Scientific uncertainty, technological unreliability and public policy

Presented at Workshop on “Risk Management”
5 – 7 October 2000, Naples

European Science Foundation
Standing Committee for the Humanities

Dr Chris Elliott
chris.elliott@pitchill.com

Contents

Acknowledgements and disclaimer

Section 2 of this paper, dealing with scientific uncertainty, is based on work carried out by the author with OXERA (Oxford Economic Research Associates Ltd) and Safetycraft Ltd for a group of Departments and Agencies of the UK government. The group was chaired by the Health and Safety Executive, and included the Cabinet Office (Regulatory Impact Unit), the Department of the Environment, Transport and the Regions, the Department of Health, the Environment Agency, the Food Standards Agency, the Office of Science and Technology, and the Scottish Executive.

Section 3 of this paper, dealing with technological unreliability, is based on work carried out with Risk Solutions Ltd for the Safety and Standards Directorate of Railtrack plc, the company responsible much of the UK rail infrastructure.

The author is grateful to his colleagues and sponsors for their help, but none of the opinions and assertions of this paper necessarily represent their views. The author is solely responsible for the content of this paper.

1 Introduction

Context

It is a truism that the application of science, in the form of technology, brings dangers as well as benefits. From the first use of fire through to nuclear power, from primitive agriculture through to genetically modified seeds, society has had to decide whether the benefits are sufficient to outweigh the dangers.

We are used to taking risk decisions for ourselves and our families. We decide whether to overtake when driving, whether a piece of chicken is cooked enough, or whether to insure a valuable clock. We understand (perhaps unconsciously) the need to balance the risk that we may suffer harm with the benefits of accepting that risk.

There however very many risks where others must make that decision for us. We may not have the expertise to take the decision, or we may not be free to exercise choice independently from other members of society. Someone else must decide whether to license a new drug, which rail safety system to install or whether to permit GM crops to enter the food chain.

The challenge to those in authority is to take a decision which is:

- rational
- defensible; and
- equitable.

This challenge is made much harder because the decision must be made in the absence of certain knowledge of the consequences. **Scientific uncertainty** exists where the science underpinning the decision is not complete, or is disputed. **Technological unreliability** remains when, although the science is not in doubt, the limited reliability of systems leaves a statistical probability of harm.

This paper is concerned with the **governance** of science and technology. Proper governance takes due account of the costs and benefits of actions, and of the need for decisions to be rational, defensible and equitable.

Issues

A rational, defensible and equitable decision on matters of risk needs to address two questions:

1. how great are the benefits and dangers?
2. do the benefits outweigh the dangers?

These are very different questions. Question 1 is in principle amenable to objective and quantitative assessment, although the accuracy of that assessment will be limited by scientific uncertainty. Question 2 is essentially subjective, rooted in the culture and attitudes of society to technological risk. It involves judging the relative merits of options (including “do nothing”) and balancing the possible positive and negative outcomes of each option.

Democratic society looks to its governments either to take the decisions or to regulate those who take them. Section 2 of this paper explores two aspects of the governance of science and technology by governments:

- how governments should obtain and use scientific advice as an input to Question 1, and
- where responsibility for Question 2 lies with government, how should it take the decision.

The decision to exploit science in the form of technology is often taken by industry. Governments can delegate responsibility for answering question 2 to industry, and their role is solely to set a framework of social values within which industry decides what is acceptable. Section 3 of this paper examines some of the issues that should be included in a framework.

2 Scientific uncertainty

Background

This section of this paper is a summary of a major research project. The full text of the final report of that project may be downloaded from:

http://www.hse.gov.uk/research/crr_pdf/2000/crr00295.pdf

Scientific advice

Government has to take decisions in which scientific uncertainty contributes to the risks that might flow from those decisions. The person who takes the decision will need to take account of the science (and its uncertainties) but will also need to consider many other issues. However, the scientific input is special, because:

- science is subject to disciplines, including the use of evidence, and
- it offers reliable techniques for predicting future outcomes based on rational models that follow well-tested natural laws.

Scientific advice which contributes to risk decisions may be based on:

- **Observation**: empirical evidence that is unambiguous and not contentious, although may be open to different interpretations
- **Formal analysis**: which should lead to a consistent result, regardless of who conducts the formal analysis;
- **Reasoned judgement**: the outcome of a disciplined approach to a problem, where deductions are made by extrapolation or extension from formal analysis
- **Opinion**: an assertion of belief, the value of which depends solely on the integrity, competence and credibility of the person expressing it.

There are many mechanisms for the provision of scientific advice, including advisory committees, academic experts and scientists employed by government, industry or pressure groups. All are valuable but there is need for clarity and discipline in the way that they are used.

A disciplined approach to scientific uncertainty

Differentiation of roles

A policy decision involving scientific uncertainty can use contributions from four different actors. The first three are all engaged in addressing Question 1 – "how great are the benefits and dangers?" - and the fourth then addresses Question 2 – "do the benefits outweigh the dangers?"

- **policy maker**: a person or organisation charged with assisting a decision taker in reaching a decision by providing policy analysis, generating policy options, or by conducting risk assessment;
- **scientific adviser**: a person or organisation responsible for providing scientific input to policy making or decision taking. This includes both scientists expert in narrow disciplines relevant to the problem in question, and more broadly-based scientists able to integrate several disciplines, and those within and outside the government service;
- **stakeholder:** a person or organisation representing the interests and opinions of a group with an interest in the outcome of a policy decision.
- **decision taker**: a person with the authority to take a policy decision. This may be a government Minister, or a person or body with the delegated authority to take a decision in the name of a Minister.

It may not always be possible to have each of these functions carried out by a different person or body, but, when one person is carrying out more than one of these functions, he should be qualified and authorised to do so and should understand in which capacity he is acting at any time.

A clear structure differentiating these functions would allow a more transparent, reliable and trustworthy process, which gives due recognition to the rights of those at risk from decisions.

Choice and duties of scientific advisers

There are no truly independent scientists; everyone has individual biases and personal motives. The biases of competent advisers must be identified and managed, but cannot be eliminated

Scientific advisers are to act purely in their capacity as professional scientists and should exclude all extraneous influences of which they are aware. When in doubt about where their primary duty lies, they should always remember that when giving advice to government on risk issues, their duty is towards the broad public interest, and that takes precedence over any personal or professional consideration. This empowers individual advisers to resist any inappropriate external pressures.

Scientific advisers should not seek to hide uncertainty or to apply hidden "safety margins" or the precautionary principle, since these would be straying into policy making and the judging of societal values. Neither should advisers seek closure on issues where there are genuine differences of view. They should faithfully record and present to the policy maker (and hence the decision taker) the alternative views and the evidence to support them.

Rights of stakeholders

Stakeholders should contribute to many stages of the decision-making process, including:

- identifying the issue
- defining policy questions and possible options
- framing of scientific questions
- selecting the advisory mechanisms and advisers, and
- assessing the findings.

Their involvement is especially important where the potential benefits do not fall to the same stakeholders as those who might suffer.

However, stakeholders should not be **represented** in the generation of the scientific advice. Their scientists might be involved, but as competent scientists, not representatives.

Central role of policy makers

The generation of a set of viable policy options is the final stage of a process during which scientific advice is sought, generated and used. This is a skilled professional task, which falls naturally to civil servants, especially those with a science background. They have the expertise to pose questions that are amenable to scientific answers and to build those answers into policy.

This is not an academic process. It is concerned with management of real and important issues. The policy makers should think in terms of policy options at all stages in the advisory process. The questions put to scientific advisers will often require them to consider the risks involved if different policies were followed, and to communicate any additional options that they identify.

The policy options embody the answers to Question 1, including the range of uncertainties in each of the costs and benefits.

Actions of decision takers

Decision takers are responsible for answering Question 2 and must be helped to handle complex social and scientific issues. They rely on policy makers to prepare a set of policy options and to gather and assimilate the scientific and other inputs (such as economic impacts, social acceptability and ethical considerations) which affect it.

Scientific uncertainty (and the risks it causes) must be retained in the policy options considered by the decision taker, because it is a vital and irreducible element of the issue. The decision taker should be presented with the "unvarnished" truth, as best the scientists can estimate, including uncertainties and differences of opinion.

Even though the analysis and preparation of the policy options is carried out by others, decision takers must take responsibility so must be prepared to challenge the advisers to satisfy themselves of the integrity of the advice.

Principles

This disciplined approach may be codified as seven principles, which should underpin any process of seeking and using scientific advice:

1 Scientific advice should be founded upon evidence (observation and formal analysis), and should describe both the scientific conclusions and their uncertainty. Where there is doubt, the advice should be deduced from the evidence by reasoned judgement.

2 The functions of scientific adviser, policy maker, decision taker and stakeholder should be distinguished.

3 Full information about the process of seeking and using scientific advice should be made public, and stakeholders should be encouraged to comment. Scientific advice itself should be published promptly and transparently once it is complete.

4 Scientific advisers should be selected for their competence. It is not necessary for advisers to be independent of all interests in the policy question; however, the advisory process should ensure that interests are declared and any

resulting biases are balanced or taken into account.

5 The scientific adviser’s over-riding duty is to the public interest. Advisers should not promote any special interests in preparing their advice.

6 Scientific advisers and policy makers should be candid about the limitations of scientific knowledge and should always assess the uncertainty in scientific advice and the risks associated with each policy option; this information should be taken into account by decision takers.

7 The effort expended on securing scientific advice should be proportionate to the importance of the policy issue and the difficulty of the scientific investigations required.

3 Technological unreliability

Context

Section 1 of this paper identified two questions which must be answered when deciding whether to accept a risk:

1. how great are the benefits and dangers?
2. do the benefits outweigh the dangers?

Section 2 looked at some of the ways in which scientific uncertainty may be addressed when answering Question 1, and at the responsibility of government when building that advice into a decision for Question 2.

This section looks more closely at the ways in which industry could answer Question 2, whether the decision is:

- to permit a new technology which causes risk, or
- to invest in a new technology which reduces an existing risk.

A structure for addressing technological unreliability

Rational decisions require a structured approach.

The first element of that approach is to identify **who** should take the decision. This includes the **organisation** which has responsibility for the decision and the **person** or **office** within the organisation where the formal accountability lies. Although that person may seek assistance from colleagues, or other organisations, he or she is ultimately responsible. Safety responsibility may not be delegated, although it may be shared.

The second element is to establish a procedure by which the decision will be taken. Ideally this procedure should be made public and agreed by all interested parties, and should be used consistently over many decisions. Although decisions might have to be based on the subjective judgement of the decision taker, the procedure should take the decision taker through a logical structure, for example to address his mind in sequence to three questions:

1. what **factors** should be taken into account? eg:
   - precedents of other changes
   - public policy
   - statutory duties
   - public opinion
   - recent accident history
2. what **criteria** should be applied? eg
   - financial cost and expected benefits of the technology
   - cost of harm caused or avoided
   - relative merits of avoiding fatalities, major injuries, minor injuries, economic losses
   - robustness of solution and of analysis
3. what **values** should be applied? eg
   - protection of customers
   - protection of employees
   - protection of others (eg innocent members of the general public, or trespassers)
   - economic impacts
   - public perception
   - political impacts

This approach ensures a clear separation between three different issues:

- what is the objective of the technology?
- how can we quantify the benefits and costs?
- which technology shall I support?

This separation allows the decision taker to include the external costs of a lack of safety (such as the loss of confidence by the public) and to make explicit the basis on which the decision taker has attributed relative values to different outcomes (for example, the relative merits of avoiding a small number of deaths or a large number of minor injuries). Also, where the industry is spending public money (for example rail subsidies), it offers an argument against wasting public money by seeking excessive safety.

Legal principles

Where decisions affect more than one country within Europe, the issue is made more complicated by two fundamentally different legal principles.

The countries which adopt the common law tradition (UK and Ireland in Europe, and others including the USA, Australia and Canada) follow the principle of **literal** interpretation of legal statutes. The statute is worded very carefully to say exactly what is required to be done or not done, and the courts look to see whether a defendant has complied with the **letter** of the law.

An historical example was where an English court found that someone who had voted in an election by pretending to be a person who had died was not guilty of breaking the law which forbids impersonating “any person entitled to vote”. This was because a dead person is not entitled to vote.

The countries that follow the civil law tradition, or Code Napoleon, (including France, Italy and Spain) follow the principle of **purposive** interpretation. Statutes are written simply and in an abstract way, and the courts seek to establish whether a defendant has complied with the **spirit**, or purpose, of the legislation. For example, Article 3.1 of the Código Civil of Spain states:

> *Rules are to be interpreted according to the basic sense of their words, in connection with the context, the historic background and the legislative record, and the social reality of the time in which they are to be applied, paying fundamental attention to their spirit and aim.* [1]

European law, as construed by the European Court of Justice, follows the purposive approach. This brings problems for the national courts of common law countries when interpreting national statutes which give effect to EU directives.

The two legal principles lead to different approaches to drafting legislation relevant to Question 2. This is illustrated by the wording of section 3(1) of the UK Health and Safety at Work etc Act 1974:

> *It shall be the duty of every employer … to ensure, so far as is reasonably practicable, that persons … are not … exposed to risks to their health and safety*

and by the Italian principle of “buon andamento”, which requires the administration to act:

> *to produce the best possible result at the lowest cost*

The UK law includes the concept of “reasonable”, thus allowing the courts to decide whether an action was correct in context. The Italian principle requires, for example, that the risks from the railway shall be “minimised”. That would be unworkable in the UK because the only way to comply would be to stop the trains – this is the only way to **minimise** the risk. The Italian courts however will interpret the spirit of the law, to apply an implicit requirement for “reasonable” behaviour and not to demand perfection.

It can be seen from this that the two legal codes end up at the same point – either via explicit reasonableness and literal interpretation or explicit perfection and reasonable (purposive) interpretation.

Some approaches to answering Question 2

No increased risk

The simplest method is to permit any technological innovation which either does not increase risk, or increases it by a negligible amount. This is codified in France as GAME (Globalement Au Moins Équivalent)) and in Germany as MEM (Minimum Endogenous Mortality).

This sets a threshold or acceptability, but does not provide a basis for deciding whether a technology which passes this elementary test is acceptable.

Economic assessment

One could accept any new technology which:

- brings greater economic benefits than the economic value of the harm that it causes; or
- reduces the economic value of existing harm more than its financial cost.

In practice, the principle of no increased risk precludes the acceptance of a new technology which brings **any** significant risk of harm, and the simple financial approach is only applicable to deciding whether a change to reduce risk can be justified. Even then, it is not enough to say “I will accept all changes that exceed a threshold level of benefit/cost ratio”. The budget is finite and all changes that can be justified on simple grounds may not be affordable.

All parties have to agree the economic value of harm. This is not contentious when the harm is itself economic, but is harder when the harm takes the form of death or injury to people. A crude method is to attribute a financial value to preventing a fatality (VPF), and to express major and minor injuries in terms of equivalent fatalities.

There is general agreement in Europe that the VPF is around 1M euro. At least two models exist for expressing accidents as equivalent fatalities:

- in the UK, 10 major injuries or 200 minor injuries are taken as equivalent to one fatality when expressing human harm in economic terms;
- in Germany, 10 major injuries or 100 minor injuries are taken as equivalent to one fatality when calculating whether a change meets the MEM requirement.

A higher VPF may be needed to reflect the greater social impact of multiple deaths or injuries from a

[1] Las normas se interpretarán según el sentido propio de sus palabras, en relación con el contexto, los antecedentes historicos y legislativos, y la realidad social del tiempo en que han de ser aplicadas, atiendiendo fundamentalmente al espriritu y finalidad de aquéllas

single event, or where members of the public place trust in industry. For example, in the railways a VPF in excess of 1.5M euro is used, rising to several million euro for major accidents, and the airlines appear to use a VPF in excess of 10M euro. The MEM approach weights the impact of each fatality or injury in a major accident more heavily than it would be weighted if it occurred in isolation.

ALARP – As Low As is Reasonably Practicable

I travel frequently for work and often need tools, for cutting paper, sharpening pencils, connecting modem wires, opening equipment cases and even occasionally opening a bottle of beer. It would be useful to keep a knife in my briefcase, and I would be prepared to pay up to 20€ for a suitable knife.

That was the starting point of my consideration – the recognition of a **benefit** that might arise from access to a technology, in this case a knife. However, it is not a difficult step to realise that a sharp blade presents a **hazard** (defined as any situation that could contribute to an accident). I assessed the **risk** (defined as the likelihood that an accident will happen and the harm that could arise) and concluded that a bare blade presented an **unacceptable** risk.

However, I looked for ways to mitigate the risk. Pocket knives have a blade which folds into the handle when not being used. I concluded that the risk from such a knife was **tolerable**. I looked at pocket knives and found that I could buy a cheap one for 5€ or a Swiss Army knife for 15€.

Having established that there were several knives that cost less than the value to me of the benefits and for which the risks were tolerable, I looked at the **residual risks** to see if they could be further reduced. I judged that the cheap knife was likely to break, which might expose the blade, but that the Swiss Army knife was unlikely to break. I considered enclosing the knife in another container and even locking that container in case someone tried to take the knife to threaten me. I concluded that both of these additional measures were **unreasonable**, given their cost and the low level of risk that they removed, but that choosing the Swiss Army knife was a **reasonable** measure to reduce risk.

Another consideration was the need to comply with **regulations**. I checked that the airlines would allow me to carry my Swiss Army knife (provided that the blade is less than 10 cm long).

I also considered **alternatives** (such as a pencil sharpener – not much use on a beer bottle) but concluded that the disadvantages of each outweighed the slightly lower risk than that of the Swiss Army knife.

I concluded that the risks of carrying a knife had, by selecting a Swiss Army knife, been reduced to a level which was **as low as reasonably practicable**, and that, by restricting the blade length, it would comply with "legal" constraints.

The knife case illustrates all of the key concepts of what is probably the most complete formalism for answering Question 2. The principle of ALARP (As Low As is Reasonably Practicable) has gradually emerged from decisions in legal cases starting in 1938, and it has been extended to underpin the entire health and safety system of the UK and, at European level, other domains such as medical devices.

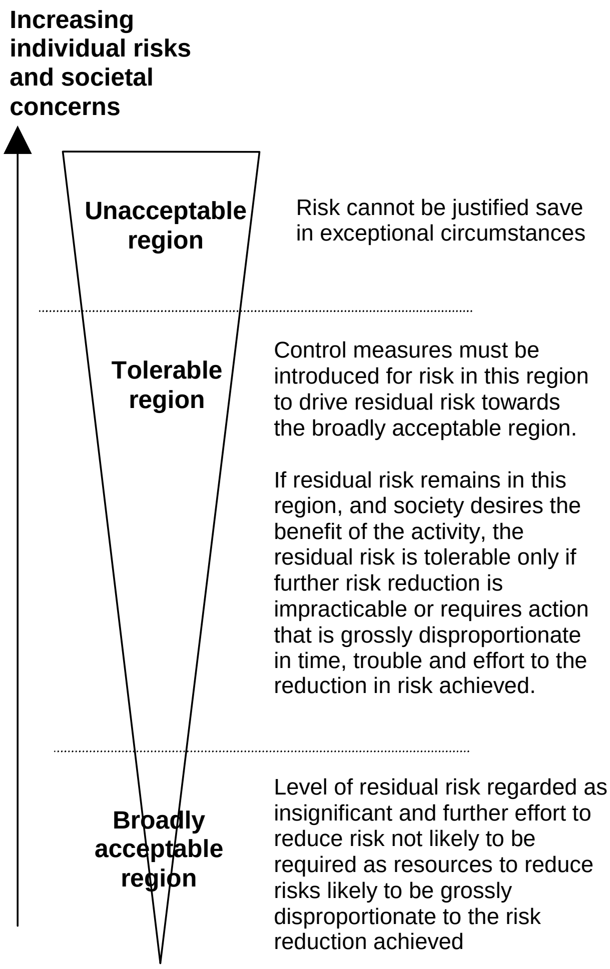

An illustration of ALARP[2]

ALARP applies once it has been decided that a proposed course of action brings benefits. Those benefits, and the risks associated with it, must first be quantified. The risks are then considered.

[2] taken from "Reducing Risk, Protecting People", Health and Safety Executive 1999

Taking the decision – the example of railways

This Section has illustrated three principles by which a decision may be taken:

- No increased risk
- Economic assessment
- ALARP

Each offers a rational basis for approaching the decision as to whether to accept a single technological unreliability. However, only ALARP addresses the case where there are several possible solutions, and even then it does not provide a complete answer. This is illustrated clearly by the issue of rail safety.

Railways are already very safe. By any objective measure, rail is the safest form of passenger transport within Europe. Despite this, there is great political and social pressure to make railways safer. Railways are also becoming more autonomous, either by privatisation or by becoming independent companies owned by the State, and divided, by independent companies having to share out the costs of safety. As a result, railway managers are increasingly required to decide where they should invest to improve safety, and to be able to justify their decisions.

It is political unacceptable for railways to introduce changes which increase risk. The “No increased risk” test is a minimum requirement.

Economic assessment is a key part regulated railway funding. Railways might not be permitted by their financial regulators to make investments in safety that do not pass an economic assessment. That can lead to conflicts, where the economic assessment fails to capture all of the political requirements for safety.

Railways have finite budgets. They cannot afford to make all investments which pass the economic assessment test. It is then necessary to decide whether, for example, to reduce the risk of high-speed accidents (which are rare but cause fatalities) or of low-speed accidents (which are more common but cause mainly minor injuries). ALARP provides a framework within which to take that decision, but leaves for the decision taker the subjective valuation of different kinds of harm.

A further complication for railways within Europe is the need to inter-operate. A train may travel from Berlin to Glasgow, passing through Netherlands, Belgium, France and England. The same train has to be acceptably safe in each country. Similarly, a signal control computer that is judged safe in Italy may have to be acceptable in Sweden, to comply with competition law. “Safe” has to be agreed between those counties, despite fundamentally different legal codes.

4 Conclusions

Summary

This paper has sought to set out the logic by which decisions on matters of safety of the public may be taken so as to be rational, defensible and equitable. It approached the problem by breaking the decision into two stages. The first deals with quantifying the risks in the presence of **scientific uncertainty**. The second deals with assessing the acceptability of the residual **technological unreliability**. Together, a formalism which addresses these establishes a **public policy** for the governance of science and technology by society.

The seven principles that should underpin the way in which government obtains and uses scientific advice here should lead to legitimate and democratic behaviour by governments. Their key consequences are:

- the decision will be taken by the appointed decision taker who is then responsible for that decision. It will not (either expressly or implicitly) be usurped by scientists;
- those on whom the consequences of the decision (both good and bad) will fall will influence the decision, not the scientific advice;
- the scientific advice will be provided by those with the correct expertise, without necessarily requiring them to be independent.

These consequences are themselves controversial. For example, there is growing interest in the development of ways in which stakeholders can be part of the scientific, as opposed to decision taking, process. This paper argues that this is irrational. There is also a tendency for decision takers to seek to offload the responsibility for the decision onto the scientists. That too is irrational and arguably illegal; science is only part of the decision and the scientists have no democratic mandate.

The ALARP principle is offered as a formalism which goes some way towards a way of deciding which risks to accept. It is not perfect, both because it still requires subjective judgement to implement and because it is difficult to translate into a legal culture based on the civil law. It does however provide a framework within which the decision taker can structure his thinking and ensure that, even when it is necessary to be subjective, the lines of argument are clear and may be reproduced and reviewed, for example if new information or requirements were to emerge.

Way ahead

Proper governance of science and technology will accommodate:

- uncertain science; and
- technology that can never be wholly reliable

within an overall framework of public policy which is shared across Europe. At a minimum, “Europe” must imply the Member States of the European Union, but decisions taken in one country have risk implications for other countries, even outside the EU, and a shared approach is essential if risks are to be properly managed.

This Workshop provides an opportunity to explore the governance of science and technology in many different cultures and subject to many different legal constraints.