupational Safety and Health Administration.² But no inspection of the plant was undertaken by that federal agency: individual complaints do not automatically lead to agency inspections.

In February of 1975, the Virginia Air Pollution Control Board cited Life Sciences for failing to obtain an air permit. The Board had determined that release of sulfur oxides from the Life Sciences plant were in excess of permitted release levels. One month later, the Virginia Water Control Board linked the malfunctioning of sludge-digester equipment at the Hopewell sewage treatment plant to excessive levels of Kepone in Life Sciences' water-borne effluent.³ The company's discharge permit did not cover industrial discharges,⁴ implying a violation of the stipulations of the permit. Nevertheless, the Board did not revoke that permit. What the Board did do, with the approval of the Environmental Protection Agency, was to set more restrictive pretreatment standards and limit Kepone concentrations in effluent discharges to the Hopewell treatment plant.

But the acute symptoms of Kepone exposures among Life Sciences' workers ultimately focused government attention. On July 23, 1975, Dr. Robert S. Jackson, Virginia's chief epidemiologist, heard of a case of diagnosed Kepone poisoning, and promptly visited the plant. The conditions he found led him to order the plant closed immediately. The next day, Life Sciences management agreed to close the plant and voluntarily comply with conditions set by the Virginia Health Department.⁵ In addition, Jackson ordered physical examinations of all present and former Life Sciences employees.

The examiners concluded that at least seventy to seventy-five workers and ten spouses and children had been poisoned by the pesticide. Twenty-nine of the victims were hospitalized for ailments which included brain and liver damage, sterility, tremors, blurred vision, skin discoloration, joint and chest pains, stuttering, anxiety, involuntary movement of the eyeballs, loss of memory, twitching eyes, and slurred speech (Regenstein, 1982).⁶

On August 19, 1975, the federal Occupational Safety and Health Administration investigated worker exposure to Kepone at the Life Sciences plant. The company was charged with four violations of the Occupational Safety and Health Act of 1970, and a \$16,500 fine was imposed.

The Kepone incident subsequently became the subject of extensive environmental postmortems and legal proceedings. The environmental postmortems determined that Life Sciences, in addition to exposing industrial workers to dangerously high levels of the pesticide, had also released substantial amounts of the toxin to the environment.⁷ Those releases included airborne emissions which eventually settled in the soils around Hopewell, routine wastewater discharges, releases into the sewage system from spills, malfunctioning and bad batches, and bulk disposal of liquid and solid

byproducts in sanitary landfills and sludge lagoons (EPA, 1978).⁸ Preliminary studies of the distribution and concentration of Kepone suggested extensive contamination of biota of the James River and of soils near the Hopewell plant. Further investigation revealed the presence of Kepone in samples of frozen fish taken from the James as early as 1967. Thus Allied Chemical had almost certainly been releasing Kepone into the environment before Life Sciences had assumed control of production operations. In the wake of those findings, Governor Mills Godwin ordered the entire James estuary closed to commercial and recreational fishing in December 1975.⁹ Research on the health effects of Kepone led to the setting of "action levels" for finfish, shellfish, and crabs: the action level is the maximum amount of Kepone in edible portions of fish which the U.S. Food and Drug Administration believes to be safe for human consumption.¹⁰ Public concern over Kepone contamination of the James was heightened by the associated threat of contamination of the Chesapeake Bay located only 120 kilometers downstream.

As the scale of the problem became clear, state and federal governments created extensive monitoring and assessment operations. Virginia and Maryland formed state task forces to monitor ambient Kepone concentrations and Kepone levels in shellfish and finfish. The Virginia Task Force, for example, established a monitoring and surveillance program (Bellanca and Gilley, 1977). The agenda included inquiries into methods and costs of cleaning up the Life Sciences production plant and the Hopewell primary sewage treatment plants, disposing of wastes from the plants, and decontaminating the James River. The same task force also coordinated marine and epidemiological studies to assess the impact of Kepone on the river system. As part of the effort, the Virginia State Water Control Board and the Virginia Institute of Marine Science created a long-term comprehensive sampling and monitoring program of the water, sediment, and biota from the James River. Virginia's Division of Consolidated Laboratory Services and the Allied Chemical Corporation designed and promulgated test and monitoring protocols for determining Kepone concentrations in the air, water, soil, sediment, and biota.

The Maryland Task Force supervised Allied Chemical in containment and storage of Kepone at Allied's Baltimore, Maryland, facility. That group tested residents of the neighborhood of the Baltimore plant for Kepone and found no detectable levels of contamination. When tests of soils near the plant indicated the presence of Kepone in an adjacent park, the park was subsequently closed, stripped, and resodded with uncontaminated soil (U.S. EPA, 1978). The Maryland Task Force also sampled for Kepone in the Chesapeake Bay. Commercial oyster harvesters seed oysters in the James River and then transplant them into oyster bars in the Chesapeake Bay. Oyster bars were therefore carefully monitored, and those bars in which oysters exceeded the Food and Drug Administration's action levels were closed. The bars were reopened a year later when the level of Kepone in sampled oysters had fallen below the detectable level. Both Virginia and Maryland continue to monitor for Kepone.

In August 1976, the Federal/State Kepone Task Force recommended that a feasibility study be undertaken to evaluate mitigation proposals for the

James River. In response, the governors of Virginia and Maryland requested that the U.S. Environmental Protection Agency conduct such a study. In March 1977, that federal agency began a \$1.4 million effort called the Kepone Mitigation Feasibility Program. That project's primary objectives were measurement of the extent of contamination, calculation of the fate and transport of Kepone in the James River system, assessment of current and long-range effects of Kepone contamination on biota, and appraisal and exploration of mitigation and removal alternatives (EPA, 1978). The project was an interagency effort coordinated by the U.S. Environmental Protection Agency, Criteria and Standards Division, Office of Water and Hazardous Materials. Participating agencies included the Department of Energy (in a task agreement with Battelle Pacific Northwest Laboratories), the United States Army Corps of Engineers (coordinating with the United States Fish and Wildlife Service), the Environmental Protection Agency's Gulf Breeze Laboratory, and the Virginia Institute of Marine Sciences. Table 2 summarizes the responsibilities of the various agencies. Data collected by the state and federal monitoring programs constitute a substantial information base for modeling the incident.

Concurrent with these activities were investigations to determine legal responsibility for the incident. The U.S. Environmental Protection Agency discovered that Kepone and two other polymers that Allied had been producing at the time were not included in Allied's 1971 compliance form filed under the Refuse Permit Act of 1899. Thus Allied had failed to receive a permit for these three chemicals, and had been illegally discharging effluents containing them. Subsequent federal criminal investigations traced that omission to a deliberate attempt by some in management to avoid incurring the purchase cost of expensive water treatment equipment, estimated at \$700,000.¹¹ A memorandum written by Virgil A. Hundtofte, later a co-founder of Life Sciences, to ten other Allied executives at the Hopewell plant was introduced into evidence in the criminal proceedings. In that memorandum, Hundtofte wrote, "It was felt that this effluent might go unnoticed by the EPA until we tied into the R.W.T.P. (Regional Water Treatment Plant) or, at worst, interim treatment would not be required" (Zim, 1978). The Justice Department was unable to show, however, that corporate management had endorsed the permit evasion.

On May 7, 1976, Allied Chemical, Life Science Products, the City of Hopewell, and a number of individuals were indicted by a grand jury in the federal district court of Richmond, on 1,097 counts of violation of federal anti-pollution laws.¹² Allied Chemical pleaded nolo contendere (no contest) to 940 charges of violating water pollution regulation, and was convicted on all counts. The largest pollution penalty levied against an American company was then imposed: \$13.2 million. In a negotiated agreement, that fine was reduced by \$8 million. The latter amount was set aside by Allied Chemical to create the Virginia Environmental Endowment. Since the \$8 million was thus transferred to a nonprofit entity, Allied was able to take a tax deduction on the contribution, reducing their net costs by about \$4 million. Life Sciences was convicted on 153 counts and fined \$3.8 million; since the firm's net worth was zero, the fine was a meaningless gesture (Goldfarb, 1978). A \$10,000 fine was imposed on the City of Hopewell which pleaded guilty. The

Table 2. Kepone Mitigation Feasibility Project Responsibilities

DEPARTMENT OF ENERGY/BATTELLE

Sampling and analysis of suspected Kepone contamination to complement existing data.

Acquisition of water quality, sediment, hydrologic, and other data in the James River of coordination with Virginia Institute of Marine Sciences.

Modeling of transport and fate of Kepone in the James River.

Evaluation of nonconventional mitigation techniques.

Assessment of the overall impact of current Kepone contamination and possible mitigation approaches.

U.S. ARMY CORPS OF ENGINEERS/U.S. FISH AND WILDLIFE SERVICE

Analysis of worldwide sediment removal/dredging techniques and applicability.

Engineering studies to contain, stabilize, or remove Kepone-contaminated sediments.

Evaluation of environmental impact of selected engineering alternatives.

U.S. ENVIRONMENTAL PROTECTION AGENCY GULF BREEZE LABORATORY

Effect of Kepone on estuarine biota, including biological accumulation, distribution, and fate.

VIRGINIA INSTITUTE OF MARINE SCIENCE

Field data on biota, sediments, and hydrology of the James River.

U.S. ENVIRONMENTAL PROTECTION AGENCY HEADQUARTERS

Program management and report.

Source: EPA Mitigation Feasibility for the Kepone-Contaminated Hopewell/James River Areas, 1978, p. IV-3a.

two owners of the Life Sciences plant, Virgil A. Hundtofte and William P. Moore, each were fined \$25,000 for polluting the James River.

Former employees of the Life Sciences plant decided to sue Allied Chemical for compensation for their health-related damages, since Life Sciences was nearly bankrupt. Those workers sued for a total of \$186.3 million, but most of the cases were settled out of court for approximately \$3 million. The state of Virginia received \$5.25 million as a partial settlement for violation of its water pollution control law. Virginia reserved the right, however, to sue Allied for costs necessary to cleanup the James River and dispose of contaminated residuals. About 400 individuals, whose primary occupation involved the sale of fish and oysters, sued Allied for \$24 million. A class-action suit for \$8.5 billion for lost income on behalf of 10,000 people working in marine-related businesses was also filed against Allied (Pruitt v. Allied Chemical Corp.).

The environmental impacts of the Kepone release will extend over many years. Laboratory studies at Gulf Breeze indicate that Kepone neither biologically nor chemically degrades in simulated estuarine systems. Estimates suggest that it may take decades for natural dispersion mechanisms to reduce concentrations of Kepone in the James River below those required by Food and Drug Administration action levels. Uncertainties regarding the effect of Kepone on living organisms led to extensive research projects which have attempted to determine maximum safe dosage levels for exposed biota. Action levels have been revised in accordance with laboratory findings. In March 1977, the Food and Drug Administration relaxed action levels for finfish from .1 to .3 parts per million. Experiments on rats indicate that Kepone is a cumulative poison, so that toxic levels result when small amounts are ingested over a long period of time. Studies also showed that since Kepone is fat soluble (lipophilic), it tends to bioaccumulate in the fatty tissues of the body. Transport of increased concentrations of Kepone through both aquatic and terrestrial food chains thus has been an additional source of concern.

Apprehension over the possible transport of large quantities of Kepone into the Chesapeake Bay has led to extensive studies and to the construction of sophisticated models of Kepone transport in the James estuary. To trace the distribution, concentration, and effect of the Kepone releases, an understanding of the interrelationships between the component parts of the estuarine system is necessary, for the hydrodynamic and geological features of natural water systems, in conjunction with the physical, biological, and chemical characteristics of both the system and the chemical, determine the pathways of substances (O'Connor, Farley, and Mueller, 1981). Models have been developed that include spatial, temporal, and chemical transformation and transport processes for chemicals in waterways. Laboratory studies of partitioning reaction in estuarine systems have found that Kepone tends to attach itself to particulate matter. As a result of the settling of particulate matter, elevated concentrations of Kepone are found in bed sediments.

Remaining questions have centered on the long-term fate of Kepone in the James estuary and in the Bay. The U.S. Environmental Protection Agency

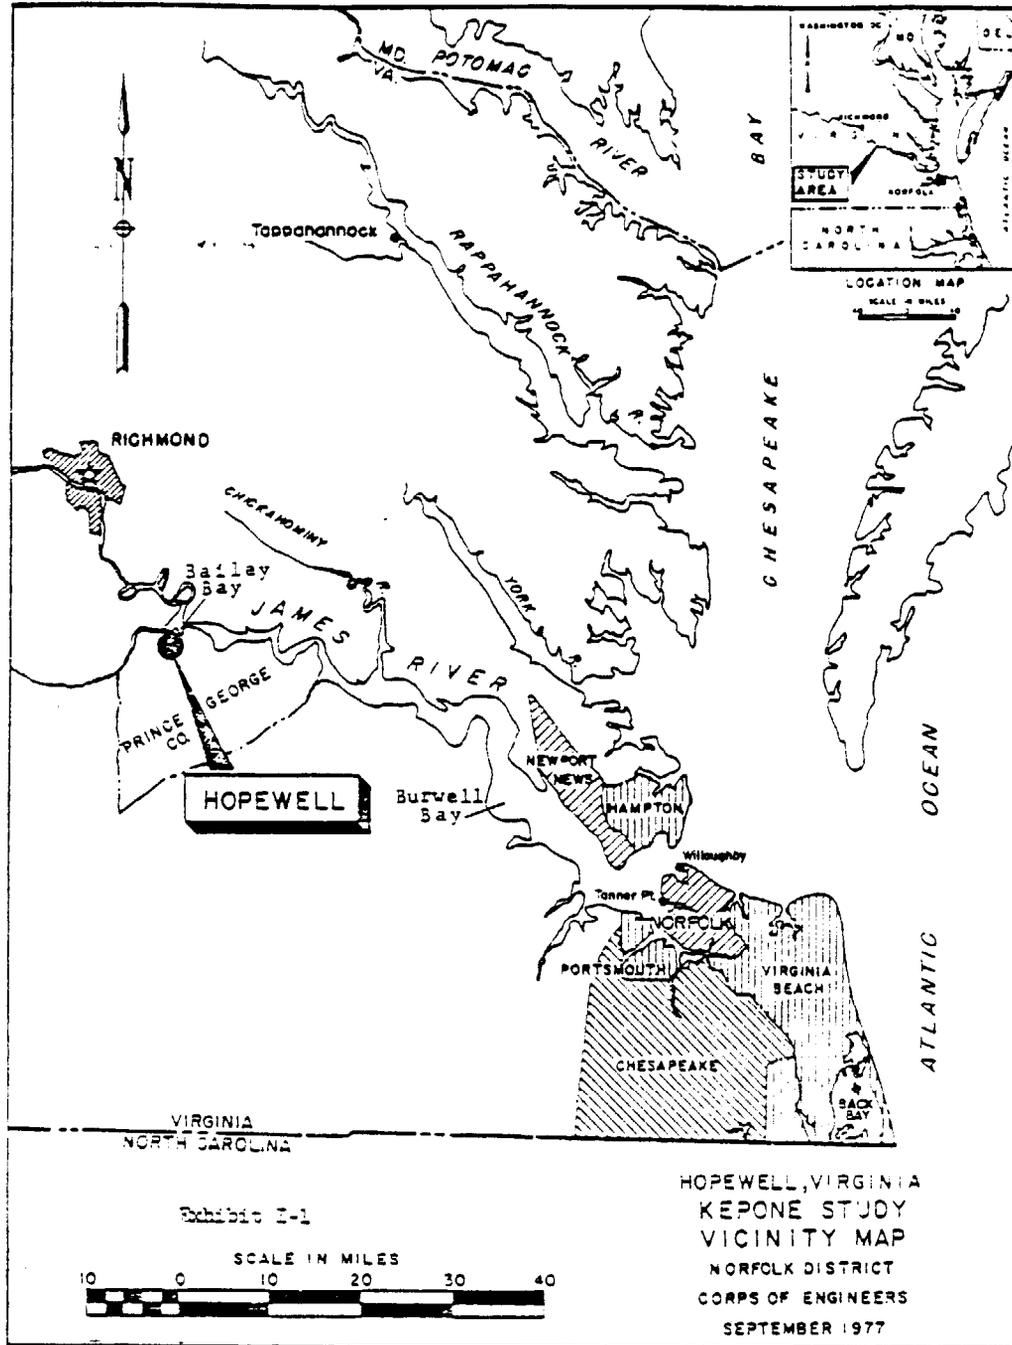
Mitigation Feasibility Project ruled out a full-scale cleanup of the James River as prohibitively expensive (the lowest cost estimate was over \$3 billion). In any event, that cleanup action might in fact further disperse the Kepone in bottom sediment. Remedial mitigation efforts, such as dredging limited to areas of high Kepone concentration, were rejected by the Environmental Protection Agency on the grounds that the resultant reduction of Kepone would not have immediate effects but would be evident only in the longer run (National Wildlife Federation, 1981). Thus there will be continuing low-level contamination of the James estuary, for perhaps several decades, until Kepone is either buried in bottom sediment or washed into the Bay. For some substantial portion of that period, there will continue to be a possibility of a high-flow incident in the James transporting substantial amounts of Kepone into the Bay. And finally, we take note of another possibility in Kepone-like incidents: that of human health damages from drinking water exposures. Because the James is estuarine and brackish further upstream than Hopewell, there were no such exposures in the actual Kepone incident. But in a Kepone-like incident in a river which serves as the source of drinking water, human exposures would be possible.

KEPONE MODELING: A SEQUENCE OF EXPERIMENTS

Coming face to face with the problems of modeling an estuarine ecosystem is a humbling experience, and one which we wish to share. The James River and the James estuary are divided by an imaginary fall line marking the point at which the river becomes nontidal: that fall line (figure 2) lies above the Hopewell, Virginia, location of the Life Science plant from which Kepone was released into the James. Thus all of the actual "Kepone problem" belongs, and must be treated in, an estuarine setting.

That setting complicates our task because the tidal variation must be modeled if we are to obtain an accurate representation of Kepone ambient concentrations and transport. This issue arose relatively early in the debates about estuarine modeling, and we believe that it has been decisively settled in favor of those who, like Donald Harleman, argued this position, in Ward (1971). There is a simple and compelling argument underlying that position. The equations of hydrodynamics can be made tractable only by approximating the naturally-occurring coefficients of diffusion. But premature averaging--averaging over the tidal cycle before solving the equations--in effect leads to the wrong set of diffusion coefficients, and thus ultimately to erroneous predictions. There is nothing wrong with averaging over the tidal cycle after computation: tidally-averaged summary statistics may be easier to interpret and understand, and will not be misleading.

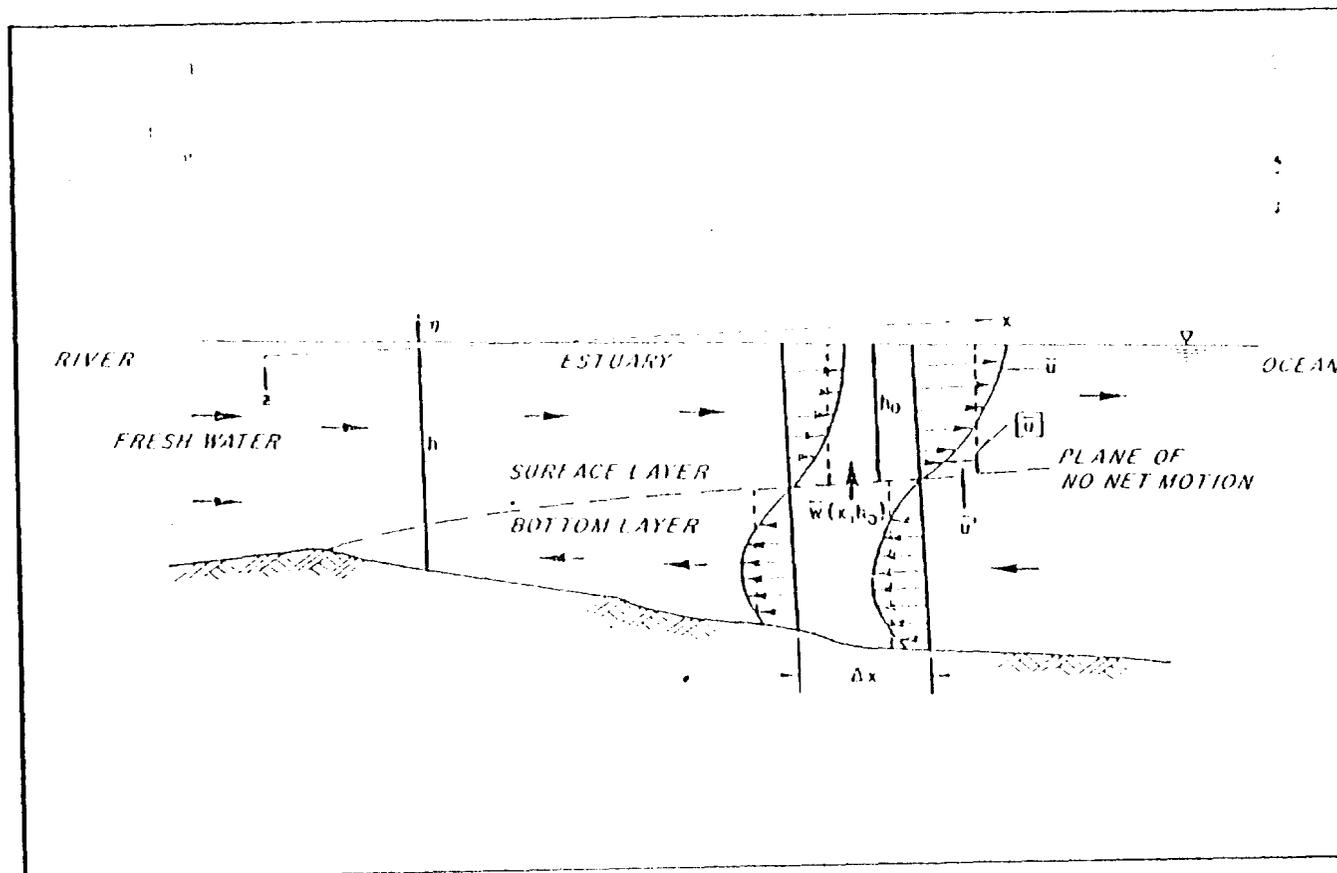
The above argument applies to all estuaries, stratified or unstratified. In the particular case of the James estuary, there is considerable stratification, with a saline "wedge" carrying water upstream, and a freshwater upper layer carrying water downstream (figure 3). Over the tidal cycle, the average velocity in the wedge is upstream and the average velocity in the upper layer is downstream. Of course, the net flow is downstream and equal to the freshwater runoff into the James. But what is remarkable is how large the tidally-averaged layer velocities are relative to the net flow



Source: U.S. Environmental Protection Agency. 1978. Mitigation Feasibility Study for the Kepone-Contaminated Hopewell/James River Areas, EPA-590/9-77-033 (Washington, D.C., U.S. EPA).

Figure 2. The James River Estuary and Chesapeake Bay

Figure 3. The Salinity Intrusion and the Freshwater Layer of the James



Source: Donald J. O'Connor, Kevin J. Farley, and John A. Mueller. 1981. "Mathematical Models of Toxic Substances in Estuaries with Application to Kepone in the James River," Grant No. R-804563 (Bronx, N.Y., Manhattan College, Environmental Engineering and Science Program).

(table 3). For later reference, examine the first row of table 3, with net freshwater inflow of 1,000 cfs: the top and bottom layers are flowing at rates on the order of 40,000 cfs, a factor of 40 higher. That differential is particularly important because the relevant flow velocity for calculation of the resuspension of Kepone in bottom sediment is of course the bottom layer flow velocity "seen" by the sediment, and not the net, layer-averaged freshwater outflow.

That brings us to the role of sediment in Kepone transport in the James. Superimposed upon the complex hydrodynamics of the stratified estuarine system we have the coupled sediment transport system, critical because Kepone is relatively insoluble and is relatively strongly adsorbed onto particulate sediment. By adsorption we refer to the binding, by intermolecular forces, of Kepone molecules on the surfaces of sediment particles. Because that binding is, in the case of Kepone, relatively strong, and because Kepone is relatively insoluble in water, most Kepone in the James is presently resident in James bottom sediment.

Taking these elements together, we can sketch a rough picture of the major features of Kepone transport in the James. Kepone released into the James at around Hopewell will either adhere to sediment in the water column, or remain dissolved in James water. Dissolved and adsorbed Kepone above the estuary bottom will come into equilibrium. Similarly, sediment particles in the water column are continually being deposited on that bottom, and bottom sediment is continually being resuspended. Taken together, those processes gradually transport Kepone downstream and toward the mouth of the James on the Chesapeake Bay.

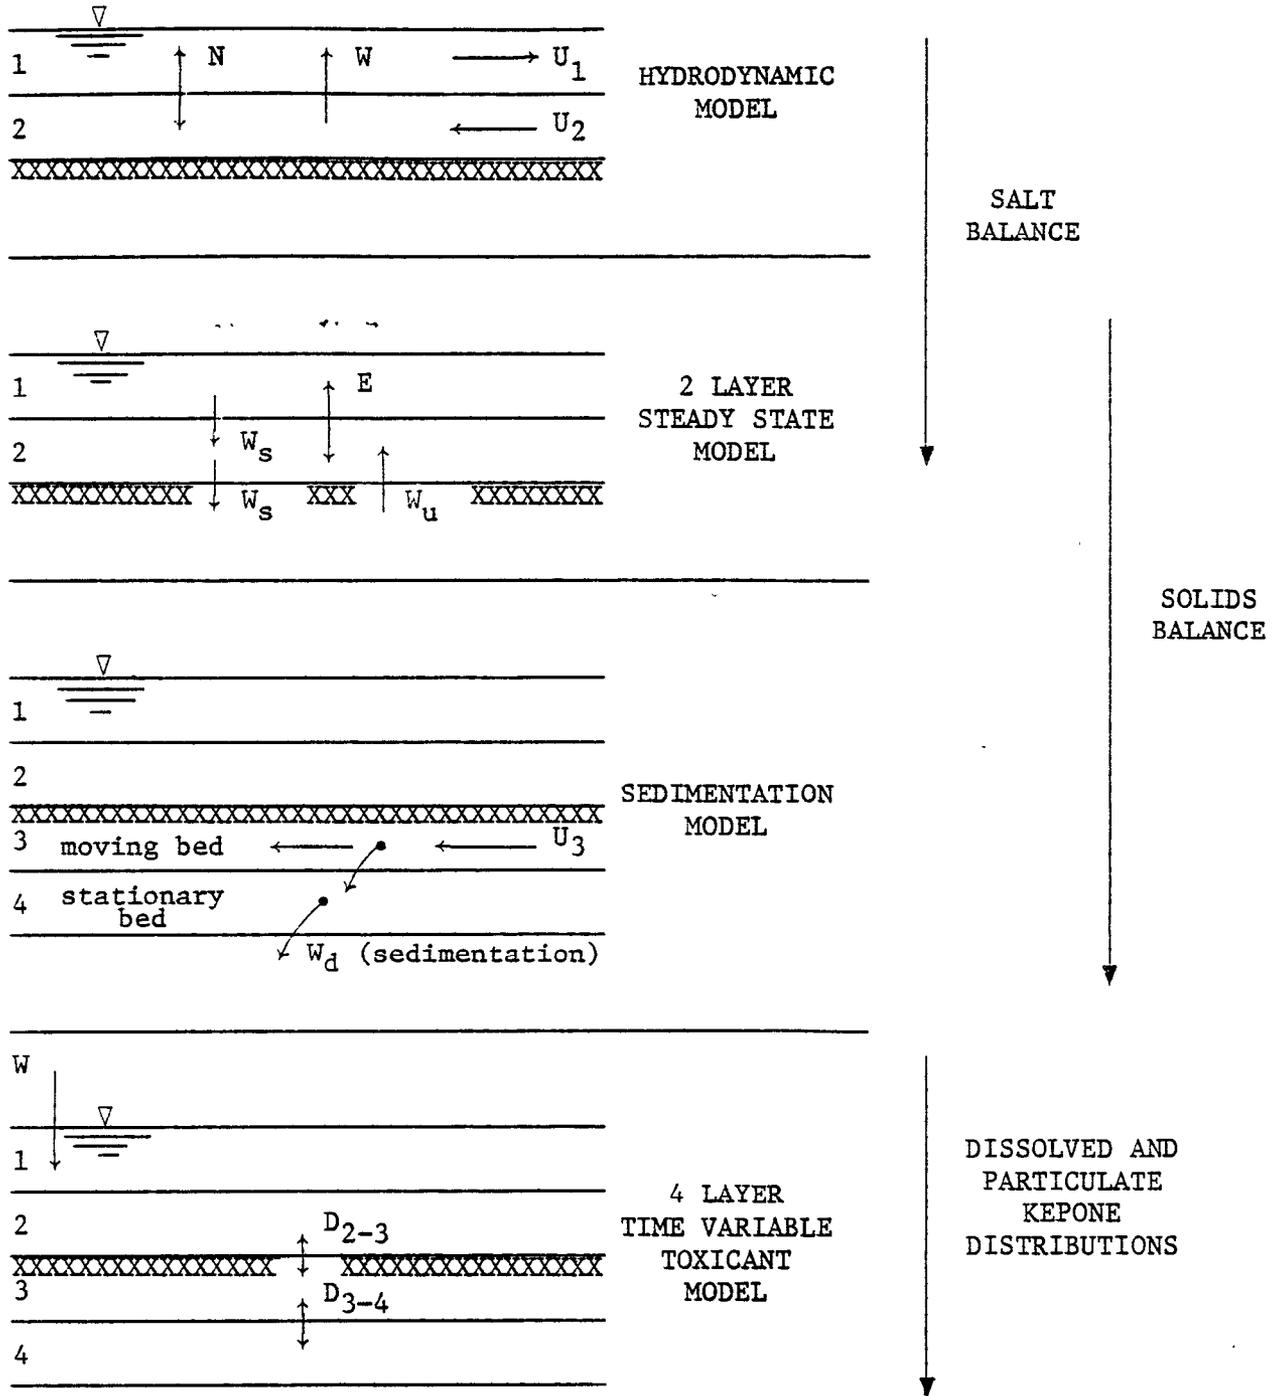
A general conceptual framework for estimating Kepone concentrations and Kepone transport rates is sketched in figure 4. Though we are critical of some of the work from which that figure is drawn, the general framework is unexceptional. At the top level is the hydrodynamic model, since hydrodynamics drives the entire system. At a second level lies a model of the sediment system, which is coupled to, and driven by, the hydrodynamic model. Finally, at a third level, Kepone transport is coupled to the sediment model by the processes of adsorption (of Kepone onto sediment) and resuspension and deposition of sediment. The latter processes provide the mechanism for exchange of Kepone between the water column and the bottom sediment of the estuary.

Figure 4 is of course a schema, and far from implementation. For implementation, a whole series of critical choices must be made. To some extent, those choices can be loosely summarized as: choose appropriate levels of spatial and temporal disaggregation for each of the model levels of figure 4. "Appropriate" means with reference to some purpose: we must decide exactly what we are trying to estimate. In preparation for the discussion and those decisions, let us compare, along a few relevant dimensions, several well-known models of contaminant transport in estuarine environments.

Table 4 summarizes those comparisons. Of the four models listed, only one, the Harleman, Holley, and Huber (1966) model, is an analytical model;

Table 3. Net Flow and Layer Flows in the James

Flow Pattern	Flow (Cubic Feet per Second)				
	James River	Appomattox River	Chickahominy River	Downstream Net Tidal Circulation	
				Top Layer	Bottom Layer
1	1,000	152	56	41,600	40,400
2	3,200	495	182	43,380	39,500
3	4,380	657	263	73,300	68,000
4	7,044	1,334	407	74,790	66,000
5	11,500	1,770	651	89,920	76,000
6	18,500	2,775	1,110	126,400	104,000



Source: Donald O'Connor, Kevin J. Farley, and John A. Mueller. 1981. "Mathematical Models of Toxic Substances in Estuaries with Application to Kepone in the James River," Grant No. R-804563 (Bronx, N.Y.L, Manhattan College, Environmental Engineering and Science Program).

Table 4. A Hierarchy of Submodels for Kepone Transport

Table 4. Comparison of Models of Contaminant Transport in Estuaries

Model	Analytic or Numerical	Time Scale	Spatial Scale	Model of Deposition and Resuspension
I Harleman, Holley, and Huber (1966)	Analytic	Intratidal	Continuous (analytical model)	None explicit
II Dailey and Harleman (1972)	Numerical	Intratidal	1 kilometer horizontally	None explicit
III O'Connor, Farley, and Mueller (1981)	Numerical	Tidally- averaged	1 kilometer horizontally, 1 foot vertically; stratified	Deposition at constant settling rate; resuspension at a rate related to lower- level flow velocity
IV Onishi (1977)	Numerical	Steady- state	1 kilometer horizontally, vertically- averaged (not stratified)	Deposition and resuspen- sion tied to net flow velocity (or equivalently, to net freshwater dis- charge)

more precisely, it reduces the problem of computing contaminant concentrations to the problem of performing one numerical integral. The remaining three models are numerical, requiring varying degrees of numerical computation. They vary widely in their computational cost: very roughly, models II and III cost about \$20 per run, while model IV costs about \$500 per run.

How does one decide upon the best model, or more appropriately, upon a plan for using the existing models? As we have remarked above, it depends upon the purpose of the exercise. Our own purposes have dictated the dimensions of comparison listed in table 4. Of the four models, only I and II are intratidal, an advantage we have argued is crucial. Models III and IV have what we consider other troublesome features: of these the most critical is the representation of the sediment deposition and resuspension processes in those models.

Recall that much Kepone in the James is adsorbed onto sediment particules, either suspended in the water column or resident in the estuary bottom; and that one of the serious potential problems posed by Kepone in the James sediment is contamination of the Bay. For now, simply refer back to figure 2. Clearly, the rate of Kepone transport into the Bay must depend upon the flow pattern in the James, and clearly that flow pattern must be faithfully represented in a computation of Kepone transport into the Bay.

But now refer back to table 4. In model III, deposition at a constant rate is assumed (the rate actually used in the model is four feet per day). If one thinks back to the underlying physics, here is what is happening. Sediment, and Kepone adsorbed on that sediment, is being resuspended and, on balance, carried downestuary. That is true even though, on balance, the (tidally-averaged) flow in the bottom, or saline-intrusion, layer, is landward. The seeming paradox is no paradox at all because there is some interchange of water, and hence of Kepone, between the two layers, and because on balance, the flow in the top layer is downestuary, or seaward.

As sediment is advected downestuary, it is subjected to three kinds of forces: gravitation, impulsive forces associated with turbulence, and intermolecular forces. The latter need not concern us here: the former act in opposite directions, with gravitation inducing a net downward drift, and turbulence contributing to continued suspension of sediment in the water column. If the balance between these two countervailing forces is velocity-dependent, then the overall transport rate of Kepone into the Bay may be sensitive to the flow pattern in the estuary.

We have similar concerns about model IV, which builds on one of the most detailed hydrodynamic model of the James estuary. In that model, the stratification of the estuary is neglected, and the Kepone transport model is calibrated to net flows. But since flow velocities in the bottom layer can be as much as a factor of forty greater than net flow velocities, and since resuspension phenomena really depend upon the flow velocities "seen" by bottom sediment, Kepone transport predicted by the models may be artificially velocity-insensitive.

We shall shortly see that the velocity dependence of Kepone transport is, for our purposes of benefit estimation, a key issue. We summarize this section with a strategy for bending models I through IV to our particular purposes of risk assessment. We will begin with model I, because it is easily manipulable, computationally inexpensive, and intratidal in time scale. But it has no sediment transport "submodel," and so we will have to invent one. The same will be true of model II. The "invented" submodels will be useful for placing upper bounds on benefits; we will argue that those bounds are better as the distribution of flow velocities shifts upwards. To go further, we must identify the damage or benefit categories that we think may be significant, and choose a Kepone modeling strategy for each.

AN OVERVIEW OF THE ANALYSIS

Our goal is to derive damage estimates corresponding to these three very different kinds of Kepone-related problems. Each of those damage estimates will require a distinctive approach to modeling the Kepone-contaminated estuary-Bay system, and it is to that task of modeling that we turn in the next chapter. But before we do, it may be helpful to take an overview of what is to come. Our assumptions about the source term associated with Kepone in James bottom sediments are discussed in chapter 5 in the section on Evidence of Contamination in Drinking Water and in the appendix; for the most part we have imputed the source term from measurements of ambient Kepone concentrations and from previous calculations of Kepone transport into Chesapeake Bay under various hydrological conditions. We have also placed a detailed description of the hydrological data we have used in the appendix; for the most part this is United States Geological Survey data.

In chapter 5, the section on Toxicological Significance of the Contaminants describes our construction, from that data, of synthetic James flows: that construction is necessary because hydrological data is sparse, and synthesis of the full distribution may be necessary to capture tail events with large associated damages. From that basic synthetic distribution of James flows we can derive two related distributions we will need: a probability distribution for ambient Kepone concentrations in the James and a probability distribution for Kepone transport into the Chesapeake Bay, (constructed in this section and described in tables 14 through 16).

From those two basic derived probability distributions, and from two kinds of additional assumptions, the cost of risk bearing estimates that we are after follow easily. The additional required assumptions cover dose-response and value of health risk parameters (in chapter 5's section on The Likely Effect of the Contamination on Water Supplies) and values for James and Bay service flows (in chapter 5's postscript).

NOTES

¹An estimated 90% (Sterrett and Boss, 1977) to 99.2% (U.S. Environmental Protection Agency, 1976) of the Kepone that was produced in the United States was exported to the Caribbean, Central and South America, Africa, and Europe. Domestic uses of Kepone were for ant and roach control only.

²Production standards requiring workers to use respirators and gloves as well as shower and change clothes before leaving the plant were ignored by Life Sciences management. Workers frequently ate their lunches at the workplace where Kepone dust was found to lay sometimes as much as several inches deep on the floors.

³Normal bacterial action required for sewage treatment was inhibited by the Kepone in the effluent. As a result untreated sewage discharges polluted a large part of the James River (Sterrett and Boss, 1977).

⁴Allied discharged untreated toxic effluent directly from its Semi-Works Kepone production plant into a tributary of the James River called Gravelly Run. In 1970, restored provisions of the Refuse Permit Act of 1899 required that all industries which discharges wastes into navigable waterways obtain a permit from the U.S. Army Corps of Engineers. Since toxic discharges were unlikely to be approved, Allied filed a short-form permit application that allowed unmetered and unsampled "temporary discharges." The temporary permit enabled Allied to avoid the purchase of costly pollution control equipment. When Allied's permit expired in 1972, the Environmental Protection Agency was given discharge permit authority under the Federal Water Pollution Control Act Amendment of 1972. Allied filed to renew their permit again under temporary status to buy time until the completion of a municipal Regional Water Treatment Plant. When Life Sciences assumed production in 1973, they neglected to file with the Environmental Protection Agency. Instead, Life Sciences applied to the Virginia State Water Control Board for a permit to discharge sanitary wastes. Although the treatment plant was not capable of degrading Kepone, they were granted a permit through an arrangement with C. Jones, the Director of Hopewell's Department of Public Works and former plant manager of Allied's Semi-Works Plant. Permission was granted provided Life Sciences would meet a pretreatment standard of three parts per million of Kepone. The permit, however, made no explicit mention of industrial wastes.

⁵Although the Virginia State Health Department ordered production to cease at the Life Sciences Products plant on July 25, 1975, production continued into September.

⁶If longer-term ailments, such as cancer, reproductive impairment, or latent neurological or psychological effects of the exposure develop in later years, it will be more difficult to link those effects with exposures from ingestion of contaminated seafood.

⁷Limited information has impeded a determination of the amount and timing of Kepone releases into the environment. Estimates of releases have been based on assumptions of correlations between releases and production levels (O'Connor, et al., 1981).

⁸An additional source of residue resulted from the dismantling of the Life Sciences plant. The plant's machinery was taken apart and placed in a sealed pit at the Hopewell city dump. Water from the site was collected in thirty-three tank cans to be filtered and sprayed on some of Allied's property in Hopewell. Since Life Sciences was nearly bankrupt, Allied agreed to dismantle the plant, which cost them approximately \$394,000. Allied believed it would be futile to attempt to detoxify the plant (Fortune, 1978).

⁹To date, the James River remains closed to fishing. During the past seven years, Virginia's governors have lifted, and then reimposed, parts of the ban, but federally-mandated levels of Kepone allowed in catches have restricted commercial fishing in the James River. The commercial seafood industry has urged Virginia state agencies to relax government safety standards for Kepone substantially. A group of Virginia state agencies have proposed a three-fold increase in the action levels, from .3 to .9 parts per million for all finfish. Virginia Secretary of Commerce and Resources Betty Diener claimed that the level of the increase was calculated to permit Virginia to lift the fishing ban, since no fish caught in the last five years has ever exceeded that level (Isihoff, 1982).

¹⁰Federal action levels are based on laboratory experiments which determine exposure levels of intake necessary to produce malignant tumors in animals. To establish the federal "action level," the intake that produced the tumors is multiplied by a traditional safety factor of 10^{-3} . Action levels are usually measured in parts per million (ppm). In March 1976, the Food and Drug Administration adopted action levels for Kepone-contaminated shellfish (.3 ppm), finfish (.1 ppm), and crabs (.4 ppm).

¹¹In comparison to the projected costs of billions of dollars to decontaminate the James River, the expense to purchase necessary equipment would have been insignificant.

¹²For discharging Kepone without a permit from its Semi-Works plant, Allied was charged with 940 violations of the Refuse Act (Rivers and Harbors Act of 1899) and of the Federal Water Pollution and Control Act of 1972. At the time, there were allegations that Allied was conspiring to provide false information to the United States. Allied, Life Sciences, Hundtofte and Moore (the co-owners of Life Sciences), and the City of Hopewell were charged with unlawful discharges of Kepone into the sewer system. No criminal charges were made regarding the conditions at the Life Sciences plant because the Federal Occupational Safety and Health Act did not provide for them in this circumstance.

CHAPTER 4

THE BENEFITS OF AVOIDING A KEPONE-TYPE INCIDENT

A LOTTERY REPRESENTATION OF A KEPONE-TYPE INCIDENT

To begin, it may be helpful to form a picture of a simplified topography capturing the essential features of the real-world James estuary and Chesapeake Bay. Figure 5 presents what we will refer to as the topography of the "didactic" Kepone incident.

DAMAGE ESTIMATION: A FRAMEWORK

Figure 6 then sets out a compatible damage-estimation framework. The didactic Kepone incident is viewed as a compound lottery. If that lottery is accepted, then the first branching point corresponds to the release subplottery: either there is, or there is not, a substantial release of Kepone into the James River. If a release does occur, a second subplottery is imposed: there may, or there may not, be both contamination of drinking water and closure of the James to recreational and commercial fishing, depending upon flow conditions. A third (and final) imposed subplottery indicates the possibility of substantial transport of Kepone into the Bay, and thus of substantial contamination of the Bay.

This lottery representation is a convenient way of identifying the major damage components, and a prescription for combining those components into an estimate of the benefits of actions or policies aimed at preventing a Kepone-like incident. The major damage components identified are health-related damages from drinking water contamination, the loss of the services of the James fisheries to recreational and commercial users, and the loss of the fishery services of the Bay. The prescription for damage estimation is then simple: compute willingness to pay to avoid bearing the risk associated with the lottery.

We would argue that this lottery representation is a natural one, in the sense that it captures the significant characteristics of the Kepone problem in the James. As stated, the Kepone suspended in James River water is in equilibrium with a much larger amount of Kepone in James sediment. Therefore, at current rates of transport out of the river and estuary, and into the Bay, it will take about 200 years before the James is cleansed of Kepone by natural processes. This fact is necessary background for estimation of two damage categories: recreational use of the James (in which fish are taken to be eaten) and human health effects from contamination of drinking water. It can be reasonably assumed that once contaminated, the James will be closed to commercial and recreational fishing (in which the

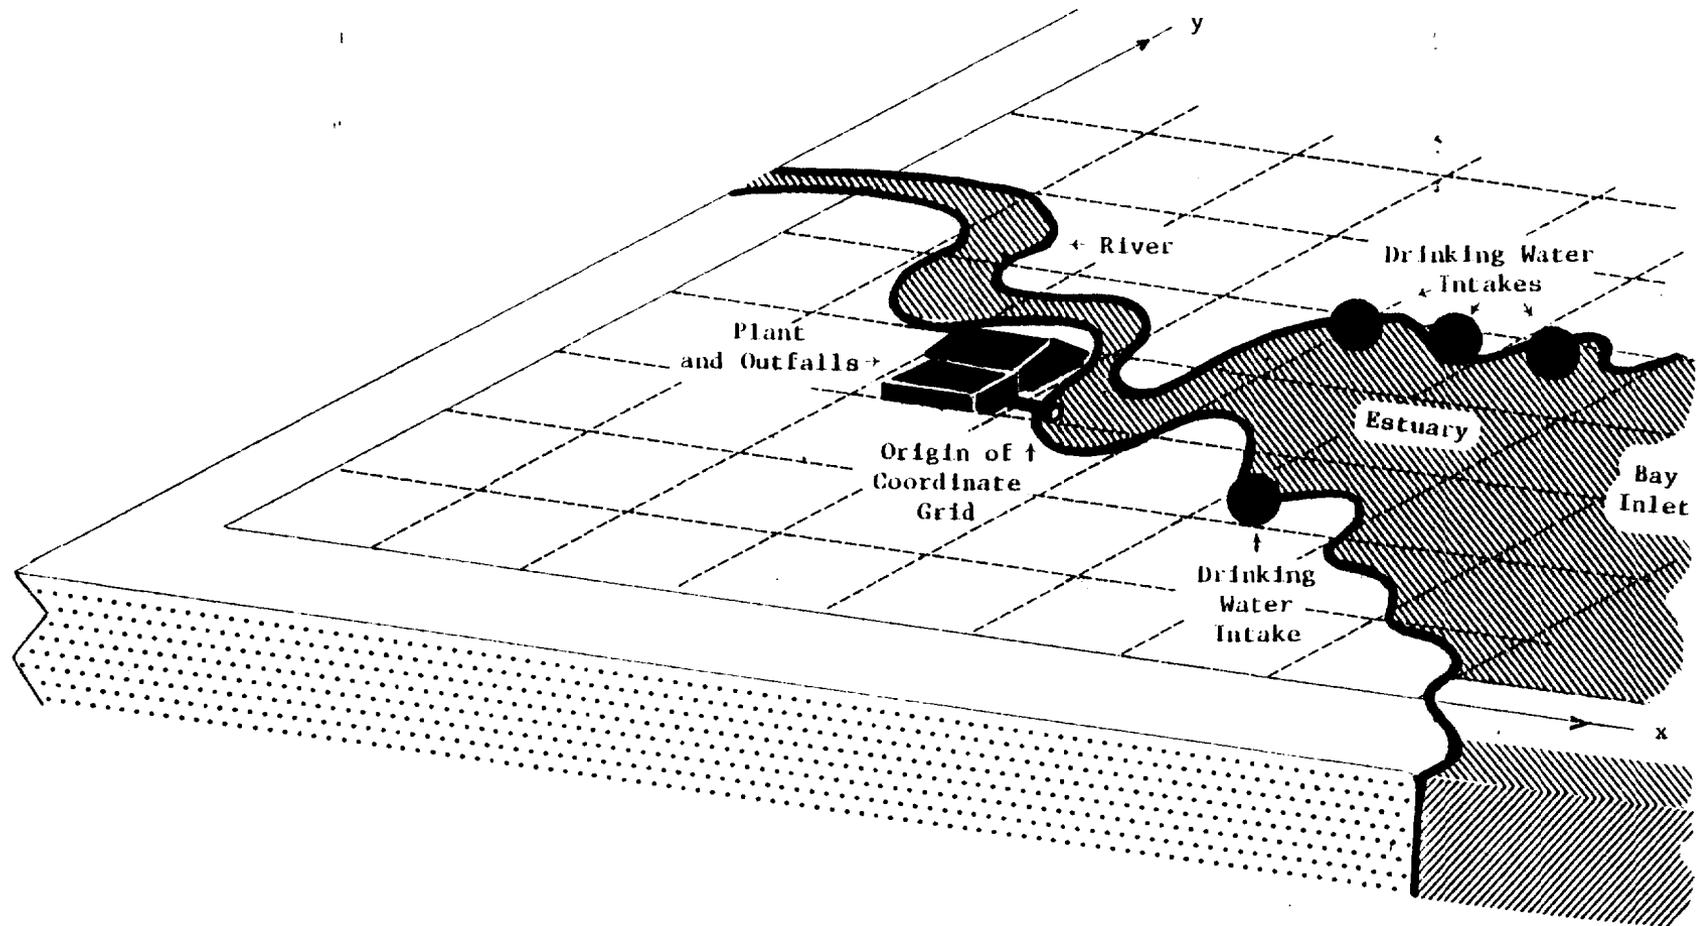
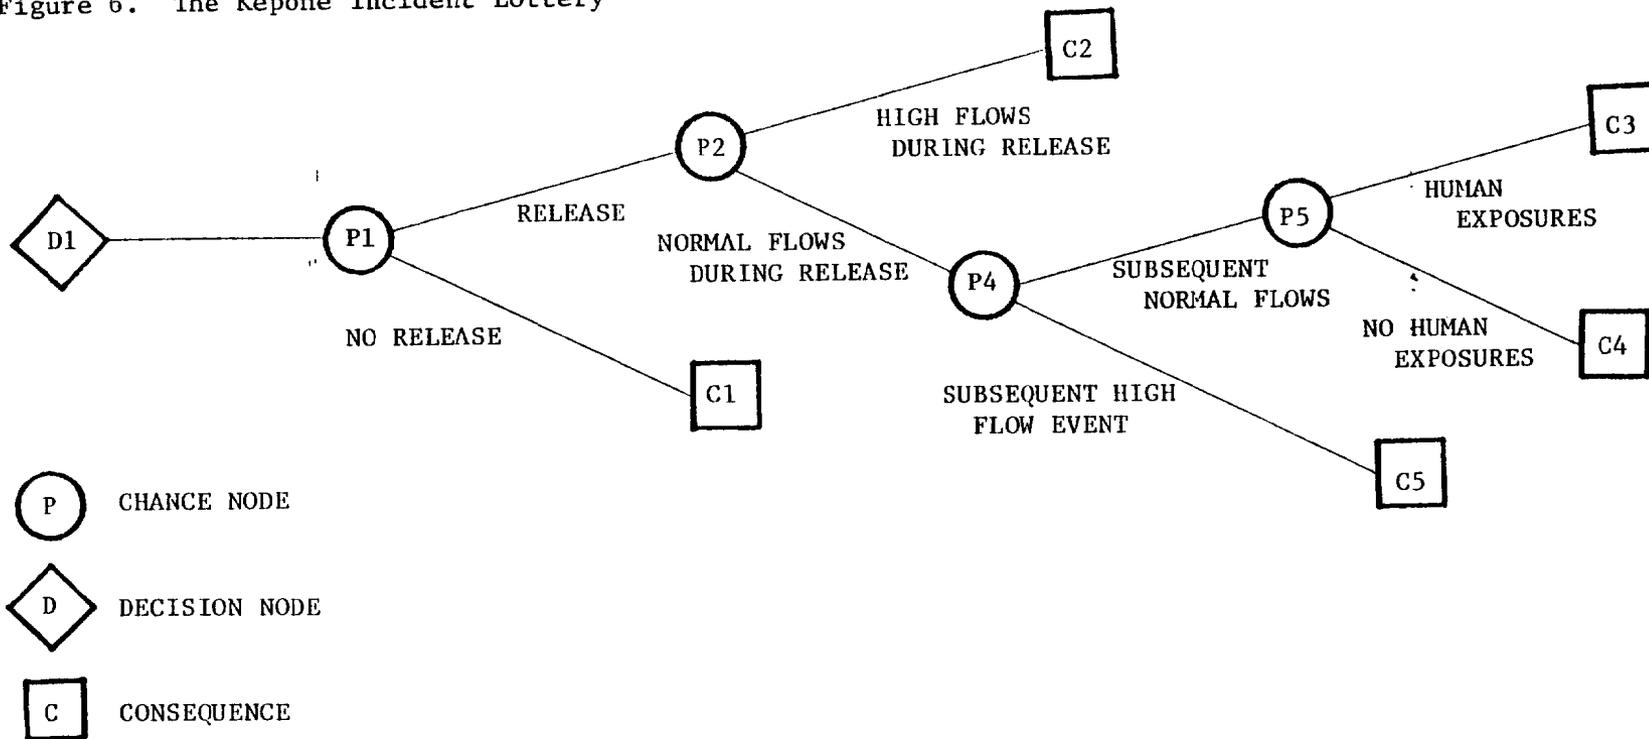


Figure 5. Topography for the "Didactic Kepone Incident" Case Study

Figure 6. The Kepone Incident Lottery



59

- D1: Lottery accepted; plant operated
- C1: No release from plant
- P2: Flow during release lottery
- C2: Transport into Bay during release
- P4: High-flow event lottery
- P5: Human exposure lottery
- C3: Human exposures and remediation
- C4: No human exposures and remediation
- C5: Transport into Bay during subsequent high-flow event

fish and shellfish taken are consumed) for the indefinite future. One ex post damage component is therefore the discounted present value of those particular service flows of the James River and estuary.

As indicated in figure 5, drinking water intakes may be located along a river subject to Kepone-like incidents. In the case of the real-world Kepone incident, the counties adjacent to the James are sparsely populated, and brackish estuarine water is in any event not suitable for human consumption. For that reason, human health damages associated with general-population exposures to contaminated drinking water from the Kepone incident are almost certainly negligible. Nevertheless, in keeping with the didactic use of the case study incidents we examine here, with research primarily directed toward methodological improvement, estimates are made of how large those exposures might have been in an incident in which the riparian counties were populous and the water affected otherwise potable.

Finally, consider the possibility of significant Kepone contamination of the Chesapeake Bay. We have noted that measurements of Kepone contamination of the Bay taken after the original incident imply low transport rates. The associated damages to the recreational and commercial fishery services of the Chesapeake Bay are almost certainly modest, since the Bay dilutes, well below Food and Drug Administration action levels, routine current Kepone inflows.

But it would not be correct to estimate the Bay-related component of Kepone-incident damages only from our ex post observation of current Kepone contamination levels. We have argued for an ex ante perspective and for the representation of our didactic Kepone episode as the lottery of figure 6. The conceptually correct measure of the cost of accepting that lottery is willingness to pay rather than accept the lottery. The relevant willingness to pay measure extends over all ex ante possible outcomes of the lottery, and not only over the particular outcome that has been observed ex post.

Certainly there is a real possibility of much worse ex ante outcomes, arising from much higher rates of Kepone transport into the Bay. Kepone transport rates might be substantially increased in at least two ways. Releases from the Life Sciences plant at Hopewell might have been very high during some period coinciding with a high-flow period for the James River. And very high flows in the James and James estuary might transport Kepone already resident in bottom sediment into the Bay.

In either of those two cases, the substantial fishery services of the Bay might be in jeopardy. For the actual Chesapeake Bay (figure 2), only the lower portion of the Bay is at risk; there is relatively little mixing upwards toward Baltimore. But remember our distinction between the real and didactic Kepone incidents: had the incident occurred in a river emptying into the upper Bay, the whole of the Bay might have been at risk. In what follows, when we say Bay, we mean "that portion of the Bay mixing with flows from the relevant tributary."

The Bay is a unique natural asset providing a long list of service flows. Some, in principle, are excludable: among those are recreational and commercial fishery services. Others are more intricate: for example, the

Bay may serve as a hatchery for many of the species taken further north, perhaps as far north as the New England coast.

This service flow mix presents a challenge to applied cost-benefit analysis. In principle, the excludable service flows should be managed to maximize rents: were they so managed, we would have a measure of the loss suffered in an episode which destroys the Bay fisheries. Because the Bay fisheries are not so managed, such rents are (in principle) dissipated. Were catches from the Bay a negligible portion of the catch on the Eastern seaboard, and were capital in the Bay fisheries industries perfectly mobile, there would be no economic losses.

The issues are subtle: rather than face them here, we have, for our didactic purposes, chosen to summarize the situation with the assumption that the Bay is a unique natural asset to which a definite rental accrues. Destruction of the Bay fisheries in a didactic Kepone episode then implies a loss equal to the discounted present value of those rents. Such a loss is of course noninsurable.

ESTIMATING THE PROBABILITY OF HIGH-FLOW EVENTS

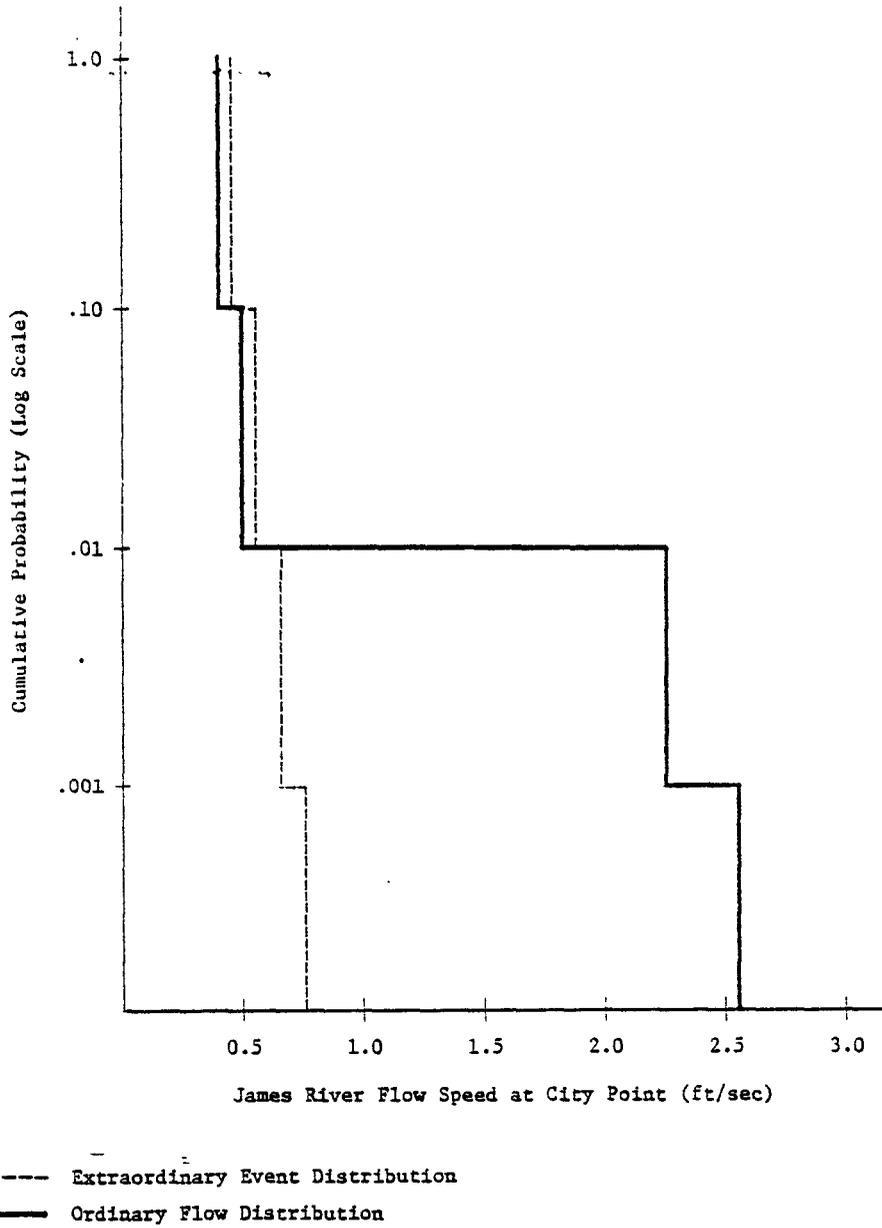
High-flow regimes may arise in two ways: from the high-flow tail of the usual distribution of flows generated by runoff from the watershed draining into the James River and estuary, or from extraordinary events exogenous to, and statistically independent of, that watershed runoff distribution.

First consider the ordinary events, the high-flow events associated with normal runoff from the James watershed. There is abundant data on flow in the James River and James River estuary tributaries: hydrologists call that time series the "historical trace." Among the central problems of the subdiscipline of theoretical hydrology is reconstructing, from the historical trace, the stochastic process of which that trace is one realization. To answer questions like "What is the probability of seven successive days of net freshwater inflow greater than 1,000 m^3/sec at Richmond?", we need the parameters of the stochastic process describing net inflow.

We have found that the following conclusion is insensitive to the method of reconstructing the James River flow distribution: the dominant contribution to the high-flow event "tail" of the flow distribution comes from the normal James flow distribution, and not from extraordinary events exogenous to that distribution, such as hurricanes. To see why, make the heroic assumption that extraordinary events like hurricanes are not reflected in the monthly-averaged flow data for the James. Then treat those monthly-averaged flows as observations drawn on independent, identically-distributed lognormal variables. Finally, estimate the parameters of those distributions. The solid-line cumulative distribution of figure 7 summarizes the results.

A similar distribution for hurricanes can be reconstructed from National Oceanographic and Atmospheric Administration estimates of a probability distribution for hurricane-related floods of various sizes along the Eastern seaboard. That flooding distribution can in turn be translated into a

Figure 7. Cumulative Probability Distributions for Ordinary and Extraordinary High-Flow Events



distribution of "transient" high-flow events on the James. To do this well, so-called flood models are necessary. For present purposes, simply assume that the level of the James is raised by the full level of the flood height, and that subsequent runoff occurs "normally" in a flow regime undisturbed by the storm. The broken-line cumulative distribution of figure 7 presents the resulting probability distribution for hurricane-related high-flow events. That figure indicates the dominance of the ordinary high-flow distribution: the five-hundred-year hurricane, for example, contributes an effective incremental four-day flow of less than one cubic foot per second.

We therefore proceed to a framework for computing the normal flow-distribution related contribution to Kepone transport into the Bay. For that purpose, we choose to work with model I of table 4 and with a "didactic" and simplified version of the actual Kepone case that is amenable to treatment by model I. Return to the schematic of figure 5: let us choose the following specific representation of that hydrogeography. We take the James estuary to be a constant width rectangular channel, and we place the possible source of secondary contamination--the location of Kepone in the bottom sediment--at thirty kilometers upestuary from the mouth of the James, the place at which the James empties into the Bay.

That geometry allows use of the Harleman, Holley, and Huber (1966) analytical solution for concentration ratios: we have

$$\frac{c(x,t)}{c_0} = \int_0^t \frac{U_f}{\sqrt{4\pi E(t-\tau)}} \exp \frac{-[x - U_f(t - \tau) + \frac{U_T}{\sigma} (\cos t - \cos \sigma)]^2}{4E(t - \tau)} d\tau \quad (4.1)$$

In that equation the variables and notations are as follows:

- x Distance from contaminant source (measured downestuary from the source)
- t Time
- c(x, t) Contaminant concentration at (x, t)
- U_f Advective channel flow velocity
- U_T Amplitude of tidal velocity oscillation
- Tidal frequency
- E Dispersivity
- c₀ Mixed contaminant concentration at source point
- I Mass injection rate of contaminant

The last two of these variables are related by the identity:

$$c_0 = \frac{I}{\rho A U_f} \quad (4.2)$$

where ρ is the density of water.

These parameters and variables, and their units, are listed in tables 6a and 6b below. Table 6a has an additional column for the parameter values used in our initial runs.

Specification of the initial experiments can be completed with two further sets of assumptions. We need an empirical relationship between advective velocity and the mass injection rate of the Harleman model. That relationship must capture the physical relationship between flow velocity and resuspension rates, perhaps the critical relationship for our purposes. And we need a probability distribution for advective velocities.

For the empirical relationship, we have taken

$$I = I_0 + I_1 U_f + I_2 U_f^2 \quad (4.3)$$

a simple quadratic in advective velocity. In our initial runs we will take values of the coefficients I_0 , I_1 , I_2 estimated from current best-practice estimates of Kepone transport into the Bay. This choice is central to the transport computation. The most obvious benefit of this exercise with a simple, "relatively analytical" transport model is that that sensitivity is highlighted and not disguised. In principle, a planner can run this model (very inexpensively) to see how bad things may be if his subjective estimates of the coefficients are bad. The coefficients themselves are "summary statistics" of the net effect of suspension and deposition.

For the probability distribution of James net freshwater flows, we begin with the available data, which gives twenty-two years of daily flows. We start by treating individual monthly flows as drawings on independent lognormal variables, using the twenty-two years of data to construct, for each month, the parameters of the corresponding lognormal distribution.

PRELIMINARY EXPERIMENTS IN THE SPIRIT OF THE HARLEMAN, HOLLEY AND HUBER ANALYTICAL MODEL

Because of the difficulties inherent in the numerical computation of the integral (4.1), we begin by using the injection-site concentration (4.2) to compute upper bounds on concentrations and transport. Note that the source-point mixed concentration depends on I , the mass injection rate, A , the estuary cross-sectional area, U_f the estuary advective-flow velocity, and ρ , the density of water. The last factor is a constant and of no further interest here. What we need is a way of constructing the probability distribution of the ratio on the right-hand side of equation (4.2). For Kepone mass transport is simply Kepone concentration times flow velocity, so that the probability distribution for mass transport can easily be constructed.

Table 6a. Parameters and Units, Initial Kepone Experiment Runs

Parameter	Units	Meaning	Assigned Values(s)
E	Square feet per second	Longitudinal dispersion	1.3×10^4
	Cycles per second	Tidal frequency	7.272×10^{-5}
U_F	Feet per second	Advective velocity	0.1 to 10.0
U_T	Feet per second	Tidal velocity	2.0
A	Square feet	Estuary cross-section area	1.0×10^4
	Kilograms per cubic foot	Density of water	28.3

Table 6b. Variables and Units, Initial Kepone Experiment Runs

Variable	Units	Meaning
I	Kilograms per second	Effective mass injection rate of contaminant
I0	Kilograms per second	Coefficients in equation describing the velocity-dependence of the mass injection rate
I1	Kilograms per foot	
I2	Kilogram-second per square foot	
CRATIO	Pure number	Ratio of contaminant concentrations at source observation points
KTRAN	Kilograms per hour	Intra-(tidal)cycle rate of Kepone transport into Bay
DKTRAN	Kilograms per year	Daily transport of Kepone into Bay
EKTRAN	Kilograms per year	Annual transport of Kepone into Bay
PROBUF	Pure number	Flow velocity probability distribution

Consider first the numerator of that ratio, the mass injection rate I . That parameter summarizes a critical relationship for estimates of Kepone transport into the Bay: the advective-velocity dependence of the Kepone resuspension rate. Because that relationship is both central to our transport calculation and uncertain, we want to explore the sensitivity of the transport estimate to our ignorance of advective velocity dependence. To do so, we will work with a family of distributions for I , with each member of that family depending upon U_f , and therefore having a frequency distribution generated by the frequency distribution of U_f .

The particular family we choose is constrained by two requirements. The first is that Kepone transport into the Bay under steady-state advective flow velocities typical of those observed in annual averages be of the same magnitude as observed Kepone transport rates into the Bay: between 10 and 100 kilograms per year. The second is that a range of "allocations" between the linear and quadratic terms in advective flow velocity U_f be explored. In particular, we choose a parameterized form of the equation (4.3). With I_0 taken equal to zero, define

$$I(L, U_f) = I_1(L) * U_f + I_2(L) * U_f^{**2} \quad (4.4)$$

$$L = 1, \dots, 10$$

where

$$I_1(L) = 1 - L \quad (4.5)$$

$$I_2(L) = 10^{**(-1)} * L$$

$$L = 1, \dots, 10$$

Now return to the remaining nonconstant factors, A and U_f , in equation (4.2). A , the cross-sectional area of the estuary, will of course depend upon flow velocity, since at higher flows the estuary level and cross-section area are higher. Thus, for a given value of U_f and a particular member of the family of distributions $I(L)$, the mixed concentration at the source point, c_o , defined by equation (4.2) is completely determined if the U_f -dependence of A is known. Further, from the probability distribution of U_f , the probability distribution of c_o can be derived.

As explained in the data appendix, available data on James flows gives daily discharge data over a twenty-year period, and data on cross-sectional area for about thirty of those days. From that latter data, the dependence $A(U_f)$ of cross-sectional area on advective flow velocity can be reconstructed: a regression of A on U_f and U_f^{**2} (without intercept) gives an embarrassingly high R^2 value.

After converting monthly discharge data to monthly flow data we can begin our construction of a probability distribution for James River flows. Recall how that construction works. Assume that each observed monthly flow represents a drawing on an underlying lognormal distribution, and that the monthly flows are independent lognormal variables. Then the twenty recorded

observations for each month give us an estimate of the mean and variance of the underlying lognormal distribution.

Having constructed those distributions, we are prepared to construct a distribution of annual transport of Kepone into the Bay. That construction proceeds as follows. First, we draw some large number of sequences of flows from those distributions: each of those sequences consists of twelve flow values, each one drawn on one of the monthly distributions. Then for each such sequence, begin by computing Kepone transport into the Bay, based on the assumption that transport proceeds at the rate of initial concentration c_0 given in (4.2). Finally, we use the set of all such sequences to impute (or estimate) the parameters of an effective lognormal distribution for annual Kepone transport into the Bay. All we are doing is moving the source point mixing concentration down the estuary and into the Bay. The justification for this procedure, which must substantially overestimate transport into the Bay and which violates our own judgment in favor of intratidal models, will be evident once we have the numerical results in hand. But in anticipation, we will see that the transport risk results are small enough, even when thus overestimated, to make any much more detailed computation pointless. Our strictures in favor intratidal models apply to the computation of ambient concentrations in the estuary, but not to transport into the Bay.

The results of those calculations are recorded in tables 7 and 8. In table 7, we present the effective lognormal distribution parameters for annual Kepone transport into the Bay: note that there are ten such sets of parameters, one for each value of L , the integer-valued index of the family of distributions corresponding to the function $I(U_p(L))$. We list the resulting monthly transports, given in terms of equivalent annual transport rates, in table 8. To convert the numbers of table 8 to actual monthly rates, simply divide by 12.

Finally, from the constructed annual Kepone transport distributions, we can get at the numbers we want: the probabilities that transport will exceed some critical value. The computation is a simple matter of computing the area in the tail of the corresponding lognormal distribution of annual transport. For the family of distributions indexed by L , and for four arbitrarily chosen critical values--50, 60, 70, and 80 kilograms annual transport--we obtain the results of table 9.

We can now move to actual "unweighted" damage estimates, since the probabilities of table 9 are the essential ingredient in those estimates. Remember, in what follows, that "unweighted" means unweighted by the initial release probabilities; that weighting is taken up in the following section.

Recall from the discussion of chapter 2 that the risk to the Bay is a risk to an asset which provides an essentially unique set of service flows. Put another way, inframarginal returns accrue to that natural asset, if properly managed, in the form of rentals; those rentals reflect the uniqueness of the service flows from the Bay.

Table 7. Lognormal Distribution Parameters, Annual Transport of Kepone into the Bay

L	Mean	Variance
1	0.298E 02	0.714E 00
2	0.338E 02	0.119E 01
3	0.379E 02	0.179E 01
4	0.420E 02	0.251E 01
5	0.460E 02	0.337E 01
6	0.501E 02	0.430E 01
7	0.541E 02	0.537E 01
8	0.582E 02	0.664E 01
9	0.623E 02	0.794E 01
10	0.663E 02	0.948E 01

Table 8. Equivalent Annual Kepone Transport Rates by Month, in Kilograms

Month	L = 1	L = 2	L = 3	L = 4	L = 5
1	0.304E 02	0.346E 02	0.389E 02	0.431E 02	0.474E 02
2	0.306E 02	0.349E 02	0.392E 02	0.435E 02	0.478E 02
3	0.315E 02	0.361E 02	0.407E 02	0.453E 02	0.499E 02
4	0.304E 02	0.346E 02	0.389E 02	0.431E 02	0.474E 02
5	0.299E 02	0.340E 02	0.381E 02	0.422E 02	0.463E 02
6	0.297E 02	0.337E 02	0.378E 02	0.418E 02	0.458E 02
7	0.288E 02	0.326E 02	0.363E 02	0.401E 02	0.439E 02
8	0.289E 02	0.327E 02	0.365E 02	0.403E 02	0.441E 02
9	0.288E 02	0.326E 02	0.364E 02	0.401E 02	0.439E 02
10	0.293E 02	0.332E 02	0.372E 02	0.411E 02	0.450E 02
11	0.293E 02	0.332E 02	0.371E 02	0.410E 02	0.449E 02
12	0.298E 02	0.338E 02	0.379E 02	0.419E 02	0.460E 02

Month	L = 6	L = 7	L = 8	L = 9	L = 10
1	0.516E 02	0.558E 02	0.601E 02	0.643E 02	0.685E 02
2	0.521E 02	0.564E 02	0.607E 02	0.650E 02	0.693E 02
3	0.545E 02	0.590E 02	0.636E 02	0.682E 02	0.728E 02
4	0.516E 02	0.558E 02	0.601E 02	0.643E 02	0.685E 02
5	0.504E 02	0.545E 02	0.586E 02	0.627E 02	0.668E 02
6	0.499E 02	0.539E 02	0.579E 02	0.619E 02	0.660E 02
7	0.476E 02	0.514E 02	0.552E 02	0.589E 02	0.627E 02
8	0.479E 02	0.517E 02	0.555E 02	0.593E 02	0.631E 02
9	0.477E 02	0.515E 02	0.552E 02	0.590E 02	0.628E 02
10	0.489E 02	0.528E 02	0.568E 02	0.607E 02	0.646E 02
11	0.488E 02	0.527E 02	0.566E 02	0.605E 02	0.644E 02
12	0.500E 02	0.541E 02	0.581E 02	0.622E 02	0.662E 02

Table 9. Probabilities of Annual Kepone Transport Greater Than Specified Critical Values, for the Indexed Family of Distributions; Annual Kepone Transport in Kilograms

Critical Level	L = 1	L = 2	L = 3	L = 4	L = 5
P > 50	0.0	0.0	0.596E-07	0.161E-05	0.180E-01
P > 60	0.0	0.0	0.0	0.0	0.596E-07
P > 70	0.0	0.0	0.0	0.0	0.0
P > 80	0.0	0.0	0.0	0.0	0.0

Critical Level	L = 6	L = 7	L = 8	L = 9	L = 10
P > 50	0.508E 00	0.967E 00	0.100E 01	0.100E 01	0.100E 01
P > 60	0.578E-05	0.769E-02	0.239E 00	0.786E 00	0.984E 00
P > 70	0.596E-07	0.596E-07	0.138E-04	0.448E-02	0.118E 00
P > 80	0.0	0.0	0.596E-07	0.596E-07	0.241E-04

For that reason, the proper method of valuing risks to that asset must be based upon an analogy with a risk-averse individual: society must decide how cautious it wants to be in placing that asset at risk by permitting the operation of facilities which may damage that asset beyond repair, for long periods or in perpetuity. Because the asset is unique, the loss is essentially noninsurable (meaningful contracts for compensation could not be made), and society can only choose (ex ante) how much prevention it wants, to buy. The maximum amount society is willing to pay for prevention is essentially the benefit associated with avoiding the risk.

Thus the value of the risk in question depends upon how risk averse society chooses to be. But that does not mean that "nothing can be said": it is helpful to exhibit the values implied by specific degrees of risk aversion, and thereby to inform judgment. Thus if society effectively values risks to the present value R of Bay-derived rentals by using a logarithmic utility function $\ln(W_0 + R)$ with initial wealth parameter W_0 , then willingness to pay V to avoid a risk threatening destruction of the fishery service flows of the Bay will be defined implicitly by the equation

$$\ln(W_0 - V) = (1 - p)\ln(W_0 + R) + p\ln(W_0) \quad (4.6)$$

The V so defined will of course be a function $V(W_0, R, p)$ of W_0 , R, and p. The interpretation of W_0 is initial wealth; heuristically, it plays the role of a risk-aversion parameter, translating social aversion to the particular risk involved into an equivalent initial wealth. The logarithmic utility of wealth function is particularly convenient because it admits easy solution for V as

$$V(W_0, R, p) = (W_0) \left(1 - \left(1 - \frac{R}{W_0} \right)^p \right) \quad (4.7)$$

Some numbers may help here: table 10 below presents values of the "risk-aversion correction factor" multiplying W_0 . For low values of p (low probabilities), and for values of R/W_0 small compared to 1 (losses small compared to total wealth), V differs little from expected loss pR. For larger values of p (high probabilities) and higher values of R/W_0 (losses which are a significant fraction of total wealth), V can be substantially greater than expected loss pR.

Now return to our derived distribution for annual transport into the Bay. Introduce two additional assumptions: that the relevant total wealth (risk-aversion) parameter is \$10 billion, and that the extreme high-flow sensitivity distribution (L = 10) is the relevant one. Then, letting the annual rental R derived from the Bay range from \$1.0 billion to \$10.0 billion, we obtain the results of table 11. Suppose, for example, that we have determined that the annual rentals accruing to the service flows provided by the Bay amount to \$8 billion, and that Kepone transport into the Bay of 80 kilograms or more in any one year will effectively destroy those service flows in perpetuity. Then the cost of risk as a fraction of initial wealth is 0.119E-06. When multiplied by initial wealth W_0 equal to \$10 billion, the cost of risk is found to be \$119, an insignificantly small

Table 10. Some Representative Values of the Collective Risk Adjustment Factor

Loss as Fraction of Initial Wealth (R/W_0)	Probability of Loss (p)			
	10^{-3}	10^{-4}	10^{-5}	10^{-6}
.1	1.053×10^{-4}	1.055×10^{-5}	1.072×10^{-6}	1.192×10^{-7}
.2	2.231×10^{-4}	2.229×10^{-5}	2.205×10^{-6}	2.384×10^{-7}
.5	6.929×10^{-4}	6.932×10^{-5}	6.914×10^{-6}	7.153×10^{-7}
.9	2.300×10^{-4}	2.302×10^{-5}	2.301×10^{-6}	2.324×10^{-7}

Table 11. Fractional Cost of Risk, for Several Critical Values and Annual Rental Values

Critical Level	R = 1	R = 2	R = 3	R = 4	R = 5
P > 50	0.0	0.0	0.0	0.834E-06	0.124E-01
P > 60	0.0	0.0	0.0	0.0	0.596E-07
P > 70	0.0	0.0	0.0	0.0	0.0
P > 80	0.0	0.0	0.0	0.0	0.0

Critical Level	R = 6	R = 7	R = 8	R = 9	R = 10
P > 50	0.372E 00	0.688E 00	0.800E 00	0.900E 00	0.100E 01
P > 60	0.530E-05	0.921E-02	0.319E 00	0.837E 00	0.100E 01
P > 70	0.596E-07	0.596E-07	0.222E-04	0.103E-01	0.100E 01
P > 80	0.0	0.0	0.119E-06	0.119E-06	0.100E 01

figure. Again, we remind the reader that this is "unweighted": weighting by an initial release probability will even further depress this estimate.

COST OF HEALTH RISKS IN A DIDACTIC KEPONE INCIDENT: UNWEIGHTED AND RELEASE-PROBABILITY WEIGHTED ESTIMATES

Now let us turn to some brief consideration of a damage category that, while not relevant for the actual Kepone incident, is a possibility in similar incidents: human health damage. To emphasize the status of these estimates, we refer to a "didactic" Kepone incident. Returning to figure 4.2, the relevant consequence category here is C3. And to emphasize that our first calculation is done assuming a release has already occurred, we call these first estimates "unweighted" estimates.

These calculations are relatively simple. We need some assumptions: that individual exposure to drinking water drawn from the contaminated river is 1 liter per day, or 365 liters over the duration of the period before contamination is detected; that period is taken as one year. Finally, we assume a population at risk of 100,000 individuals.

Then we can readily compute, for each of our family of distributions indexed by L (L = 1, ..., 10), the mean annual concentration of Kepone in the river during that one year. Those concentrations, in parts per billion, are recorded as the first row of table 12 below.

To go from concentrations to the costs, to individuals, of bearing the incremental mortality risk is straightforward: compute exposure, then multiply by a dose-response coefficient and by a value of incremental mortality risk estimate. Because both of those multipliers are subject to some uncertainty, we begin by taking a range of values for each. We let the dose-response multiplier, which has the units of incremental annual mortality risk (over one's lifetime) per parts per billion incremental annual exposure (over one's lifetime), take the value $10^{(-4)}$ and $10^{(-7)}$. Similarly, we let the incremental mortality risk, which has the units of dollars per incremental annual mortality risk (over one's lifetime), take on the values $10^{(4)}$, $10^{(5)}$, and $10^{(6)}$. Finally, computation of an aggregate health risk figure requires only that we multiply the individual health risk figures by the population at risk, assumed here equal to $10^{(5)}$. The resulting entries give the undiscounted present value of health risk damages, in dollars.

Those estimates are still "unweighted" by release probabilities, and must be so weighted to arrive at what we have called the standard theory estimates of the cost of risk bearing. For illustrative purposes, let us select what we believe are plausible values of incremental annual mortality risk, 10^{-5} , and value of incremental annual mortality risk, 10^{+5} . Then the unweighted cost of risk estimates lies between $\$0.131 \times 10^{+11}$ and $\$0.289 \times 10^{+11}$; that range is associated with our range of uncertainty about the flow-velocity dependence of Kepone concentrations in the "didactic" James.

Again for illustrative purposes, suppose that individuals assign the subjective probability 10^{-4} to the event "a major release of Kepone into the

Table 12. Undiscounted Present Value of Health Risk Damages, Didactic Kepone Incident, in Dollars

	L	1	2	3	4	5	6	7	8	9	10
MEAN CONCENTRATION		0.358E 00	0.406E 00	0.454E 00	0.502E 00	0.550E 00	0.598E 00	0.646E 00	0.695E 00	0.743E 00	0.791E 00
INDIVIDUAL HEALTH RISK											
RISK VALUE											
10E-7 10E4	0.131E 02	0.148E 02	0.166E 02	0.183E 02	0.201E 02	0.218E 02	0.236E 02	0.254E 02	0.271E 02	0.289E 02	0.289E 02
10E-7 10E5	0.131E 03	0.148E 03	0.166E 03	0.183E 03	0.201E 03	0.218E 03	0.236E 03	0.254E 03	0.271E 03	0.289E 03	0.289E 03
10E-7 10E6	0.131E 04	0.148E 04	0.166E 04	0.183E 04	0.201E 04	0.218E 04	0.236E 04	0.254E 04	0.271E 04	0.289E 04	0.289E 04
10E-6 10E4	0.131E 03	0.148E 03	0.166E 03	0.183E 03	0.201E 03	0.218E 03	0.236E 03	0.254E 03	0.271E 03	0.289E 03	0.289E 03
10E-6 10E5	0.131E 04	0.148E 04	0.166E 04	0.183E 04	0.201E 04	0.218E 04	0.236E 04	0.254E 04	0.271E 04	0.289E 04	0.289E 04
10E-6 10E6	0.131E 05	0.148E 05	0.166E 05	0.183E 05	0.201E 05	0.218E 05	0.236E 05	0.254E 05	0.271E 05	0.289E 05	0.289E 05
10E-5 10E4	0.131E 04	0.148E 04	0.166E 04	0.183E 04	0.201E 04	0.218E 04	0.236E 04	0.254E 04	0.271E 04	0.289E 04	0.289E 04
10E-5 10E5	0.131E 05	0.148E 05	0.166E 05	0.183E 05	0.201E 05	0.218E 05	0.236E 05	0.254E 05	0.271E 05	0.289E 05	0.289E 05
10E-5 10E6	0.131E 06	0.148E 06	0.166E 06	0.183E 06	0.201E 06	0.218E 06	0.236E 06	0.254E 06	0.271E 06	0.289E 06	0.289E 06
10E-4 10E4	0.131E 05	0.148E 05	0.166E 05	0.183E 05	0.201E 05	0.218E 05	0.236E 05	0.254E 05	0.271E 05	0.289E 05	0.289E 05
10E-4 10E5	0.131E 06	0.148E 06	0.166E 06	0.183E 06	0.201E 06	0.218E 06	0.236E 06	0.254E 06	0.271E 06	0.289E 06	0.289E 06
10E-4 10E6	0.131E 07	0.148E 07	0.166E 07	0.183E 07	0.201E 07	0.218E 07	0.236E 07	0.254E 07	0.271E 07	0.289E 07	0.289E 07
TOTAL HEALTH RISK											
RISK VALUE											
10E-7 10E4	0.131E 08	0.148E 08	0.166E 08	0.183E 08	0.201E 08	0.218E 08	0.236E 08	0.254E 08	0.271E 08	0.289E 08	0.289E 08
10E-7 10E5	0.131E 09	0.148E 09	0.166E 09	0.183E 09	0.201E 09	0.218E 09	0.236E 09	0.254E 09	0.271E 09	0.289E 09	0.289E 09
10E-7 10E6	0.131E 10	0.148E 10	0.166E 10	0.183E 10	0.201E 10	0.218E 10	0.236E 10	0.254E 10	0.271E 10	0.289E 10	0.289E 10
10E-6 10E4	0.131E 09	0.148E 09	0.166E 09	0.183E 09	0.201E 09	0.218E 09	0.236E 09	0.254E 09	0.271E 09	0.289E 09	0.289E 09
10E-6 10E5	0.131E 10	0.148E 10	0.166E 10	0.183E 10	0.201E 10	0.218E 10	0.236E 10	0.254E 10	0.271E 10	0.289E 10	0.289E 10
10E-6 10E6	0.131E 11	0.148E 11	0.166E 11	0.183E 11	0.201E 11	0.218E 11	0.236E 11	0.254E 11	0.271E 11	0.289E 11	0.289E 11
10E-5 10E4	0.131E 10	0.148E 10	0.166E 10	0.183E 10	0.201E 10	0.218E 10	0.236E 10	0.254E 10	0.271E 10	0.289E 10	0.289E 10
10E-5 10E5	0.131E 11	0.148E 11	0.166E 11	0.183E 11	0.201E 11	0.218E 11	0.236E 11	0.254E 11	0.271E 11	0.289E 11	0.289E 11
10E-5 10E6	0.131E 12	0.148E 12	0.166E 12	0.183E 12	0.201E 12	0.218E 12	0.236E 12	0.254E 12	0.271E 12	0.289E 12	0.289E 12
10E-4 10E4	0.131E 11	0.148E 11	0.166E 11	0.183E 11	0.201E 11	0.218E 11	0.236E 11	0.254E 11	0.271E 11	0.289E 11	0.289E 11
10E-4 10E5	0.131E 12	0.148E 12	0.166E 12	0.183E 12	0.201E 12	0.218E 12	0.236E 12	0.254E 12	0.271E 12	0.289E 12	0.289E 12
10E-4 10E6	0.131E 13	0.148E 13	0.166E 13	0.183E 13	0.201E 13	0.218E 13	0.236E 13	0.254E 13	0.271E 13	0.289E 13	0.289E 13

James"; for specificity, think of a relatively large and sudden release. Then the cost of risk-bearing estimates, when weighted by this factor, lie in the range $\$0.131 \times 10^7$ to $\$0.289 \times 10^7$. That range may be roughly translated into the statement: an annual expenditure of about one million dollars is warranted if it can prevent an incident like Kepone contamination of our "didactic" James River. Discounting that stream of expenditures at 10%, that statement becomes: a capital expenditure of about ten million dollars is warranted if it can prevent an incident like Kepone contamination of the James River.

These are large numbers, and would almost certainly justify a rather substantial and rigorous monitoring, or early warning, system for "novel" and low-level contaminants of drinking water.

IMPLICATIONS AND CONCLUSIONS

Now let us look back over the calculations we have done. We computed two kinds of damages: those associated with a sudden transfer of Kepone into the Bay, and those associated with human consumption through the drinking water exposure route. We take those up in turn.

The damages associated with a sudden transfer of Kepone into the Bay are surprisingly small. The numbers are so small that we feel content to rest with the very rough method, bound to give overestimates, which we used. To recall only the essentials of that method, we assumed that source concentrations are transported undiminished downstream and into the Bay, and that there is effectively no process of deposition and resuspension. The critical sensitivity, the sensitivity of rate at which Kepone in bottom sediment at the source point is resuspended to flow velocity, is captured, in a very rough and ready way, by our method. The relatively small resulting variation in Kepone transport into the Bay is tribute to the relatively small variation in flow velocity. That is true for both the normal flow distribution and the extraordinary flow distribution, where by the latter we mean the distribution associated with hurricanes, floods, and the like.

The damages to human health that can result from a Kepone-type incident in a water body from which drinking water supplies are drawn are very much larger. Again, there is a more or less simple reason. If Kepone is reasonably carcinogenic, and if several years elapse from the time at which such exposures begin to the time at which taking drinking water from the river is halted, then the risk can be quite large. And thus the valuation of the risk posed by those exposures can be quite large.

The cost of risk figures are in fact so large that they would justify a substantial program for monitoring drinking water supplies for novel and possibly dangerous pollutants. Similarly, they would justify substantial expenditures to guarantee that plants situated on waterbodies linked to drinking water supplies do not contaminate those waterbodies with low levels of dangerous materials over long periods. This is a kind of monitoring that is different in character from the kind generally in place now: typically, we measure the levels of the more familiar bacterial and mineral components of drinking water.