

THE PUBLIC SERVANT AS GOD: TAKING RISKS WITH THE PUBLIC

Notes for an Address

by

A. R. Dobell

at the

National Conference on Management in the Public Sector
Victoria, British Columbia
April 22, 1986

This address is based on a forthcoming working paper jointly prepared with Ted Parson, whose original contribution is gratefully acknowledged. Responsibility for any errors rests with me.

You Can't Avoid Playing God

When considering our own acts, most of us acknowledge that we live in an uncertain world and that all activity involves some risk. It is clear that through even the most commonplace decisions we impose risks on ourselves and others. Sometimes these risks are so small and ordinary that we ignore them, as we do with the risks of driving to the shopping centre or considering different modes of transport for a business trip; at other times we acknowledge and weigh them explicitly, as in deciding whether or not to take up skydiving or to increase life insurance coverage. But acknowledged or not, risks are always present.

Public policy decisions are also commonly accompanied by risk. This is not surprising, for there is no reason to expect that decisions made collectively should be exempt from the risk consequences that characterize private decisions. Moreover, even if an issue of public risk is not initially an area of government activity, it is likely to become one because of the strength of the public's perception that one of government's duties is to protect them from danger. Government is naturally drawn into any area where there exists a real or perceived public risk, in order to regulate it; government is then necessarily involved in the subsequent decisions of how and how much to reduce risk, and at what cost.

There is much resistance to the idea that government decisions impose risks on people. Public servants or politicians may respond with indignant denial if accused of taking risks with the health or safety of the public. Perhaps they do so because to admit otherwise would involve too great a personal risk. But a thoughtful examination reveals that many decisions in the life of a public servant do just that. Even if we neglect economic risks and consider only risks to life and limb, areas of public policy as diverse as public health, environmental regulation, product safety, occupational safety, incomes policy, investment in industrial facilities, criminal justice, and foreign policy all have risky consequences.

The truth of this claim is most evident at the political level, where major policy directions are determined. But it is just as true that managers in the public service routinely make decisions involving public risk, both in their roles as advisors of ministers and as administrators of programmes and policies.

When a city official decides whether to install new street lighting; when a highways official chooses on which of several possible improvement projects to spend a limited budget;

when a parole officer decides on the status of a released prisoner; when an environmental official chooses how forcefully to act against a factory violating emissions standards; or when a provincial official considers whether to permit herbicide spraying on hydro rights of way, all are making decisions with direct consequences for public risk.

Since you cannot avoid taking risks with the public, you presumably should ask how you can best do it in a rational, fair, and defensible manner. In this paper, I hope to offer some insight into why public decisions on risk are often so difficult and so divisive. The difficulties are of several kinds, with their roots in the ambiguous nature of individual decisions, in the environment of scientific uncertainty that often prevails, and in the difficult distributive issues that are raised. I will also offer some tentative suggestions as to how decisions on risk might be improved. For those readers who wish to pursue this subject in more detail, a forthcoming IRPP working paper, prepared with Ted Parson, will offer a fairly substantial literature survey.

Individual Decisions on Risk

The most reasonable place to begin an enquiry into public decisions on risk is with an examination of how individuals make risky decisions on their own behalf. Amazingly enough, this brings us into severe trouble at the outset. There exist two large bodies of research into individual decision-making on risk. One of them is prescriptive, and has its roots in economics and statistical decision theory; it extends the notion of economic logic to uncertain situations, and proposes axioms governing how an ideally rational person should act in the face of risk. The other is descriptive, and has its roots in psychology; it begins with experimental observations on how people actually make decisions under uncertainty. The difficulty is that these two approaches yield results that are fundamentally in conflict. The normative theory of how people should behave (or of how ideally rational economic people do behave) does not even approximate how people really do behave in experimental situations.

The normative theory represents an attempt to extend the definition of the economic person to situations of uncertainty. The economic man or woman, you will recall, is the one who knows precisely what he or she wants, and is willing and able to undertake the calculations necessary to make the best choices possible in pursuit of those known objectives, within the relevant constraints such as a limited income or limited resources more generally.

When the world is risky, the economic person must make choices not knowing what the outcomes will be, but having only (at best) estimates of the probabilities of various outcomes. The theory of rational choice under uncertainty dictates how such a person should weigh evidence in order to estimate probabilities, and how he or she should choose an action once probability estimates are derived. Two cornerstones of the theory are the assumptions that only the values of various outcomes--the actual consequences--and their probabilities of occurrence should matter, and that preferences should be linear in probabilities. (This means that if a certain event is desirable, you should rate a 20% chance of that event occurring as exactly twice as desirable as a 10% chance.) If you follow the prescribed rules, you will maximize your welfare in the long run. This theory could be described as a theory of ideal gambling.

It offers an elegant and attractive theory, particularly because violation of any of its axioms can be shown to lead to inconsistent or ridiculous choices. But unfortunately, the evidence of psychological experiments points very strongly to the conclusion that people do not act this way. And the ways in which people violate this code of behaviour are not minor and haphazard, but large and systematic.

The principal workers in this second field of research have been the psychologists Daniel Kahneman of UBC and Amos Tversky of Stanford. They have performed a large number of experiments in which they ask people to estimate probabilities or to make choices between risky propositions when the probabilities are given. Again and again they have found that even the most basic requirements for the normative theory of decision-making on risk are violated by people's real choices.

Most crushing for the normative theory has been their observation that people's preferences depend strongly on what they call "framing", the way that risky propositions are expressed. For example, they find that most people, if forced to choose between a certain loss of \$50 and a 25% chance of losing \$200, prefer to chance the \$200 loss; but when the certain loss of \$50 is called an insurance premium, most people prefer it.

In another experiment more directly related to risks to life and limb, they asked a large number of physicians whether they would recommend treating lung cancer in a 45-year-old man with radiation or with surgery. To one group they showed probabilities of dying during treatment and of dying within 5 years--e.g., with radiation you have a 0% chance of dying during treatment but a 78% chance of dying within 5 years, while with surgery you have a 10% chance of dying

on the operating table but only a 66% chance of dying within 5 years. To the other group they showed the same results expressed in terms of probability of survival, i.e., 100% chance of surviving radiation treatment, etc. The shocking result was that the fraction of physicians preferring radiation to surgery was strongly dependent on which of these two forms of expression was used. Only 16% of the doctors recommended radiation treatment when the figures were expressed in terms of survival, while 50% recommended it when given mortality figures. This influence of framing is particularly remarkable when you consider that these are experts making a professional judgment on a potentially critical matter falling squarely within their domain of specialization.

Other experiments show other strong and systematic violations of the normative theory. One example is that people apparently cannot properly handle small probabilities, say of one in a thousand or smaller. Our evaluation of such probabilities shows a strong threshold effect. Until something brings it to our attention, we tend to ignore an event with such a low probability, that is, we treat it as if its probability were zero. But once we have acknowledged it as possible, we then greatly over-estimate its likelihood. An example is the widely observed tendency of people to over-estimate the probability of certain rare and dramatic causes of death, such as botulism or radiation exposure. This phenomenon has been demonstrated by Slovic, Fischhoff and Lichtenstein, who asked people to estimate the number of fatalities per year in the US due to various causes, after informing them that 50,000 per year die in auto accidents. Their results show that rare deaths are over-estimated by factors as high as several hundred--botulism, for example, which kills 2 people per year in the US, is estimated to kill 500--while the most frequent causes of death are all under-estimated. The same bias in a different context has been found by Bill Ziemba of UBC in his studies of betting at racetracks. He finds that people bet too often on long-shots and not often enough on favorites, and has exploited this market imperfection to develop an effective betting system.

Kahneman and Tversky have argued that we assess probabilities through the use of heuristic devices, or rules of thumb, that are biased or irrelevant. They name three heuristics: availability, representativeness, and anchoring. We use the availability heuristic when we estimate the probability of an event by the vividness of our images of it or the ease with which instances of it can be called to mind. It is exemplified by the higher probability that people assign to airline crashes immediately after a large well-publicized crash.

The representativeness heuristic causes us to estimate the probability of an event by how representative it is of the class of events to which it belongs. We are employing this heuristic when we count the observation of a license plate number 22222 as less likely than 43781, or when we succumb to the "Gambler's Fallacy" and expect that a tossed coin is more likely to come up heads after a long run of tails.

Anchoring is the process by which people select an initial value for their estimate of a probability or other unknown quantity, and then incorporate further information by making small adjustments to their initial value. The initial guess continues to exercise a significant influence on the final estimate, even in the presence of strong opposing evidence.

The anchoring heuristic can be regarded as a tendency to place too much confidence in the first guess, and consequently is related to the often-observed over-confidence that lay subjects and experts alike place in their estimates of unknown quantities. In an experimental demonstration of this overconfidence, researchers at MIT gave seven engineers all the relevant technical information about an earthen embankment, then asked them to estimate the height at which it would fail. Two of the engineers substantially over-estimated, and five under-estimated, the height of failure, but more remarkably, not one gave a confidence interval wide enough to contain the true (experimentally measured) value.

Perception of risk has also been shown to depend on several other factors including the credibility of the source of information, whether or not the risky activity is undertaken voluntarily, and the distribution of risks. For many sources of risk, the public does not have access to authoritative information, but must learn about the risk from newspapers, television, government announcements, or other intermediate sources. The way people view any particular risk depends in some complex way on the perceived credibility of the source of information. An Ontario Gallup poll concerning credibility of sources of information about nuclear waste disposal found that university scientists and the media were rated the most credible, while government sources were rated rather low in credibility. An official of the Ontario Ministry of the Environment has pointed out that there is a perverse asymmetry in the public's reaction to information from the Ministry; bad news is believed, but good news is not. The principle of the media that bad news makes more dynamic copy than good may contribute to this phenomenon. Another curious asymmetry in the public's attitudes to government is pointed out in a 1984 University of Waterloo study on hazardous waste transport. People generally perceive

government regulatory bodies as having been too cozy with the industries they regulate and having failed in the past to protect adequately the public interest. Yet their expectations of government's protecting their interests in the future, and their assessment of government's power to do so, are very high.

Chauncey Starr was the first to emphasize that whether or not an activity is undertaken voluntarily is an important determinant of what associated risk will be regarded as acceptable. He hypothesized that people will accept risks approximately a thousand times greater from voluntary activities (smoking, skiing, hang-gliding) than from involuntary sources (air and water pollution, pesticides in food, living near a refinery).

More recent surveys suggest that many factors other than voluntarism are also relevant. Slovic et al surveyed how the overall risk level of a number of sources was related to other perceived attributes, and found that total perceived risk could be accurately predicted on the basis of three measures: "Novelty", which increased as the risk was less similar to known risks, less evident to those exposed, and more immediate in its effects; "dread", which increased as the risk was less controllable, less equitable in effect, more catastrophic, involuntary and fatal, and as the risks were borne more by children and future generations; and "extent", the number of people exposed.

Kahneman and Tversky have developed a consistent descriptive theory of people's choices that accommodates these results, called Prospect Theory. These two theories of people's decision-making under risk, the normative Bayesian decision analysis and the predictive prospect theory, are unquestionably in conflict. People do not act as the theory of rational choice says they should. Furthermore, it is unlikely that any other theory will be found to reconcile the two, for any theory to be descriptively valid must reflect the dependence of people's preferences on framing, but surely no normative theory could be susceptible to framing. That is, no rational decision process could let the decision rest on the way it is described rather than on the objective outcome or characteristic features of the choice.

This conflict between rationality and people's observed behaviour must, however, be treated with caution, particularly if one's reaction is to conclude that people do not know what is best for them and need to have their risk choices made for them by scientists and experts. The first caution is that scientists and experts appear to be as prone to these errors of judgment as anyone else. Much of Tversky and Kahneman's experimental work was performed with scientists

and physicians as subjects, and the biases and heuristic devices described here have been observed to be employed by lay people, by scientists making judgments outside their area of expertise, and by scientists in their own fields when the data presented was insufficient to calculate a precise solution. Moreover, the problem of over-confidence in predictions has been observed to be the most pronounced in the judgements of experts. We are clearly all prey to these biases, errors, and misperceptions.

At the individual level, the conflict between these two models of choice merely suggests that if we wish to act like economically rational people, we must police our own personal decisions on risky matters very carefully. But for the public servant with official responsibility for making decisions that will impose risks on others, the conflict makes for a profound dilemma: should you choose for others as they would choose for themselves, knowing that such a choice does not represent their best interests, or should you choose as an ideally rational person would, knowing that the affected people may oppose the choice that you are making on their behalf?

Problems of Collective Choice

As if the foregoing were not problem enough, there are several other aspects to public decisions on risk that make them even more difficult. Many public decisions on matters of risk-- certainly most of the contentious ones-- must be made with very poor information. For certain issues such as the siting of energy facilities and the transport of hazardous goods, the reason is that the resultant risk is principally due to the chance of disastrous accidents, whose probability of occurrence is very low. The total public risk from a decision may depend very strongly on whether a particular kind of accident will happen on average once per thousand years or once per hundred thousand, but there is no reliable way to distinguish between these two values. Theoretical attempts to do so are filled with arbitrary and questionable assumptions, while attempts to do so on the basis of evidence are thwarted by the fact that the small number of occurrences of such events to date are consistent with a wide range of true probabilities. On observing the first occurrence of such an event, one does not know whether one has simply

observed a rare event, or has just seen cause to revise the probability estimate dramatically upward.

For other decisions such as whether or not to allow herbicide spraying on hydro rights-of-way, the size of the risk may depend on the toxicity of low public exposures to a hazardous substance. Information on the magnitude of these hazards is also hard to obtain. Evidence on the toxicity of chemicals usually comes from epidemiological studies of large human populations, in which a statistically significant increase in the incidence of cancer, for example, is sought in the group with higher exposure to a chemical. Such studies are limited in their ability to detect other than the most acute effects by the difficulty of finding two groups whose only significant difference is their level of exposure, and the difficulty of distinguishing a true underlying increase in the cancer rate of one group from random fluctuations in the incidence among groups. Improvement in our knowledge on these issues through experiment is usually not available. The experiments that are performed, in which large numbers of animals are exposed to high doses of the suspected hazard, are costly, slow, and possibly of limited relevance because of the problems of drawing inference from animals to humans and from large doses to small ones. If animal rights activists prevail, even this imperfect source of information may cease to be available. Experiments on humans would necessarily involve exposing a group of people to a suspected hazard, and so are ethically and legally proscribed.

With so little opportunity to settle disagreements through evidence, it is not surprising that there exists vigorous scientific disagreement on questions of risk. The common tendency to exaggerated confidence in the accuracy of one's predictions reported above exacerbates such disagreement. So does the pressure that the political use made of the answers to such trans-scientific questions creates for advocacy research.

Obscure and complicated scientific questions also provide ample opportunity to project one's values into one's scientific opinions. A study by Manne and Richels provides an interesting example. In a survey of the opinions of American scientists on a variety of issues in energy policy, they found a very strong correlation between low estimates of Uranium reserves, high future demand growth for electricity, and high costs for photovoltaic electric cells. The correlation between such apparently unrelated opinions was apparently explained by the fact that these scientists favored the development of the breeder reactor. Those who opposed the breeder had an equally strong clustering of their opinions around the opposite extremes. In each case, it

appears that a policy opinion on the desirability of the breeder had organized a disparate collection of opinions on issues of fact to support it.

One topical example of all of these phenomena is the current controversy over the Federal government's financing the removal of urea formaldehyde foam insulation from homes in which it had financed the installation a few years ago. Apparently, no one knows how dangerous this product really is. When it was in the homes, people's fears over its possible carcinogenicity and other health effects led to a public outcry that the government pay for its removal. Now that the government is doing so, it is being criticized for wasting the taxpayer's money. These events illustrate that where risk to the public is concerned, you can be denounced for being too careful as well as for being careless. And this is so even if no one can tell for certain what level of care is appropriate.

It has been argued that the government is right to pay for the removal of the insulation even if it represents no significant risk, because people's subjective perception of the risk has caused large objective losses in property value of the affected homes. But this argument does not resolve the problem, for if the market is merely reflecting people's exaggerated fears, such an argument would suggest that public decisions should be made on the basis of perceived risk, even when the perceptions are incorrect. And that does not seem defensible.

Public decisions that have both an economic component and a risk component always involve the assignment of a monetary value to a human life. This assignment is rarely done explicitly, but it is always there. Whenever you decide how much money to spend on the reduction of a risk, you have implicitly assigned a value to life. A number of studies have been done comparing the levels of spending on risk reduction in a number of programmes or ministries, in order to find the implied values of lives. These studies have found enormous ranges in the implied value of a life, from a few thousand dollars (in automobile safety, for example) to tens of millions or even billions (in the safety features of nuclear power plants, for example). While a strong argument can be made that disparities of this order are too high, there is a great deal of controversy over the question of what an appropriate value for a life is.

A large academic literature attempts to address this question, either by evaluating the loss of earnings resulting from an early death or by examining individuals' willingness to pay for small reductions to the risk they face. Studies observing the market for risky occupations have calculated values for a life ranging from \$200,000 to \$2 million.

Such studies purport to measure individuals' willingness to pay to save their own life. But when risks and benefits accrue to different people, as they almost always do, we must address questions such as, should some be made less safe in order that others be made safer? By how much? Or, more poignantly, should some be made less safe in order that others can be made wealthier? By how much?

A utilitarian would argue that to maximize total benefit we must assign an appropriate value to a life and take all measures available to save lives at less than this cost. But perhaps not all lives should be valued equally; perhaps the lives that are most at risk should be valued highest. This stance corresponds to the increase in our willingness to pay for risk reduction as the risk to our own life increases, and also reflects the social values inherent in our often-observed willingness to spend far more money (and even to put additional lives at risk) in order to rescue particular individuals from peril than we are to avert a peril before the victims are known. Alternatively, perhaps we should be willing to pay more to save a young life than an old one.

The difficulties discussed so far are particularly associated with decisions on risk. But decisions on risk are also beset by a host of other difficulties that are common to a broad range of political decisions. Issues of risk usually involve unequal distribution of risk and monetary costs or benefits to different people, and so create winners and losers. This raises difficult ethical questions in deciding when a benefit for society as a whole justifies the imposition of a risk that falls especially heavily on a small target population, with realized serious adverse consequences for even fewer. It also has the consequence that risk decisions are made in a charged political environment in which the prospective winners and losers use all the means at their disposal to try to influence the decision in their favour.

There are abundant examples of public decisions on risk with distributive consequences. A recent decision regarding compulsory reporting of pilots' medical conditions reduces the right to privacy of one group in order to achieve a reduction in risk for society as a whole. Recent controversy over compulsory reporting or even quarantine for victims of AIDS concerns the same issue. Decisions on the transport and disposal of hazardous materials or the siting of large industrial facilities such as chemical plants, refineries, LNG facilities or power plants all involve the imposition of elevated risks on the people who live nearby for the economic benefit of society as a whole.

The intrinsic difficulties of decision-making under risk cited above make the political decision process for such issues particularly vulnerable to tactical maneuvering. The shortage of compelling scientific evidence on most contentious risk issues and the accompanying widespread expert disagreement make it possible for almost any position to find credible expert support. And the widespread adherence to the homily that "we the government do not take risks with the public" makes any proposed decision into an open target for its opponents; any decision that results in public risk, as all do, can be portrayed as government's failing in its responsibility to protect the public.

All of these problems make the resolution of public policy issues of risk extremely difficult. But they are also among the most important questions we face, as is indicated in part by the controversy they generate. Many of the most hotly contested policy disputes of recent years have had major risk components--for example, nuclear power, pesticides, and the transport and disposal of hazardous chemicals. Moreover, both government policy and public opinion have been profoundly inconsistent in their treatment of various risks.

What Can Be Done?

This list of problems besetting the making of public decisions of matters of risk is certainly formidable, but it does not justify throwing up one's hands in despair. Although we really cannot know what risks we face, we still have to make decisions, both personal and professional. And there is some guidance that can be offered to help a public servant make the best decisions on matters of risk that are possible under the circumstances.

First, there are a couple of positions that one would do well to avoid. The contentious and uncertain environment that characterizes issues of risk can give public officials a strong incentive to protect themselves from censure. There are two patches of political high ground available, either of which might seem attractive under some circumstances. The first is occupied by exaggerating the clarity and certainty of the scientific facts relevant to the decision and pretending that these facts make the best course of action so overwhelmingly apparent that there is only one possible decision. Uncertainty is suppressed, values rolled up with facts, and personal responsibility is successfully avoided with the words, "The numbers made me do it". This stance is more readily available the further one is from the political level.

The second patch of high ground lies in the opposite corner. It is occupied by exaggerating the lack of factual knowledge available, asserting that questions of public risk are entirely issues of values, and referring every difficult risk decision to polls or referenda for resolution. Uncertainty is exaggerated, facts are rolled up with values, and personal responsibility is successfully avoided with the words, "The people have spoken". This stance is more readily available the closer one is to the political level.

These two stances are of course caricatures, and neither represents responsible behaviour for a public servant. But between these patches of high ground the water is very muddy. Public servants must somehow balance their greater access to relevant knowledge and time devoted to thinking about issues of public risk against the preferences of the public as they perceive them.

The first issue that has to be addressed is the disparity between people's observed choices on risk and rational choices. It seems clear that responsible conduct by a public servant must favour the normative approach rather than the descriptive one. People may misjudge their interests, but if you are charged with serving their interests and acting in the public interest, you surely have a responsibility not to base your decisions on such misperceptions. The obvious difficulty with this approach is that it moves dangerously close to paternalism, or disenfranchisement of the populace. If we presume that the public servant knows what the people need better than they themselves do, then how is an informed and responsible public to make its views prevail when a public servant acts arbitrarily?

Amos Tversky, one of the architects of Prospect Theory mentioned above, insists strongly that this theory should not be taken as a guide for decisions on behalf of others, but that in such contexts only the normative economic approach makes any sense. He compares his theory to the theories of perception that describe how people are susceptible to certain kinds of visual optical illusions, and argues that we should no more make prospect theory the basis of public policy on risk than we should make the theory of visual illusions the basis of our system of measurement of lengths.

A way of balancing the proxy and paternalistic approaches to public decisions was suggested by a participant in a Royal Society conference on risk a few years ago. "If the public, knowing that botulism causes only 2 deaths per year, wishes to devote disproportionate resources to protecting against it, then I accept the democratic decision. But if the public opinion is based

on a misperception of 500 deaths per year, then I believe I should appeal the decision." And indeed contest it.

Suppose that you have decided to try to be rational in a decision on a risky matter and not be swayed by personal or public misperceptions. What do you do? Although much of the information you need may not be available, you can get some help from taking a systematic approach to the problem and describing the possible consequences of your decisions in terms of numerical probabilities. There may be evidence available in the form of relevant experimental or theoretical scientific results, that can help you estimate both the likelihood of certain things happening, such as certain kinds of accidents, and the likely magnitude of their consequences.

There are techniques, gathered under the rubrics of risk-benefit analysis or quantitative risk assessment, that can help combine relevant information from a wide variety of sources to estimate the total risk resulting from any choice you may make, and to indicate how much uncertainty there is in the estimate. These techniques have vigorous defenders who proclaim that they are the ultimate answer to difficult decisions involving risk, and vigorous opponents who claim that they are of no value whatever. In fact, their mathematical sophistication cannot finally make up for the profound uncertainty that may prevail in the fundamental data that go into them, and so the estimates of risk that they generate are often rather crude. But they nevertheless have strong advantages over the alternative methods of making decisions in the face of risk, most of which fall under the categories of intuitive expert judgement, precedent and current practice, or uninformed guessing.

To estimate the consequences of a particular event or exposure level, these techniques can combine data from epidemiological studies, animal experiments, and known biological or physical mechanisms of injury. These data are combined with estimates of the number and sensitivity of people exposed, and physical models of the dispersion and dilution of hazards to give estimates of the total consequences in terms of numbers of injuries or lives lost. To estimate the probability of accidents of specified severity occurring, a technique known as Fault Tree Analysis is used, which allows reliability data from many industries to be employed. Probabilities are estimated by considering the chains of events that can lead to an accident. One asks first, in order to have an accident what must have happened? Then, in order for each of those events to happen, what must have happened? By reasoning backwards from a hypothesized accident to the simpler events and logical connections between them that can have caused it, the

problem of estimating the total probability of an accident can be reduced to the combination of a large number of simpler probabilities for such events as failed switches, stuck valves and broken pipes, which are well known from a long history of operating experience.

These techniques can expose and clarify the assumptions about unknown probabilities that go into your estimates of risk. They allow you to combine whatever information is available on a given feature to make a best guess, and then let you test how sensitive the total calculated risk is to your guess. This is a strong advantage over other approaches, which also require guesses but usually bury them in an informal argument so that they can neither be seen nor tested. By making your thinking about an issue of risk more systematic, these techniques can also call to your attention significant sources of hazard that you would have otherwise overlooked.

A formal risk analysis does not tell you what the best choice is, but merely clarifies the consequences of each choice to allow a well-informed political decision. The value issues are not resolved by the analysis. Explicit decisions may be required on such matters as how much you are willing to pay to save a life, or how much disparity in risk between two groups of people is acceptable, or how much you are willing to pay to reduce such a disparity. But these decisions are explicit and subject to political debate, and consequently a formal analysis separates as much as possible this process of value-based decision from the process of describing the consequences of various possible decisions.

The most significant drawback to formal risk analysis lies in its capacity for misuse. Its numerical results lend it a sense of authority and exactitude, even though they are always based on a large number of technical, economic, and social assumptions; it can thus be used to give a spurious image of certainty to a truly uncertain issue, to present value judgements in the guise of facts, and to bolster the "numbers made me do it" stance caricatured above. Used correctly, it represents the best chance to separate issues of fact from issues of values in the decision-making process. But the extent to which this goal is achieved depends strongly on both the quality of the analysis and the quality of its communication.

By providing estimates of how much various groups of people are put at risk by a decision, a formal risk analysis can also help in determining appropriate government measures to mitigate, compensate, or share risks. When the public benefit requires the imposition of a risk that falls especially heavily on a few, then the resultant inequity can to some extent be reduced

by providing monetary compensation or insurance for those who bear the increased risk. People who live near a proposed industrial facility, say a LNG terminal, that exposes them to some increase in risk may be more willing to accept it if the industry or the government offers a lump-sum payment to compensate for the decrease in the value of their property, or insures them for losses resulting from the plant.

Compensation for the loss of property value that results from a risky or unpleasant facility assumes that price changes in real estate markets accurately reflect people's attitudes to the imposed risk. A similar example of ex ante compensation is the payment of higher wages for riskier employment. The assumption is that workers may knowingly accept higher risks in return for a higher wage, but a number of factors call this assumption into question. First, although in some industries such as logging the riskiest jobs are the highest paid, this might not be risk compensation because these jobs are also the most skilled; and in many industries jobs with higher risk are paid less than jobs with lower risk. Second, the information available about individuals' tendencies to misjudge risks suggests that workers may be no better able than anyone else to make a rational informed decision of how to trade off risks against income.

The two difficulties associated with ex ante compensation are that the appropriate compensation levels are hard to determine--markets may not do it, and with non-market schemes the affected individuals have an incentive to over- state their claims--and that blanket compensation schemes cannot adequately reflect individual differences in required compensation. Some members or officials of a community may in fact prefer to see a risky facility built for economic or employment reasons while others perceive the facility as all risks and no benefits. But attempting to find the fair level of compensation is certainly preferable to letting an imposed risk go uncompensated.

Insurance for losses suffered from the imposed risks can take the form of compulsory liability insurance held by the person or company carrying out the risky activity, or insurance by the government. Historically, it has often happened that insurance was not agreed on in advance, but the government has stepped in with compensation after a disaster. This has been the case both with technological hazards and natural ones.

The main problem with insurance for imposed risks is determining the appropriate level of coverage, particularly since, if losses occur, very high liability payments may be determined by jury verdicts. Whatever level is chosen, though, equitable distribution of risks would be

assisted if the nature and size of compensation or insurance schemes were publicly established at the time that the risk is first imposed.

It has also been suggested that the expectation that disaster relief will be provided, combined with people's misestimation of risks, leads too many people to choose to put themselves in risky situations, e.g., moving in to live behind dikes in flood plains.

It may be that the best government response to a matter of public risk is regulation or information. If the consumption of a product carries risks not easily assessed by the consumer, perhaps a government programme informing them of the magnitude of risk will allow them to make better decisions. But if obtaining the necessary information is too difficult for consumers, regulation may be required. Such regulation could take the form of required standards of product safety, licensing of certain professionals such as physicians, or the restriction or banning of the sale of a product.

Regulation can also encourage equity of risk distribution by exploiting different sensitivities to risk. Uranium miners suffer a high risk of lung cancer, but the cancers are very slow in developing (with a latency period of 20 to 30 years), and the risks are vastly greater for smokers. These risks might thus be borne most fairly if only nonsmokers over 50 years of age were allowed to work in uranium mines, but would such discrimination be acceptable?

When the government regulates the availability of goods or services on the basis of risk, it enters an area of strong conflict between different groups in society. People differ greatly in the extent to which they wish to take responsibility for the risk consequences of their decisions. Consequently in many areas there is vigorous conflict between those who wish to have a product or service available because in their view its benefits outweigh its risks, and those who wish the government to protect them by ensuring that no product with significant risk is available in the open market.

Perhaps the clearest examples of these conflicts are in health care, and drugs. Consider the conflict over the licensing of non-medical healers such as acupuncturists or homeopaths. Some people argue that prohibiting such services deprives them of access to their treatment of choice, while others claim that licensing such services would expose the uninformed public to dangerous quacks. There is currently a similar controversy over the proposed licensing of the new oral contraceptive, Depo-Provera.

Since virtually no commodity can be consumed without some risk, a legal system that allows people to sue for harm caused by products raises similar difficulties and illustrates the same conflict between those who wish to choose for themselves and those who wish to be taken care of. The recent withdrawal from sale of the Copper 7 IUD in the United States illustrates this conflict. Although the Copper 7 was widely regarded as the safest of the IUDs, its manufacturer faced large numbers of personal injury lawsuits-perhaps because of the publicity generated by the claims against the badly designed and excessively dangerous Dalkon Shield. But the withdrawal of the Copper 7 means that virtually no IUDs are available to American women.

The American legal system does generate some remarkable ironies on issues of liability for public risk. Contrast the following recent judgements: a manufacturer of football helmets was ordered to pay \$12 million to a high school player paralyzed in a game because the helmet did not carry a label warning that it was dangerous to ram opponents with it; while a Philadelphia Court of Appeal ruled that the Surgeon-General's warning on cigarette packages absolved cigarette manufacturers of responsibility for damages due to smoking-induced cancers, despite the enormous advertising campaign mounted by the manufacturers to induce people to smoke.

While the Canadian public and courts have been much more restrained, there is growing evidence that Canadians are being influenced by the torrent of large product liability and malpractice suits they read about in the US. The recent and well-known Brampton case is illustrative. While the wish somehow to compensate someone who has suffered grievous injury is natural, there are two associated questions that are exceedingly difficult, both in the legal setting and in the making of public policy: what measure of responsibility falls on people for the consequences of their own actions? And if no product or activity can be perfectly safe, how do we react when an apparently rare unfortunate event does occur? By saying that it is unfortunate that we have observed such an uncommon accident? Or by saying that the occurrence of an accident is sufficient evidence that someone has not been careful enough? Is there no line that can be drawn between negligence, incompetence or crime on the one hand, and simple bad luck on the other?

On the part of the public, what is required is acknowledgement that perfect safety is not attainable, and a lot of the responsibility for each person's safety must necessarily fall on his or her own shoulders. Clearly there is a right to be informed of the risks of what we do or consume, to the extent they are known at all, and to expect honesty and competence from those we deal

with in the market and the government. But risks and bad luck do exist, and we do not have the right to demand that we always be lucky.

If such an attitude of responsibility were more widespread, it would likely be less dangerous for a public servant to admit to taking risks with the public. Decisions regarding risk would not have to be made in such an environment of secrecy, and perhaps there would be increased opportunity for responsible informed public input into collective decisions on matters of risk.

Conclusion

For the public servant, then, what advice can we offer? Clearly a major part of the message is that you do take risks with the public and you cannot avoid doing so. Nothing you can do will permit zero risk, or even the finding of an incontestable best choice on any matter of risk.

But a formal systematic analysis of the magnitude of risks, using estimated numerical probabilities and consequences, can help you clarify just what risks you are working with, and help you make the best decision under the circumstances of incomplete information, distributional inequities, conflicting agendas, and time pressure that you have to endure.

Programmes of compensation and insurance, of regulation and information, can help mitigate unequal risks that must be imposed, and can help people take responsibility for their own choices of risks.

Finally, an understanding of the systematic departures from rationality that affect people's judgements of risks can help you anticipate what kinds of public reactions your decisions will create, and can remind you of the paramount importance of presenting issues of risk to the public as transparently as possible, so that their judgements of risks will be less subject to misperceptions and bias.

The responsibility of the public servant could be summed up as reconciling what the analyst sees as the public interest with what the politician feels as public pressure. This requires framing decision problems as clearly as possible, and taking seriously the educational role of the public servant. It will often require programs of information and market regulation to support a market environment in which people can express well-informed preferences and find their most rational choices. But in the final analysis, your responsibility as a public servant is to base your

decisions and advice on the normative, scientific analysis of the public interest, not on the public opinion polls.