

DRAFT FOR COMMENTS

PUBLIC POLICY IN A RISKY WORLD

Rod Dobell and Ted Parson

DRAFT, JUNE 30 1986

1) Introduction.

One hears a lot about risk in the news these days -- radiation from Chernobyl, sabotage and accident in air travel, unknown health risks of pesticides and dangerous goods, AIDS, the closing of emergency wards, even roller coasters. But public statements on risk have a familiar sound, as if the same questions were being asked and the same arguments made again and again, with little evidence of progress. Perhaps the time is right for a different kind of public debate on risk, one that entails stepping back to take a more general, long-term view and asking about broad social attitudes to risk.

Discussion of risk in public policy proceeds at several levels. The most common and most specific concerns the imposition of particular risks. In debates of this kind, the lines are most often drawn between those who wish others to accept a particular risk, arguing that the social benefits associated with it are sufficient to justify it, and those who would forbid or restrict its imposition, arguing either that it is intrinsically intolerable to impose risks on others, or that the associated benefits are insufficient. Alternatively, when an accident has occurred or a hazard is recognized, an equivalent debate may occur over whether or not the activity or product concerned is safe enough, and whether it should be restricted or disallowed.

Many decisions on risk are of this type. Should Victoria city council permit rail cars containing dangerous goods to travel through the harbour area? What regulation or inspection activities are appropriate for amusement rides in light of the roller coaster accident in the Edmonton West Mall? What pesticides should be licensed for use in public recreation areas, and with what restrictions? Should more nuclear generating stations be built in Ontario? If so, where, and subject to what safety standards?

A second, more general level of risk debate concerns what means society ought to employ to make decisions about particular risky activities or products. In this debate, one finds several schools of thought represented. One group argues that no risk is acceptable, and that society must strive to attain a state of zero imposed risk; consequently, any activity for which a risk can be detected must be curtailed. A little thought reveals, though, that this is an unattainable goal; we are all subject to a great variety of risks already, and every act and decision carries with it some associated risks.

Of those who acknowledge that some non-zero quantum of risk must be socially acceptable, some ask “How safe is safe enough”, seeking a single measure of total risk that can be deemed acceptable. Others argue that a single measure cannot reflect the diversity of people’s activities and preferences and that it is clear from individual behavior that different degrees of risk may be acceptable in different situations. They argue that the acceptability of a risk may be determined by many factors other than its magnitude alone, and seek the conditions that determine the acceptability of a risk in each situation. Both of these groups, however, argue about safety in numerical or probabilistic terms -- the question “how safe?” can in these terms be answered, for example, “on average, two deaths per year”. They both also often advocate analytical techniques to calculate as best we can the total risk imposed by a particular decision, and compare it to the benefits and costs flowing from the same decision.

A fourth group, by contrast, argues that in real decision, issues of risk are not the only or most important dimension, and it is consequently neither necessary nor desirable to base decisions on an explicit calculation of magnitudes of risk. Rather, they argue that public policy decisions should be made through the implicit balancing of interests and rights of various affected parties that has traditionally characterized legislative or judicial decisions.

There is yet another still more general possible level of debate on risk -- one in which the broad issues of social attitude to risk are considered, in which the concern is with questions such as how we can best deal with the inevitably risky nature of the world. Can the costs of risk bearing be shared by society, and if so, how much can they be or ought they to be shared? Is it necessarily preferable to achieve the lowest possible level of risk, or can there be advantages to risk bearing? This level of discussion has been insufficiently recognized in public policy debate, and several factors are making it more and more urgent that it be addressed. Preoccupation with specific questions of risk has led us collectively to behavior that is inconsistent, unfair, and, in the aggregate, reckless.

2) Current Practice.

Individuals are not consistent in their behavior toward risks. This popular observation has frequently been documented in psychological experiments that investigate how people choose among particular risky prospects. It is observed that subjects can be induced to change their preferences simply by changing the language in which an unchanged prospect is presented, that they are usually very bad judges of probabilities, and that they can exhibit inconsistent choices whose net effect is to choose a certain loss. These results have most often been observed in choices among uncertain monetary payoffs, but have also been found in professional judgements by engineers and physicians -- judgements that in a real-world setting could have serious consequences.

These results are perhaps disturbing, but are surely not surprising, when we consider our own behavior and that of people we know. We often engage in risky activities without explicit consideration of the risks we face, or experience fears of activities that we know objectively to be relatively safe. Individuals who protest

vigorously against the imposition of a risk from hazardous wastes or pesticides can be observed to accept quite casually far greater risks by smoking heavily or driving without a safety belt.

A close observation of public policy decisions reveals that they are as inconsistent on matters of risk as those of individuals. Decisions imposing risks or trading off risks for economic benefits are routinely made at all levels of government, but rarely acknowledged. Risk-benefit calculations are frequently performed as a basis for policy-making, but it is hardly rare to hear a politician claim that “we do not take risks with the public”, or denounce a political opponent for apparently doing so. Moreover, large disparities can be found in the magnitudes of acceptable risks that are implied by different public policy decisions.

Inconsistency is equally evident in the retrospective evaluation of decision-makers once the consequences of their decisions become known. It seems that the taking of risks with the public is a punishable offense, if the risks are realized and people suffer harm; if they are not, then the same action may be accepted as prudent management of the public’s funds. The derision that greeted both the federal government’s original support of the Urea Formaldehyde insulation programme and its subsequent decision to provide funds for the removal of the insulation indicates clearly that policy-makers can be censured both for appearing to take too many risks with the public, and for taking too few -- even in a case where apparently no one knows how large the risk really is.

At the heart of this inconsistent attitude to risk lies an unwillingness to acknowledge that life is unavoidably uncertain. We seem to reason that anything that happens must have had a prior cause; and its cause must have been either initiated by someone or detectable by someone; therefore, explicitly or implicitly, someone is always responsible for misfortune.

A frequent reaction to a tragic accident is to appoint a Commission of Inquiry to determine the cause. The commission seeks the chain of prior events that led to the accident. Perhaps there will be one sufficient condition, one prior event or act that by itself was enough to ensure that the accident would occur; more frequently there will be a list of necessary conditions, events or acts that could not have caused the accident by themselves, but that could have prevented it had they gone otherwise. It is the identification of a chain of necessary conditions that permits us to say “if only...”. If a set of prior causal conditions can be found, then it may be possible to find a way to prevent that precise path from happening again. Finding the conditions also allows the identification of the person or persons responsible, in order to assign blame or punishment.

Eliminating that one path can not, of course, protect against all of the other possible paths -- especially the ones that have not happened yet and that no one has thought of. And there is ample evidence to indicate that many of the worst accidents and disasters occur in just that way, through chains of events that are novel and unexpected. But it may nevertheless be useful to root out the source of danger that has been realized, for the fact of its having happened once may be evidence that it is likely to happen again. If that is the case, then by removing it we are significantly increasing our overall safety.

But it is also possible that the observed chain of events leading to an accident was truly rare and unforeseeable. This possibility calls into question both the efficacy of defending ourselves against what has already happened, and the justice of seeking someone to blame. What do we say, when some surprising misfortune happens -- a dam bursts, or a train derails, or a child is gravely injured by a reaction to a vaccine? That its occurrence reveals at least a neglect of responsibility, if not reckless or malicious disregard for people's safety? Or that we have observed a dreadful but uncommon and unforeseeable event? If the former is true, then the responsible person must surely be sought; if the latter, then there may really be no responsible person, even though there may be many who might have deflected the causal sequence of events.

If, as seems likely, the degree to which misfortunes are humanly caused or foreseeable varies from 'entirely' to 'not at all', then the degree to which it is appropriate to affix blame must vary accordingly. In the wake of terrible accidents, the dominant tendency surely is to affix too much blame. We often demand in retrospect a standard of vigilance and risk avoidance that would not have been possible in advance. Not knowing in advance where or how an accident is to happen, one can not know in what direction to be vigilant, and it is demonstrably not possible to pay full attention to every possibility at all times.

5) Analogies.

These reflections suggest three general, seemingly commonplace observations about the nature of risk, each of which has thought-provoking consequences and suggests courses of action markedly different from current practice. These stem from analogies to other systems or areas of activity in which risk must be routinely handled -- to games, business, engineering systems, and biological systems such as species and ecosystems. The purpose, however, is not to trivialize or be flippant, but to advance a search for parallel structures of risk that offer insights into the limitations afflicting any attempt to cope with risk, and for attractive strategies for dealing with risk.

The first observation is that it is not possible to be free of risk, or even to know precisely the magnitude and nature of the risks you face. This is a statement of what can not be done: you can not make a system or policy perfectly free of error or risk. What then can be done? A more achievable goal is to make the system tolerant of error. If you can not ensure that nothing will go wrong, then a next-best choice is to limit the consequences of things going wrong.

This approach might be described as a "safe-fail" approach, as opposed to the "fail-safe" approach that seeks to prevent errors. C.S. Holling, an ecologist at UBC, tells the following story to distinguish the two approaches. Early in the life of the DC-10, it became apparent that there was a dangerous design defect in the locking mechanism on the cargo door. It was possible to close the door in such a way that it seemed secure but was not. When improperly closed, the door could fly off in flight, and the resultant sudden de-pressurization of the hold could cause the cabin floor to collapse, crushing the hydraulic cables that control the wing and tail surfaces. At that time, DC-10s were operated both by civilian airlines and by the military, and the approaches to the hazard

that were taken by the two groups epitomize the difference between fail-safe and safe-fail approaches.

The civilian airlines instituted training programmes for cargo handlers, attempting to ensure that the person closing the door could distinguish between proper and improper closure and would close it properly every time. The military, on the other hand, drilled holes in the floor between the cabin and the hold, so that sudden depressurization would not cause the floor to collapse. Rather than attempt to ensure that errors would not occur, they made a design change to render the consequences of errors relatively innocuous (in a military setting).

Both sets of aircraft owners were dealing with a particular, known cause of risk. It is in defending against a particular hazard, whose probability need not be known but whose mechanism of occurrence is known, that the safe-fail approach is easiest to apply. Error tolerance is frequently engineered into the control systems of complicated pieces of equipment such as aircraft, automobiles, or computers, to cope with the most likely kinds of operator error. If your word processing programme makes you ask twice before erasing a large document, it is demonstrating a safe-fail approach to risk.

But how can a safe-fail approach be taken to risks that are unknown both in size and origin? If it is not possible to anticipate every kind of error or hazard that may arise, then it is impossible to take specific measures to reduce the destructiveness of each one. In this case, the safe-fail approach is followed by limiting your possible losses through diversification.

An investment fund manager faces a problem of this kind in deciding what investments to buy. Whatever particular security is chosen has a risk of losing value, and there is no way of telling in advance which will gain and which will lose. Managers cope with this risk by spreading their funds among many different investments, expecting that factors causing losses in some will not affect all, and so the losses in some will be more than offset by the gains in others. In the absence of risk, the manager would simply put the entire fund into the optimal investment, the one with the highest overall return. But this strategy, which is optimal under certainty, is extremely foolish under risk.

Biological systems such as species and ecosystems must also cope with risk of unknown origin and size, and also manage it with strategies of diversification. In species, this can be most clearly observed in those dwelling in harsh environments. Such species often divide themselves into geographically separated sub-populations. Harsh conditions then may drive one or more of the sub-populations to extinction, but are less likely to extinguish the entire population. When one habitat is emptied by extinction, it is normally re-settled by colonization from the surviving populations, and the diversity is restored.

A similar diversification exists at the level of ecosystems. Ecologists have long been puzzled by the fact that there seem to be many more species of plants and animals than are necessary to do every "job" in an ecosystem. A simple-minded application of the theory of natural selection, without considering uncertainty, would suggest that there exists one species optimally adapted to fill each function, and that it eventually should come to dominate and drive its competitors to extinction. Not so. For each function, there are normally several different species, each of which is slightly better adapted under

certain conditions. As the environment changes, the advantage periodically passes from one species to another, but conditions are rarely stable for so long that species are driven to extinction. The consequence of this diversity is that if a species is wiped out by some hazard, the entire ecosystem does not collapse; there are other species around to do the essential work. As with investment portfolios, the strategy that would be preferred under certainty is rejected or at least modified in order to provide resilience against unforeseeable hazards.

The lessons of these analogies to public policy are clear, and should provide a cautionary warning to the more fervent disciples of optimization and specialization. If some activity is subject to unforeseeable shocks, then it is best to have several different ways of doing it. We should not keep all our eggs in one basket. Of course not, comes the reply; any fool knows that. But this reply fails to recognize the extent to which demands for efficiency and competitive pressures militate for us to put all our eggs into one basket. After all, one basket is demonstrably cheaper and more convenient than two -- as long as you do not drop it.

There are many particular examples. An economy that is dependent on one source of energy will be vulnerable to price or supply shocks in that source. An economy substantially dependent on a particular skill, process, or labour force, as was the British economy upon coal miners, could be held to ransom by them. An agricultural or forestry sector based on a small number of crops cultivated in huge monocultures, or a fishery relying exclusively on interbreeding hatchery stock, is vulnerable both to biological hazards and to fluctuations in the markets for the crops. If substances of potential health or environmental hazard are used in large quantities throughout the country or the world, then any harm that is realized might be on an enormous scale. And an economy dependent on foreign trade with one other nation is subject to disruption by changes in relations with the trading partner or its economy.

In all of these situations, the market pressures that militate for the most efficient, most profitable way of doing something in the short term -- and the static analysis which supports those pressures -- may contribute to an unstable, vulnerable system in the long term. A mix of activities -- an industry based on many different fuels, an agriculture producing many different crops, a less-than-fully specialized economy trading with many different partners -- may be less than optimal under certainty, but vastly more capable of handling risk and offering some robustness against the sharks of a dynamic trading environment.

The second observation is that by its very nature, risk will be borne inequitably. There are two kinds of inequity associated with risks. The first is that the probability of harm resulting from any particular decision is normally greater for some than for others, greater for example for uranium miners, or for the people who live near a chemical plant. But even when the distribution of prior probabilities of harm is or can be made equal, the random nature of risks ensures that the realized, or posterior distribution of suffering is always profoundly inequitable. Misfortunes come in large doses. While it may be true that the new factory imposes a risk of one chance in a million of cancer on each of the city's one million inhabitants, it will not happen that everyone gets one-millionth of a cancer. One person (on average) will get a whole cancer, and that person will have good reason to rue the decision that brought the factory to the city.

The inevitable inequity in the incidence of harm from risks argues for society's sharing some of the burden, but the possibility and the appropriate level of sharing are clearly different for different types of risk. When a risk is taken on voluntarily and knowingly by individuals, the social obligation to share the burden is limited. For economic risks, one of the bases of the market economy is the ability of individuals to take risks and bear the resultant gains or losses. It is through the discipline of gains and losses to individuals that the economy adapts to changing technology and tastes, and so remains globally competitive. Moreover, total protection against economic losses would give individuals incentive to reckless behavior, in the expectation that they can receive the benefits of their economic risks if they win but not bear the costs if they lose.

But few would argue that society should abandon even those who make bad economic decisions to starvation; there must be some measure of safety net or protection from deprivation. It is the issue of the appropriate degree of protection, obviously, that is difficult and contentious.

Does the same difficulty apply to physical risks taken voluntarily by individuals? On the one hand, no reasonable person would suggest that hikers who become stranded on mountains that are beyond their ability should be abandoned to their fate; on the other hand, if it is known that society will spare no expense, and even put other people at risk, in order to save individuals who have placed themselves in peril by reckless risk-taking, then one can expect insufficient prudence to be exercised in deciding what mountain to climb or whether to go sailing in bad weather. There is a particular tragedy in the stories of a helicopter rescue team dying in the effort to rescue a climber with a sprained ankle, or the deaths of eleven people searching for the wreck of a plane with two people aboard.

When risks are imposed on individuals involuntarily, then the obligation of society to protect those who suffer the consequences is presumably greater. If such risks are taken on by a valid social decision that the resultant benefits exceed or justify the risks, then there exists a social surplus, so society has also the means to compensate those harmed. The obligation is partly pragmatic, for no group is likely to tolerate being placed at risk unless they believe that their burden will be in some measure shared or insured. But the obligation is also in part ethical, for it is unjust that some should be forced to bear uncompensated losses to advance the welfare of others. This is particularly true when the probabilistic prior risk burden is inequitable as well as the posterior, after the fact, distribution of harm.

These two cases, risks taken on voluntarily by individuals and risks taken voluntarily by society but imposed involuntarily on individuals, may define different limits of social responsibility to the victims. But the determination of how voluntarily or involuntarily an individual took on a particular risk is very difficult; there are many possible degrees of knowledge and informed consent, and individuals seeking compensation have an incentive to over-state the degree to which the risk they faced was involuntary.

Moreover, it may not be true that any risk assumed collectively generates a social surplus through which the losers can be compensated. Even if the decision was correct in terms of expected surplus, there may not be any realized real surplus. If the "thousand-year accident" at your new LNG terminal happens in the first year of operation, then your

social surplus will be a long time coming. Alternatively, the social decision may simply turn out to be incorrect, and so there may not even be an expected surplus. And some risks -- earthquakes, floods, war, global environmental hazards and international economic crises -- obviously are involuntary even at the collective level, to the extent that these are unmitigable by the policy decisions of a single nation. In any of these cases, the losses of the victims may exceed the total ability of society to insure them.

It must also be recognized that the foregoing arguments for compensation only apply strictly in the case of purely economic risks. Economic or property losses are compensable by payments of money or goods, but losses of life and health are not. No amount of money can truly compensate for the loss of a leg, a disabling handicap, or the death of a loved one. This renders the inequity of the consequences of risks more poignant, but no less inevitable.

Recognition of the impossibility of true compensation may be behind the present tendency toward extravagant monetary awards in court cases for personal injury and death. Such awards can be interpreted as attempts to approximate the infinite payments that would be required for true compensation, or alternatively to exact vengeance or to send a message to designers and decision makers. But such awards cannot be justified on any of these grounds. Exacting vengeance for every harm suffered, given the inevitability of some harm, would involve all of society in a blood feud. And although a punitive settlement might be an appropriate means of sending a message that excessive risks were taken, it is not clear that a court has the competence or authority to determine what sizes of risks are acceptable. Even if decisions are made that impose society's ideal level of risk, there will still be victims. How can we face these people, and what can we say to them?

One thing that can be done is to recognize that while some losses cannot be compensated, some can. Medical care, therapy, and replacement of lost income are matters for which compensation to victims can be made, even if pain and suffering, disability, and loss of loved ones are not.

The third observation is that when we take risks the uncertain consequences of our actions might be benefits rather than losses. By shunning uncertainty as much as you can, you count yourself out of the gains as well as the losses. It is clear that individuals know this, for they take risks economically all the time for the prospect of resultant gains; this is the essence of innovation. They also take risks in order to learn; this is the essence of experiment. People can also be observed to accept risks to their life and health in return for financial or other gains; for many people, the benefits of sky-diving and motorcycle riding are clearly sufficient to offset the substantial associated risks.

If economic gains, learning and pleasure can all justify physical risks for individuals, can not the same argument be made for collective decisions? To pursue the minimum possible risk is to expend resources in the pursuit of a phantom, to reduce adaptability, and to miss opportunities for learning. Successfully minimizing risk in a certain policy area would require devoting all your attention to analyzing the current state of the world and fine-tuning the current optimal solution. But the world is always changing, and such a strategy leaves nothing available with which to anticipate or respond to these changes.

This may sound like a fine philosophy, but it cannot be employed if any error is catastrophic or is punished severely. If any risk that you take carries a significant probability of disaster, then you would be a fool to do other than attempt to minimize risks. But new information and learning are only gained through taking risks. How can this dilemma be resolved?

Attempting to improve policy-making through experiment and information-gathering is a problem in social learning. The problem of limiting the consequences of error also arises in many individual learning situations, particularly in the learning of skills for which the cost of failure in practice can be very high -- for example, driving, flying, or rock climbing. Teachers of these skills have developed a number of devices to increase error tolerance.

Teachers of climbing encourage students to move quickly onto rather difficult faces, with the student roped in and supported by the teacher so that it is not possible to fall more than a few inches. Under these conditions, the student can learn to feel the difference between doing it right and doing it wrong without fear of making mistakes. "Spotting" in gymnastics teaching, and dual controls and simulators in pilot and driver training, provide the equivalent environment. The student can be then encouraged to experiment, and can gain information and skills far more rapidly than in an environment in which failure must be prevented from the outset.

Nature even appears to exploit the opportunities inherent in uncertainty by making experiments. The 'Hopeful Monster' theory of evolution, one of several currently popular theories, suggests that major evolutionary advances are made by occasional large mutations that result in the birth of offspring significantly different from their parents. Most of the time the experiments fail -- the mutation is maladapted and dies. But occasionally, the adaptation is beneficial and the hopeful monster thrives.

The challenge for public policy on risk is to find the socially optimal level of risk taking, somewhere between rigid caution and outright recklessness, and to ensure that as much as possible, risks are taken in an environment of error tolerance. To a great extent, our economic and political institutions are established upon assumptions that errors can be tolerated, in the investment of decision-making power in non-experts through juries and elections, and in the frequency with which process is regarded as more important than outcome. There are, though, at least two respects in which the current environment for public policy decisions is not adequately tolerant of error. The first is small and homey, the second vast and remote. The processes of internal and external evaluation and audit to which the public servant's decisions are increasingly subject create a world in which even small failures may be punished severely, while no corresponding rewards are given for successes. The consequence is that the public servant's environment is most error intolerant in terms of personal reward, and extreme avoidance of risks is encouraged. The process of "scapegoating" which follows both Auditor-General's reports and accident inquiries has surely pushed much too far in the direction of repressing both innovation and a responsible acceptance of risk. And this must surely be one key respect in which the environment for managers in private sector activities is perceived to be much different from that in the public sector.

On a much larger scale, certain forms of human intervention or potential intervention in the natural systems of the earth are now occurring to an unprecedented extent, comparable to the scale of the earth itself. There are at least three such areas of global-scale disruption: the increase in atmospheric carbon dioxide and resultant possible climatic changes; the disruption of the ozone layer and consequent potential increase in ultraviolet radiation at the earth's surface; and the nuclear arms race. While it can be argued that any new human activity -- any new chemical, industrial activity, or intervention in a natural system -- has a small probability of having unforeseen globally catastrophic consequences (the stuff of horror films -- the mad scientist who creates a bug that destroys life on earth), it is clear that for most activities this probability can safely be regarded as vanishingly small; prior similar experience, or theoretical knowledge, or the small scale of the activity renders it overwhelmingly likely that the world can endure our puny insults.

But the usual sources of comfort do not apply to these processes just mentioned; they are all proceeding on such a scale, and in such ignorance, that they have the clear potential for globally catastrophic consequences -- the ultimate in error intolerance. And the marvellous technological advances that have made such activity possible have in no way put our activities out of reach of error; examples of major errors and failures in even the most advanced systems appear with shocking regularity. These cases of error intolerance are beyond the power of any single nation to address, but are of urgent importance nevertheless.

Our collective decisions, intentional and unintentional, are thus a confusing mixture of too much risk and too little. While the global aggregation of processes for which no overall individual responsibility can be assigned leads us to court unimaginable risks, the reluctance to acknowledge the inevitability of risk and the excessive post hoc assignment of blame leads many individuals and public servants to excessively shun risks. In some areas we even forbid individuals to assume risks that they might wish to take on in full knowledge, as in the lengthy testing and licensing process for new prescription drugs and the restriction of health care practice by non-physicians. In other areas we give incentives for reckless risk-taking by insuring individuals against economic losses, as when the federal government insures large investors in failed banks that had been offering high rates of return in order to support a high-risk lending policy.

6) Conclusion -- A Healthy Attitude to Risk.

These questions are very broad, but they are ones that must more and more be faced, even in deliberations over specific, practical policy decisions. We are being driven to face them by the increasing scale of human intervention in global systems and the apparent accompanying probability of catastrophic consequences of error; by the growing chorus of demands that individuals be completely protected by society, from physical risk and even economic losses; and by the growing sensitivity of our ability to detect hazards, which confronts us forcefully with the impossibility of eliminating completely the risk consequences of our activities.

What would be the properties of an ideal social relationship to risk? The most important one is error-tolerance. There can be no justification for pursuing courses that

have any significant probability of leading to global catastrophe. Since we have no basis for assuming that we will not make errors or be unlucky in our activities on any scale, we must seek to ensure that the potential consequences of our errors are within the capacity of society and ecosystems to cope.

Error tolerance argues for diversification in all realms of activity -- in industrial strategy, agriculture, foreign trade, and environmental intervention. When, for example, a new activity or substance of small but unknown risk is considered, it might be treated in a number of ways. If the risk can be effectively confined to those who choose it voluntarily, as with the introduction of new prescription drugs, then a less paternal approach that focuses on informing potential consumers about the current state of knowledge regarding its potential risks and benefits is likely preferable to restricting its use. There would be no doubt then be large individual variation in chosen degrees of risk-taking, with benefits both to individual welfare and social knowledge and adaptability. Perhaps this is one answer to the ethical problem of experimenting on humans; only experiment on those who choose to take the risk when fully informed of the current state of knowledge (though that in itself may be quite incomplete and subject to debate).

If the risk is environmental, perhaps the objectives of economic progress, information gathering and public protection would best be served by regional licensing, or by licensing different substances in different regions. This does amount to experimenting on people, but when the needed information about a potential (but small) hazard can only be obtained by putting it into use and seeing what happens, the only difference between this proposal and current practice is that now we make fewer but larger experiments; we experiment on everyone instead of on just one region, and we either win big or lose big. Introducing new hazards slowly and on a limited scale might delay the adoption of innovations, but it would at least provide a far better opportunity to catch significant hazards while limiting their impact.

The second property of a reasonable social attitude to risk would be a recognition, built into institutions, that some risks are inevitable and that some level of risk may justifiably be assumed for accompanying benefits. Such a recognition would result in less focus on post hoc blame for realized misfortune and more on determining whether the risks taken were of a socially acceptable size and determining where risk reduction effort has the greatest chance of doing significant good.

Finally, there must be recognition that there is an appropriate degree to which society can and ought to take the burden of risks off individuals, and that the determination of the degree of social responsibility for individual suffering is a matter that cuts to the heart of the ambiguous nature of social organization. Analogies to other human activities or biological systems are of no help here, for none of these systems exhibits the same tension between welfare of the individual and welfare of the collectivity as does human society.

In species, ecosystems, and stock portfolios, decisions appear to be made at the aggregate level for the maximum benefit of the totality. The dollar that is invested in a speculative stock and is lost has no rights and no legal recourse against the fund manager; nor has the maladapted mutation against the species. But human society values

individuals, gives them rights entrenched in charters, and endeavors to make them the principal locus of decision-making through decentralized institutions of control such as elections and markets.

The conflict between protecting individuals and benefiting the collective arises in all aspects of public policy, but is particularly poignant on matters of risk. Neither extreme position is tenable. Because individuals matter, a position of social darwinism that cavalierly dismisses individual suffering in order to maximize the welfare of the surviving members of society is clearly unacceptable. But neither is it possible for society to completely insure individuals against all losses suffered, both because true compensation for physical suffering is impossible, and because the attempt to compensate all losers may exceed the adaptive capacity of society as a whole.

The resolution of this dilemma must acknowledge the existence of a “moral hazard” problem in risk-taking. Just as post hoc inquiries assigning blame for non-predictable failure influence people’s behavior to make them excessively shun risks, so does the expectation that losses will be insured or compensated influence people’s behavior to make them excessively take risks. This is clearly the case when individuals are led to expect that they can retain the benefits of success from economic risk-taking without having to bear the costs of failure. It has occasionally been suggested that large damage settlements against property-owners or municipalities for injuries sustained on their property may encourage too much risk-taking by individuals, but this seems unlikely; the prospect of monetary settlement for disability can surely not make the prospect of disability seem any less horrible if considered in advance. But they do promote too much risk avoidance on the part of the people held liable; it surely can not be socially desirable that cities close their playgrounds out of fear of liability suits.