The Arithmetic of Risk

Analytical Problems and Political Solutions

By Rod Dobell

In the Spring of 1978, the Atomic Energy Control Board published a document entitled 'Risk of Energy Production' by Dr. Herbert Inhaber. Almost immediately they found themselves embroiled in heated, if not always enlightened, controversy. By the end of the year, the document had been through two further revisions, and a third was contemplated. This contemplated fourth edition (third revision) was never published (although it is apparently available for examination in the AECB library), and the second revision is now out of print.

What got the AECB into so much hot water was an attempt to compare the estimated risks to the health and safety of workers and the public involved in the production of energy by various systems. The document emphasized the need to estimate risk over the whole fuel and facilities cycle, from mining of the fuel and manufacture of the necessary equipment and capital facilities, through to management of the resulting wastes.

Although the author takes pains to emphasize the exploratory nature of the work and the degree of uncertainty and lack of comparability in the data used, he does come to the conclusion that nuclear energy has a fairly low overall risk compared to other sources, including non-conventional systems such as solar, wind, or ocean thermal. Energy systems based on production of electricity from coal and oil were judged to offer highest risk, primarily because of the pollution problems associated with their use.

The controversial results with respect to 'soft' non-conventional sources spring from the manufacturing activity required to put the facilities in place, and, more importantly, from taking into account the need for “back-up” energy systems based on conventional sources.

This is not the place to go into the merits of the argument. What is discouraging, however, is the violent public reaction and the fact that the resulting controversy generated so much (waste) heat, and so little light. We have to find some ways to achieve a better social debate on problems involving major risks.

The last few years have seen frequent declarations of frustration with the processes by which Canadian society makes decisions involving risks and uncertainty. Prominent executives have been quoted as suggesting that we can have energy self-sufficiency, or we can have stringent safety and environmental standards established in exhaustive public hearings, but we cannot have both. Seven heads of state meeting at the Tokyo Summit of June 1979 agreed that among the measures needed to improve the long term efficiency and flexibility of western economies were regulatory policies which avoid unnecessary impediment to investment and

increased productivity. Supporting texts made it clear that much of the "new" regulation-including health, safety, and environmental standards-was implicated.

Such standards, it has been argued with vehemence, impinge adversely on prospects for economic growth and adjustment, and reduce national and personal incomes to an unacceptable degree below what they would otherwise be. Carrying the argument a step further, Aaron Wildavsky of the University of California at Berkeley has argued in a recent paper ("Richer is Safer") that the consequences are in fact perverse, that the search for greater safety through more stringent standards reduces income prospects so much that overall health and safety suffer in the end.

Wildavsky presses on to argue for less government, less regulation, more individual choice in matters involving risk, as in other economic decisions. This pluralist solution, with its roots in the emerging "public choice" literature, has echoes in quite a different stream of thought in the literature on a large-scale systems.

In the 1979 Distinguished Lecture of the International Institute for Applied Systems Analysis (IIASA), George Dantzig of Stanford University summarized the case for less "taut" technological systems. He laid out the case for a richer system with greater slack and looser tolerances, designed with substantial built-in redundancy in order to achieve resilience or robustness in the face of the shocks which are unavoidable in an uncertain world. This pressure for decentralization, for multiple experiments, for management based on trial-and-error rather than massive central planning, was seen as the higher rationality in the face of risk. What appears as wasteful redundancy to the untutored eye emerges as reasonable diversification and hedging when one looks at various contingencies. Even the most cost-conscious of business executives has seen the merits of duplicate baskets when it comes to diversifying the eggs.

(Whether all auditors-general have yet adapted their views on accountability and "value for money" to reflect this reality may be debatable, however.)

It is not only at this macroeconomic level that frustrations are expressed. The degree of apparent inconsistency amongst various public investment decisions involving safety, the alleged inability to establish reasonable safety standards in individual industries, the extent of unresolved controversy about uranium mining or the siting of nuclear plants, all seem to have given rise to growing disillusion with established processes for public decision when facing risky circumstances in an uncertain world.

One consequence has been new regulations calling for evaluation of the socio-economic impacts of proposed new regulations. At the federal level, the Treasury Board Secretariat has incorporated into its Administrative Policy Manual guidelines on the analyses and consultative processes required prior to the introduction of new regulations dealing with health, safety, and fairness.

A declared purpose of this Socio-Economic Impact Analysis policy is to encourage greater public participation in the regulation-making process, in part by issuing background studies. The third in this series of studies on government regulatory activities has just appeared, this one dealing with occupational health regulations limiting exposure to radiation in uranium

mines.

Another result has been renewed interest in the study of the distortions which occur in individual perceptions of low probability events carrying risks of highly unfavorable consequences. One such study, a joint project of the International Atomic Energy Agency and IIASA, was directed explicitly at the consequences for societal decisions related to large-scale technology. Other work on analytical methods for possible improvement in public decisions involving risk was discussed at the 1977 Denver meetings of the American Academy for the Advancement of Science, and has recently been reported in a publication in its Symposium Series under the title, "Judgment and Decision in Public Policy Formation". Both of these efforts will be discussed further below.

The purpose of this article is to reflect briefly on the issues raised in all this literature, and the implications for our present institutions. (For those interested, footnotes and detailed references-which would be out of place in this context-have been collected in a separate note available from the writer.)

More importantly, however, the purpose is to call attention to the need for discussion of this topic, with this magazine being one forum for the purpose. Canada has supported IIASA as one of the founding members; utilization of the work done by that body in the area of risk assessment is one of the ways in which that investment might have direct national payoff. For this purpose, a major conference following up the Denver meeting and attempting a synthesis of the lessons from the IIASA work and related European studies is planned for next year.

Indeed, amongst the major programs of the Institute from which Canada has derived significant benefit is the work on ecological systems initiated by Buzz Holling and his co-workers at IIASA, and now carried on at U.B.C. Some of their ideas on "resilience" and "hazard management" deserve to be debated in these pages as crucial questions of social policy.

Thus, this present article is really a call for a debate on appropriate compromises in the ancient tension between analysis and politics in government-between "scientific" thinking and "strategic" thinking in handling public decisions. It seems clear that in matters of risk assessment and standard-setting, we are increasingly heading down the road of partisan analysis, of advisers as advocates, of the competition of ideas in formal or informal "science courts".

Must we do so? Is it necessary that we live with the fact that people do not seem to deal with risky situations as "logic" dictates they should, or are the participants in social decisions "educable" on these matters? If individual perceptions of risk deviate from consistency, must collective decisions be equally idiosyncratic? What role can analysis play in such decisions?

As an aside, it is interesting that there seem to be some particular difficulties in the handling of uncertainty. Conventional wisdom in the operations research literature suggests that experience teaches pretty good solutions to optimization problems: trial and error, and rules of thumb, seem to lead close to optimal solutions to many of the very complex optimization problems that arise in industry. And, despite some recent disenchantment, it is fascinating that industrial societies have evolved market mechanisms to provide precisely the information and the incentives leading to such solutions of both static and dynamic optimization problems.

But in the presence of uncertainty, these results seem to collapse: people do not seem to be intuitively very good decision analysts, and, with few exceptions, appropriate market mechanisms for pooling risks have not evolved. How effective "good" analysts might be in arriving at “good” decisions is a question explored a little further below.

“Responsible Regulation”, the interim report of the Economic Council of Canada's regulation reference, makes it very clear that the "new” regulation will prove, in the end, an unlikely candidate for deregulation. Inextricably associated with the setting of standards-for health, or safety, or fairness-such regulation is opposed by many businessmen as both costly and arbitrary. Yet strong pressures exist to retain such regulatory authorities, even though little is known about the consequences, the alternatives, or procedures to establish appropriate levels for standards in any of these areas.

Social Decisions in the Face of Risk are Inconsistent

Present procedures do not appear to achieve anything close to a rational allocation of society's resources amongst lifesaving and other competing social objectives, or even among different ways to achieve the goal of reducing risks of life. Examples abound: investments in highway engineering appear to promise a reduction in highway deaths at a cost of only $20,000 to $80.000 per life saved; decisions on standards for air travel imply a value of at least ten times that amount. It has been suggested that some regulatory decisions in the nuclear power industry imply a cost per life saved of around six billion dollars.

One cannot quarrel directly, perhaps, with any of these numbers, although the Economist of March 22, 1980, took to task the estimates of the U.K.'s National Radiological Protection Board which had just published a cost-benefit analysis of nuclear safety involving estimates of up to 40 million pounds as the sum worth spending to save a life. But the differences mean that a transfer of existing public resources from investment in nuclear safety to air safety or from air safety to highway safety-or from highway engineering to driver education-would permit an increase in lives saved with no increase in outlays.

Similarly, on the physical evidence, it has been conjectured, partly in jest, perhaps, that a prohibition of nuclear energy compensated by an increase in solar heating might lead to greater public exposure to radon daughters in homes sealed by insulation to accommodate solar heating than any exposure to radiation from the corresponding nuclear power plants. Or, as has been often observed, a substitution of coal for nuclear energy as a power source may lead to greater long-term health hazards due to air pollution than those due to exposure to radiation.

Transport of gasoline by truck appears from all the statistical evidence to be an overwhelmingly greater threat to life than transport of radio-active materials, but public opposition to the latter is overwhelmingly greater. And public regulatory control of the latter is correspondingly more stringent.

Of course, it is necessary to recognize that one kind of life may not be a substitute for another. One may accept risk as the price paid in pursuit of other benefits, or one may pursue risk as a matter of life-style. The old joke about the patient who rejected medical advice that life could be prolonged by abstention from cigarettes, wine, fine food, and the pursuit of the opposite

sex has a lesson for analysts.

Simple calculation shows that if an hour of jogging per day will extend the life of an eighteen-year old by only two years (from 70 to 72, say) the time for living (net of jogging) is still reduced. Before jogging fans rise in anger, it should be added that if jogging is also counted as living, then no problems arise. But the serious point is that evidently all lives are not the same, and perhaps there is no reason to expect that the cost to save one kind of life will be the same as that for all others, or in all circumstances.

A different kind of inconsistency seems to arise in "rescue" situations, where the willingness of the public, or at least public officials, to support great expenditures to attempt the rescue of people in danger far outruns their willingness to invest in preventing the danger arising in the first place. What is even more puzzling is the apparent willingness to accept, or to assign others to accept, risks to the rescue teams which may in some sense outweigh the risks to those to be rescued. That is, rescue attempts may be mounted even when the expected number of lives lost is larger with the rescue attempt than without it. Interestingly, it is not clear whether, or how far, the social response differs in circumstances where the initial risks were willingly assumed (as in a mountain-climbing accident) or an occupational hazard (as in a mine explosion).

Again, however, it has to be noted that there are other benefits to be considered. For the rescue team, or the community, the benefits to be derived from successfully, in the face of great odds, pulling off a difficult rescue of another human being in trouble, may go far beyond the value of a life saved.

So social decisions are inconsistent. But these inconsistencies may reflect intangible considerations like the renewed feelings of solidarity and community which attend a successful triumph over threatened danger, or the feeling that life was not intended to be risk-free and that some kinds of life are not worth living, or the recognition that people do not handle probabilities as the textbooks say they should.

Individual Perceptions of Risk are More Personal

In some cases, it may be that the inconsistencies noted above have explanations in personal perceptions of risky situations. It is well known that expected values are not adequate measures of the consequences of chance events. A certainty of one fatality in a given year is sensibly a different prospect from a one-in-a-million chance of a million fatalities. Evidence from laboratory studies which Professor Alex Bavelas has undertaken over many years suggests also that people are not indifferent to scale or to the mode of presentation of data. An individual might rather be part of a group of fifty, of whom one-half were to be fired, than of a group of two, of whom one was to be fired, even though statistically the odds are the same. If nothing else, one may find comfort in numbers, but also perhaps a sense of security in the possibility of hiding in the crowd.

An important source of distortion in problems involving uncertainty is analogous to the well-known problem of sunk costs in project appraisal. Just as the decision to complete a partially-finished project should depend upon quite different cost estimates than the original decision to proceed with initial construction, so the assessment of the probabilities of particular

events must depend upon how much of the veil of uncertainty has been lifted thus far, and what has been revealed.

The problem boils down to the appropriate point of departure or origin of events. Pierre Simon, Marquis de Laplace, put it well almost two hundred years ago:

> *When a number in the lottery of France has not been drawn for a long time, the crowd is eager to cover it with stakes. They judge since the number has not been drawn for a long time that it ought at the next drawing to be drawn in preference to others. So common an error appears to me to rest upon an illusion by which one is carried back involuntarily to the origin of events. It is, for example, very improbable that at the play of heads and tails, one will throw heads ten times in successions. This improbability, which strikes us indeed when it has happened nine times, leads us to believe that at the tenth throw tails will be thrown. But the past indicating in the coin a greater propensity for heads than for tails renders the first of the events more probable than the second….*

Inconsistencies in social decisions arising from illusions such as these presumably can be corrected. Experience suggests that concepts such as the irrelevance of sunk costs in deciding on current actions, or the importance of appropriate discounting in appraising future benefits, can be learned. Both public and corporate decisions appear to reflect greater appreciation of these economic ideas which were deemed quite artificial- not so many years ago. Hopes for comparable improvement in the ability to avoid fallacies such as those just discussed, when taking social decisions under uncertainty, might not be misplaced.

Thus, there are two possible aspects of individual decisions or perceptions which must be studied-preferences over alternative outcomes, and perceptions of the risks themselves, which is to say, the probabilities of the different outcomes. As to the first, there exists extensive literature on methods by which individual preferences or utilities may be explored. It appears that people weight less heavily risks which are familiar and risks which are voluntarily assumed. The willingness to accept higher levels of risk in response to the promise of higher benefits seems indisputable. But in addition, there is some evidence that the degree to which risks are unknown or uncontrollable, or the extent of delay before the outcomes are known, may influence the willingness to accept risk in return for promised benefits.

The trade-off between different aspects or attributes of the outcomes in risky situations is a particularly difficult question which has led to much recent work on the subject of "multi-attribute utility functions" work in which Howard Raiffa of Harvard (and the first Director of IIASA) has again been a leader. The AAAS symposium mentioned above contains a paper describing application of this work to the problem of selecting sites for nuclear waste disposal. In particular, the argument is made that these techniques may be of use in assessing and reconciling different and conflicting values with respect to various attributes of alternative solutions.

The second problem is perhaps less familiar, but a growing body of evidence

suggests that otherwise intelligent and "rational" individuals may not have valid perceptions of the frequency of the hazardous events to which they might be exposed. Their assignment of probabilities to various possible outcomes may be subject to systematic distortion or illusion.

In particular, there appears to be a tendency to overestimate the probability of highly memorable or easily imagined events. Professor A. Tversky of the Hebrew University of Jerusalem, who for the past decade has been conducting a research program into the principles that govern human judgment under conditions of uncertainty, refers to this as the principle of availability. "In this heuristic, one judges the probability of an event, or the frequency of a class, by the ease with which an event can be imagined, or by the ease with which instances of the class can be brought to mind." Sensational media coverage of some classes of events obviously may contribute to this weighting of risks.

There appears to be an overly strong tendency for beliefs to be held in the face of the evidence. The failure to incorporate new information into one's beliefs in a sufficiently responsive way leads to perseverance of initial opinions in the face of conflicting evidence longer than any "rational" information-processing model would predict. On the other hand, Tversky refers to the contrary problem in dealing with general statistical (prior) information as compared to specific "individuating" information. "Data show that, when statistical information alone is available, people use it sensibly; once individuating information is added, the statistical information is completely ignored."

It has been argued that there are also serious threshold effects, that the human mind tends to treat all sufficiently small probabilities as the same. The differences between a one-in-a-million chance and a four-in-a-million chance may be hard to identify mentally.

Thus, some of the inconsistencies observed in social decisions as to investments in safety may reflect individual attitudes toward risk. Automobile safety is valued less because the degree of personal control seems higher, because media coverage of air disasters makes the images more vivid, even though the phenomenon is unfamiliar, and perhaps because the risks and thrills of driving have a greater element of voluntarily-assumed danger about them. But problems remain; it seems debatable whether it is acceptable that social decisions should reflect these sorts of personal attitudes and distorted perceptions of risk, when in fact a re-allocation of available resources could save lives.

Other apparent inconsistencies may perhaps be resolved by recognizing the significant differences in individual reactions to the statistical likelihood of a reduction in an expected number of (anonymous) deaths and the fact of hazard to a known, identifiable, individual person. Before the fact, a reduction in a "statistical" hazard obviously will weigh less heavily than the concerns, after the fact, with the risks to an identifiable person. "Statistical" lives saved are simply worth less than "identifiable" lives saved.

Partly this situation is understandable simply as a kind of social contract: a community which does not find it worthwhile to invest in preventive measures to reduce exposure to risk may nevertheless stand by an implicit social commitment to massive expenditures if necessary to come to the rescue of those in imminent danger. Again, however, one

may be troubled by the degree of responsibility for the weekend climber trapped on a mountain he has been warned not to attempt, as compared to that for the miner forced by economic circumstance to work miles below the surface.

Perhaps the voluntary assumption of risk for the sake of challenge or thrill should include the voluntary relinquishing of any claims for extraordinary social action in case the risks are realized?

The vexed question of the degree of social responsibility for those who bring some of the hazards upon themselves arises in other ways. Should the same social resources go into treatment of lung cancer in those who have steadfastly refused to help themselves by reducing their consumption of cigarettes as to those who have never exposed themselves in that way?

In the Treasury Board Secretariat study of exposure standards in uranium mining. the use of protective helmets as a means to reduce levels of exposure for individual uranium miners was considered. But it was argued that union officials do not like solutions to health and safety problems which require active participation and cooperation by the workers, preferring instead more general attacks on sources of damage. Under some circumstances, however, one might expect reliance on some degree of individual responsibility to be an appropriate feature of solutions to problems involving risks.

Risk Assessment

It has become common to separate the question of social choices relating to risk into two elements, as above-the problem of estimating the relevant probabilities, and the problem of appraising the acceptability of increased risk undertaken for the sake of some promised benefit.

The first problem may appear a straightforward technical matter, but we have already observed that people have real difficulties in assessing properly the risks to which they are exposed. Perceptions of risk are distorted, though this does not mean that they cannot be improved with education. Moreover, the data on which to base probability estimates are almost always absent. In any significant social decision relating to risk, judgments will have to be substituted for data in the assignment of some critical probabilities. Another way of saying essentially the same thing is that it is hard to know what is meant by the frequency of hazards associated with one-time events.

But the even more ticklish subject will be the appraisal of the outcomes, or the utilities associated with the pay-offs, or the preferences as to all the possible consequences. In estimating the social costs of projects that might increase the risk of death, or the worth of activities that might prolong life, two general lines of thought have developed.

The first, or "human capital" method, sought to value loss of life in terms of the loss in national product resulting from the increase in mortality. This notion of compensation to those affected shows up rather naturally in damage suits and court judgements, but it has been criticized, accurately, as failing to recognize the rather common personal desire not to die, even if all those around could be fully insured or compensated for the loss.

A second method has grown up to circumvent this problem, based on the idea of "willingness to pay" to avoid loss of life. A variety of approaches to assessing personal willingness to pay have been devised and discussed.

But to the extent that both of the above methods seem to be attempting to attach an explicit value to an identifiable human life, they are hotly contested as attempting to measure the unmeasurable. The fact that implicitly we do it all the time, and take for granted the necessity to do so, does not reduce the fervor of this criticism.

An alternative approach seeks to avoid asking about the "value of a life" by attempting to determine the value of a reduction in risk. Rather than evaluating the loss of a life, therefore, this method considers the price that should be paid for a reduction in the mortality curve, or an increase in the probability of survival.

The human capital method has the advantage of relying on conventional sources of data and familiar accounting or actuarial principles. Unfortunately, it is often seen as missing some essential features in the social assessment of risk. Despite this criticism, the figures so derived may be useful as lower bounds or minimum estimates for purposes of cost-benefit evaluations of proposed investments in safety.

Willingness-to-pay measures, by contrast, must rely on much less conventional data. We have, of course, some implicit evaluations or social preferences revealed by past investment decisions. We have spent, or elected not to spend, money on vast arrays of social projects having recognized impacts on lives lost or saved. Unfortunately here, too, we have noted that the record appears to show significant inconsistencies, whether one looks at implicit valuations of lives lost, or implicit valuations of reductions in risk. The Marquis de Laplace again:

> *The mind has its illusions as the sense of sight; and in the same manner that the sense of feeling corrects the latter, reflection and calculation correct the former. Probability based upon a daily experience, or exaggerated by fear and by hope, strikes us more than a superior probability (which) is only a simple result of calculus. Thus we do not fear in return for small advantages to expose our life to dangers much less improbable than the drawing of a quint in the lottery of France; and yet no one would wish to procure for himself the same advantages with the certainty of losing his life if this quint should be drawn.*

Risk Spreading

Thus, one has the problem of estimating the risks, and the separate problem of assessing the *estimating* of the risks. In the latter case, the question of risk-spreading also arises as a crucial consideration.

Consider the example of the military draft: how to determine who shall go when not all must go. It seems, before the fact, eminently fair to propose a lottery to select, say, one in six candidates for the draft in a time of military conflict. But, after the fact, the burden rests upon the one selected, while five others are unaffected. The problem is puzzling when the odds are

uniform for all concerned; it is, of course, a still more difficult social problem if the odds appear to favour the poor and native over the rich and white. Even with fully equitable selection, the problem of compensating the loser, or spreading the burden after chance has made its selection, must evidently be faced.

Richard Zeckhauser of Harvard has explored many aspects of this problem of finding social mechanisms for spreading some of the burdens imposed by nature's lotteries. In the case of the draft, for instance, arguments for a more highly-paid volunteer army arise. But in the case of the mentally retarded, or the accident victim, what compensation is possible?

While all may be exposed to the statistical risk of accident, one is selected to bear the burden. There is dramatic asymmetry in the *ex post* positions. The normal "insurance" schemes or markets for contingent claims which economists have invented to handle problems of uncertainty do not work well in these circumstances, where transactions are not voluntarily entered into, and compensation in money terms seems hardly adequate.

How much more difficult the problem, then, if the role of chance itself depends upon the decisions of a site selection panel? Those near the site selected for a nuclear plant may believe that they are forced to bear the risks of serious personal hazard, while their countrymen do not. How to spread the risks in a socially acceptable way? How to insure or compensate those under the shadow? In line with the objective probabilities? Or in line with their (genuine) distorted personal perceptions of the risks?

One possible conclusion in the face of all these factors is that there is no option but to leave decisions as to risk and safety standards to the political process, which is to say to the politicians and lawyers rather than the analysts. This, in fact, seemed to be the conclusion put forward by Michael Trebilcock of the University of Toronto Law School in a University of Victoria symposium at Pearson College last spring on the subject of deregulation.

Recalling the example of a recent California law suit involving the Ford Pinto. Trebilcock observed that one could sympathize with a thirteen-year-old who felt that 125 million dollars is not sufficient compensation for a life without use of his limbs. Indeed in this case, it was reported that the amount of damages was set by the jury's feeling that Ford needed to be taught not to rely on benefit/cost analysis in determining whether its cars were safe enough.

Yet, if not by cost/benefit analysis, then how? Though the numbers used to represent the social costs of death and injury may not seem acceptable as compensation after the fact of an accident, what other benchmarks exist to establish what is "safe enough"? If Ford engineers took the damage settlement as the basis for decisions, they would design a vehicle that was too safe-too safe relative to the dangers posed by other equally common implements, inconsistent with other social investment decisions, too safe relative to what buyers would be willing to pay for the standards achieved.

The economist's simple vision is that the incremental gain in safety from an extra dollar expended should be about the same all the way around the vehicle. There is little point in having an ultra-safe gas tank unless the steering mechanism is up to the same standard. And little point in having an ultra-safe Pinto if highway engineering or boating practices or the design of

bathtubs pose greater risks curable at lesser costs.

Conclusion

One must respect the rights of individuals to say things and do things that are not consistent with the textbooks on probability or Bayes' formulas for revising prior beliefs in the light of new evidences. Models of risk acceptance cannot dictate what risks a society should accept, but they should presumably reflect society's considered values and appraisals of the uncertainties involved.

For individuals, one may find an implicit evaluation of risks, or the value of lives saved, through preferences revealed in behaviour. One may also find expressed preferences suggesting that the acceptability of a risk is related to characteristics like its apparent familiarity, controllability, immediacy. Risks may be voluntarily assumed on weekends by individuals who would defend vigorously their right to refuse lesser risks while on the job.

But under what conditions should it be expected that a government concerned with statistical risks can arrive at decisions which coincide with, or at least command the support of, individuals concerned with individual risks to identifiable people?

The problem has been around a while. In the public debate upon the merits of inoculation against smallpox, prior to the discovery by Jenner of vaccine, Danien Bernoulli demonstrated statistically that a policy of inoculation would extend the mean duration of life significantly, even though voluntary inoculation did carry a significant risk of immediate death for some who would not otherwise have contracted the disease, Laplace reports that:

> *D'Alembert attacked the analysis of Bernoulli: at first in regard to (aspects of his statistical analysis), then in regard to its insufficiency in this, that no comparison was made of the immediate danger, although very small, of dying of inoculation, to the very great but very remote danger of succumbing to natural small-pox. This consideration, which disappears when one considers a great number of individuals, is for this reason immaterial for governments and the advantages of inoculation for them still remain; but it is of great weight for the father of a family who must fear in having his children inoculated, to see that one perish whom he holds most dear and to be the cause of it.*

I am indebted to Alex Bavelas for this reference and for putting the analytical question so precisely: why should a father and a government agree? What assumptions on the structure of the problem and on individual perceptions are required in order that one should expect identical behaviour?

What conclusions flow from this brief review?

First, it probably is true that in our behaviour we exhibit some “irrational” approaches to risk. We assume risks voluntarily; we downplay risks where we have an illusion of

control. We prefer to die in familiar, homely ways: we underplay the familiar risks and ascribe excessive importance to the unknown. We prefer to die in small accidents, while overstressing the risks of major catastrophe. We place infinite value on our lives in rhetoric, and abuse our prospects for survival with a daily diet of questionable habits and dangerous activities.

Partly as a result of such perverse behaviour, public decisions involving risk, and processes for setting standards in the health and safety area, are badly handled and need some explicit attention.

Secondly, it is clear that some tools for helping to bridge the gap between analysis and politics in public policy do exist. Their use requires considerably more study of how people handle probabilities in their thinking, and considerably more effort to bring consistency into this thinking. In the meantime, we probably do need to design with more diversity, more provision for experiment, less taut technological structures. More slack means more safety, and diversity does help.