HSE
Health & Safety
Executive

THE TOLERABILITY OF RISK FROM NUCLEAR POWER STATIONS

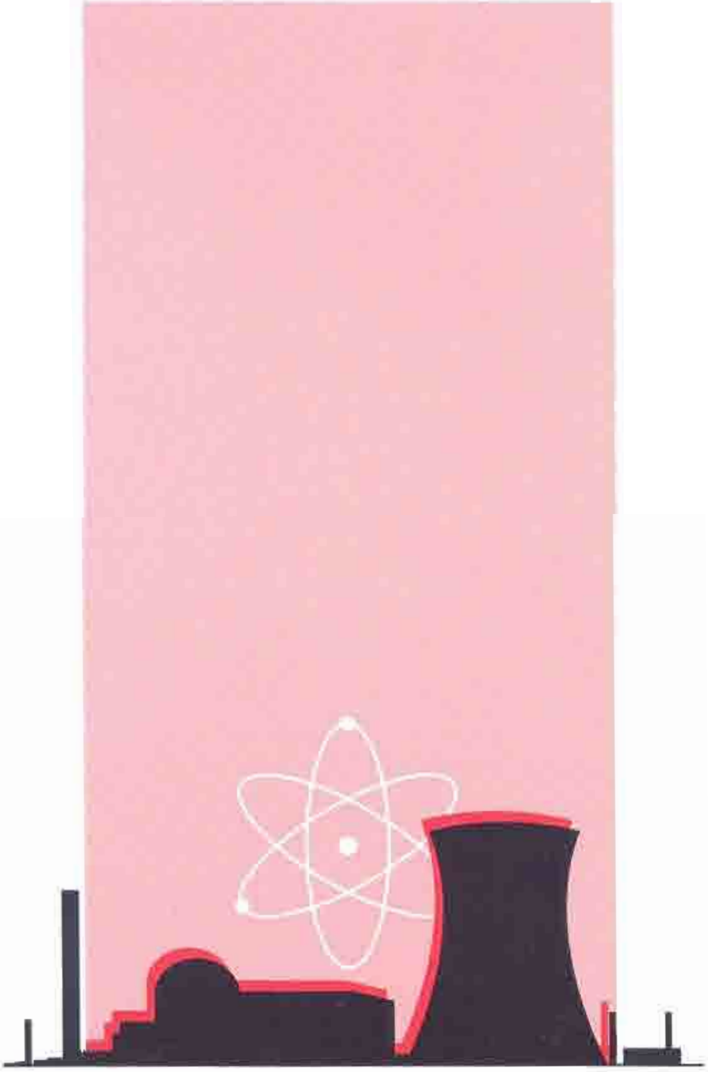

Acknowledgement

The Health and Safety Executive (HSE) wishes to acknowledge the advice and assistance received in the preparation of this document from:

Dr Geoff Ballard	Chief Executive, Safety and Reliability Directorate
Dr Donald Broadbent	Department of Experimental Psychology, University of Oxford
Dr Roger Clarke	Director of the National Radiological Protection Board (NRPB)
Mr H J Dunster CB	Former Deputy Director General of HSE and former Director of NRPB
Professor Bev Littlewood	Director of the Centre for Software Reliability, City University, London
Dr David Slater	Chief Inspector, HM Inspectorate of Pollution
and	
Mr Brian Street	Non executive Chairman of Air Products plc.

The document was produced by a working group under the chairmanship of Mr J D Rimington CB, Director General of HSE and including, besides the above, Dr S A Harbison, Head of Nuclear Safety Division and Chief Inspector of Nuclear Installations, HSE, and other HSE officials.

The tolerability of risk from nuclear power stations

Contents

Appendices

London: HMSO

INTRODUCTION

1 This document originates in a recommendation by Sir Frank Layfield, in his report of the Sizewell B Public Inquiry (1986) that the Health and Safety Executive (HSE) should 'formulate and publish guidelines on the tolerable levels of individual and social risk to workers and the public from nuclear power stations.' He said 'the opinion of the public should underlie the evaluation of risk; there is at present insufficient public information to allow understanding of the basis for the regulation of nuclear safety.' He recommended 'that as a first step, HSE should publish a document on the basis of which public, expert and Parliamentary opinion could be expressed.'

2 This attitude to industrial risk, and in particular the use of the word 'tolerable' was new. Until that time, discussion of the extent of the risks, to people or to the environment from industrial undertakings tended to be regarded as a matter for experts. Many experts believed, and some still believe, that the quantification of risks is too uncertain and too difficult a matter for people generally to grasp. Others openly said, and some still do, that most ordinary people only wish to believe that there is no risk at all from such undertakings, are probably not interested in finding out what exactly the risks are, and that to carry on a public discussion on the subject simply provokes objections to things that people are not in reality bothered about.

3 But in fact many people **are** bothered about nuclear power and other industrial risks and have become more so during the years since Sir Frank Layfield wrote his report'. The Health and Safety Commission and Executive believe that the public ought to have available to them an accurate account of important industrial risks and how they are controlled. In response therefore to Sir Frank Layfield's recommendation, HSE produced the first version of this document (the Tolerability document) in February 1988, comparing the extent of the nuclear risks with other risks and stating what measures and standards are applied to their reduction.

4 The Commission invited public comment on the document, and published all the comments received in November 1988, stating their intention to review the document and possibly to republish it in some updated form, taking account of comment.

5 Subsequently, the document received a close public examination at the Inquiry by Mr Michael Barnes QC into the application to build a new nuclear power station at Hinkley Point 2. As a result the Inquiry Report made a number of observations as to the standards to be applied to a station to be built at Hinkley. By and large, the report accepted the document as a clear and accurate exposition of the risks, and of the standards applied by HSE through its Nuclear Installations Inspectorate (NII).

6 Sir Frank Layfield had also recommended in 1986 that a review should be carried out of the Safety Assessment Principle (SAPs) applied by the NII in reviewing designs or modifications of nuclear power stations. This has since been done. Obviously, the detailed principles applied in the revised SAPs need both to reflect and follow on from the account of the standards described in the Tolerability document, which themselves describe the overall safety aims pursued by HSE. So the Health and Safety Commission have decided that the two documents should be republished in a single series, taking account both of public comment and of the conclusions of the two major inquiries.

7 This does not mean that the SAPs can never be adapted to new circumstances. On the contrary, it will be necessary to develop the detailed approach so far adopted by NII in relation to any new designs proposed for nuclear power stations in the UK having regard to the approach adopted in other countries, so far as seems sensible and safe. It seems right however, at a time when no such proposals are in immediate prospect, to bring up to date the thinking that has applied so far, and to republish, in the Tolerability document, the overall approach and standards which the Health and Safety Commission and HSE intend should apply, whatever changes may in future be made in the approach described in the SAPs.

8 In the meantime the approach and philosophy described in the original Tolerability document is increasingly applied to the regulation of other major industrial risks in the UK, and has attracted world-wide attention. That does not mean

that precisely the same limits are applied to all risks as to the nuclear risk. As the document explains, hazards and risks differ from industry to industry and some are arguably more readily acceptable, and the associated risks more tolerable, than others. But the need to be clear what the risks are, how they compare with other risks and hazards and within what limits they should be regulated, is invariable.

9 So this document, like its predecessor, is a straightforward account written for the general public. It begins by discussing how people normally approach risk. It shows how industrial risks, and in particular nuclear risks, are regulated in the UK. It explains the nature of the risk from radiation and how risk is calculated when deciding to licence nuclear installations. It is written in a spirit that final judgements about whether a given risk is tolerable are not matters for experts alone, but for the people who have to bear the risks, and who are therefore entitled to be given the best possible technical advice about them.

RISK AND TOLERABILITY OF RISK

10 'Tolerability' does not mean 'acceptability'. It refers to a willingness to live with a risk so as to secure certain benefits and in the confidence that it is being properly controlled. To **tolerate** a risk means that we do not regard it as negligible or something we might ignore, but rather as something we need to keep under review and reduce still further if and as we can. For a risk to be 'acceptable' on the other hand means that for purposes of life or work, we are prepared to take it pretty well as it is.

Risk: some basic considerations

11 Risk is the **chance** that **something adverse** will happen. To take a risk is deliberately to incur that chance; and estimating a risk involves defining that something precisely and finding a way of calculating how often it is likely to happen in particular circumstances. A more precise definition is at Appendix 1.

12 When we use the word 'chance' we simply mean the probability of something happening, as for example of a horse we fancied winning a race. 'Risk' is the chance and the consequences taken together. In common speech we quite often use the word 'risk' as referring more to the consequence than to the probability. So, for example, if we see a steeplechaser going for a high fence we may say 'that's risky' meaning not so much that there is a high probability of it being struck, but that it will have nasty consequences if this does happen. And indeed, there are three components to be considered in estimating any risk - the probability (eg, whether there is a 'high risk' or not), the event to which probability attaches, and the **severity** of the consequences. All of these components are implicit in any discussion of risk but need to be made quite explicit when considering societal risk (see paragraph 61 and beyond).

13 Whenever we do something that involves taking a risk - even stepping off a pavement when there is traffic - we usually do so because we believe we can see some benefit that outweighs it. We are likely to take the consequence for granted, to estimate the risk, however instinctively, and then see if we can reduce or avoid it. These simple principles apply also to the much more wide-ranging risks this publication discusses.

14 Besides the risks each of us willingly takes to secure benefits we want, we face a degree of risk from naturally occurring hazards. There is, for example, a chance of one in ten million* each year that any one of us will be killed by lightning. Because lightning generally kills only one person at a time, and the risk to each of us is very low, we treat it as negligible; ie apart from taking certain simple precautions the possibility of dying in this way does not influence our behaviour.

15 We all know three other important things about risk. First that there is no such thing as 'nil risk'. However we occupy our time, even if we are at home, we are taking some kinds of risk (eg a house fire or risks associated with DIY); we would be taking other kinds if we did something else. Second, we know that however remote a risk may be, it could just turn up. Remote risk is not the same as no risk at all. And third, each of us knows that our own chances may be either more or less than the average, depending on where we live,

* Written as 1 in 10^7, ie a 1 with seven zeros after it: 1 in 10 000 000.

whether we are more nimble, or younger, or have better sight, and so on.

16 Neither in the case of the risks we take voluntarily nor of naturally occurring hazards, are we usually deterred by statistics; to learn that over 5000 people are killed each year by traffic does not prevent us from using the roads, though it warns us to be cautious. A woman who wants a child will not change her mind if she learns that the average chance of her dying as a result are of the order of 1 in 10 000 (1 in 10^4). We can judge such chances by experience. And apart from anything else each of us is able to decide whether the benefit he or she has in mind is worth the risk.

Societally regulated risks

17 This paper does not deal with the risks that we consider and take as individuals, or with the naturally occurring risks. It concentrates upon certain kinds of risk that are regulated by society as a whole, with the aim of securing general benefits. Of course, society will choose to regulate 'voluntary' or 'natural' risks if they affect enough people. Thus roads are designed and laws applied to reduce traffic accidents; the risk of a freak tide flooding Central London has been reduced by the Thames Barrier. But quite apart from such events there is an increasing number of major industrial risks whose implications are so far reaching that they have to be brought under common scrutiny even though the immediate benefits may include a better life for all. Since individuals do not bear these risks of their own free will, it is important that we should all be satisfied that they are being properly controlled.

18 When risks are regulated by society, the relevant judgements cease to be in the hands of the individuals who bear the risk. The risks will be shifted around, so that some people bear more and others less of them; and the benefits may also be unevenly distributed. For instance the building of any dam imposes risk on people nearby whereas the benefits are shared by people living further aways. Societal risk may be redistributed in many other ways : for example through time, so that less risk is borne now, but more by some future generation. Or one kind of risk may be substituted for another.

19 People tend to view risk differently according to whether they can judge the hazard directly from experience, or whether the cause of the danger is not well understood or is particularly dreaded; or perhaps whether it could result in large adverse consequences from which individuals could not escape. Thus public expectations aboutthe levels of protection required, or the level of risk which can be tolerated, may well differ according to the nature of the hazard in question and people's knowledge or feelings about it. It has been suggested for example that people seem readier to tolerate the idea of sudden death by electricity in the home (40 deaths a year) than they are by the thought of some more insidious hazard, such as poisoning. Several theories have been advanced along these lines to explain the fear that many people have of radiation: the fact that it can harm without being felt; that it is capable of causing cancer; that it can harm unborn children; or simply that people do not understand and therefore dread it more than other risks.

20 There is nothing unusual about disliking one kind of hazard more than another. In the nature of things, people have their own views and feelings in these matters. There may be disagreement about the importance or incidence of any benefits. Some people may have ethical objections to particular activities or forms of harm and may in any case doubt what experts say about them. For all these reasons, the question whether deliberately to undertake major risks and regulate them societally can only be resolved in the way we settle all other matters that involve redistribution of big benefits and costs - by public discussion and through our representative institutions and legal processes.

21 When considering these matters we may try to weigh one risk against another. Appendix 2 gives comparisons between different kinds of risks that we commonly face, and says something about

---

* Actually about 1 in 13 000. In risk estimation the term 'of the order of' implies that the real figure may differ somewhat on either side of the figure mentioned.

The worst post-war accident in Western Europe was the overtopping of a dam in northern Italy which killed more than 2000 people in 1963. See Appendix 2, Table B4.

what such figures mean. When making such comparisons, we should remember that an identical risk (say a 1 in 1000 chance of dying this year) from two different hazards does not mean the same thing to everyone, since people's aversion to the two hazards may differ, and also the degree of choice we have in the matter or how we rate our personal chances.

22 From Appendix 2 we can distinguish the general levels of risk that individuals accept for a personal benefit (such as pay at work) and we can also see the level of risk we usually ignore or regard as negligible. These levels are broadly described in Table1.

Table 1

Levels of fatal risk
(average figures, approximated)

per annum	
1 in 100	risk of death from five hours of solo rock climbing every weekend
1 in 1000	risk of death due to work in **high risk groups** within relatively risky industries such as mining
1 in 10 000	general risk of death in a traffic accident
1 in 100 000	risk of death in an accident at work in the very safest parts of industry
1 in 1 million	general risk of death in a fire or explosion from gas at home
1 in 10 million	risk of death by lightning

To set these numbers into proportion, it is useful to bear in mind that one million is about the number of grains in one pound of sugar, so 1 in a million is the chance of picking out a specific grain. We also have to remember that the risks in question are risks of **being killed;** and that for each death there are additional injuries or sickness sometimes resulting in people dying sooner. When we compare with these figures the risks from radiation, which we shall be discussing later, we shall have to remember that the latter are chiefly risks not of being suddenly killed, but of having a shorter life. For the sake of comparison we have to reckon all the recordable risks of death as though they were the same thing, but to remember that this produces some bias against such risks as those from radiation.

23 To the extent that we give remote risks any thought at all, we do so knowing that each of us will ultimately die from some cause or other and that it could happen this year or next in any case. In fact, on average in Britain a man of 20 has roughly a 1 in 1100 chance of dying within a year; for a man of 40 the chance is around 1 in 600. At 60 it is 1 in 65 for a man; for a woman, 1 in 110. Each particular risk or cause of death is just one contributor to the overall risk we run.

24 All the same, we may well feel cautious about activities by others which add involuntary risks to the ones we cannot help or which we take for ourselves. People do in fact demand that such 'extra' risks be reduced to very low levels indeed compared to the ones they accept for themselves. In tolerating them, they may well want to know how the control of danger is achieved and how management in potentially hazardous industries works to reduce such risks. The purpose of the next section is to describe how industrial risks are regulated in the United Kingdom. We shall then consider how they are estimated, assessed and reduced.

THE REGULATION OF INDUSTRIAL RISK

25 The main tests that are applied in regulating industrial risks are very similar to those we apply in day to day life. They involve determining:

(a) whether a given risk is so great or the outcome so unacceptable that it must be refused altogether; or

(b) whether the risk is, or has been made, so small that no further precaution is necessary; or

(c) if a risk falls between these two states, that it has been reduced to the lowest level practicable, bearing in mind the benefits flowing from its acceptance and taking into account the costs of any further reduction. The injunction laid down in safety law is that any risk must be reduced so far as reasonably practicable, or to a level which is 'as low as reasonably practicable' (ALARP principle).

26 In some cases the ALARP principle is applied by a rapid judgement. But in the case of major risks it is often necessary to apply a much more formal process of assessment. In the case of the nuclear risk, the relevant working methods are quite complicated. This paper will turn to these in due course; but in the end, it is always a question of applying one of the three main tests outlined in paragraph 25.

The regulatory bodies

27 In the United Kingdom it is the business of the regulatory bodies, particularly the Health and Safety Commission (HSC) and the Health and Safety Executive (HSE) to make clear the principles they propose to follow in assessing or reducing industrial risks and to see that the relevant precautions are taken.

28 **HSC** is the statutory body responsible for the administration of the Health and Safety at Work etc Act 1974 (HSW Act) and its subsidiary legislation which includes the Nuclear Installations Acts. It is composed of nine members representing employers, employees, local authorities, and the public, and a chairman appointed by the Secretary of State for Employment. In nuclear matters the Commission is advised by the Advisory Committee on the Safety of Nuclear Installations (ACSNI), whose reports are published. ACSNI also advises the Secretaries of State for Trade & Industry and Scotland who are answerable in Parliament for civil nuclear safety matters.

29 **HSE** is a corporate body of three people appointed by HSC, and has some 4500 employees, mainly inspectors and technical, scientific and medical experts. Its Management Board includes the Chief Inspectors of the various Inspectorates concerned with the enforcement of industrial safety and health. The Executive is the licensing authority for nuclear installations; this function is delegated to the Chief Inspector of Nuclear Installations. Organisation charts of HSC and HSE are seen in Figures 1 and 2.

30 The regulatory system thus provides for the

Figure 1

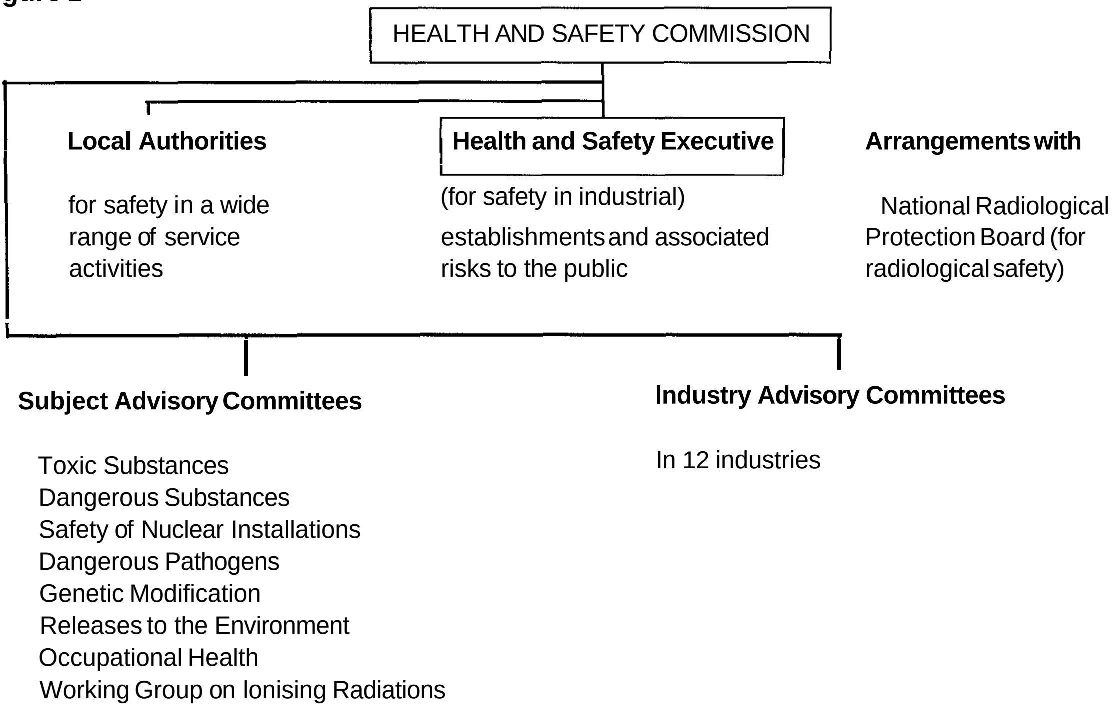

Figure 2 HSE Organisation

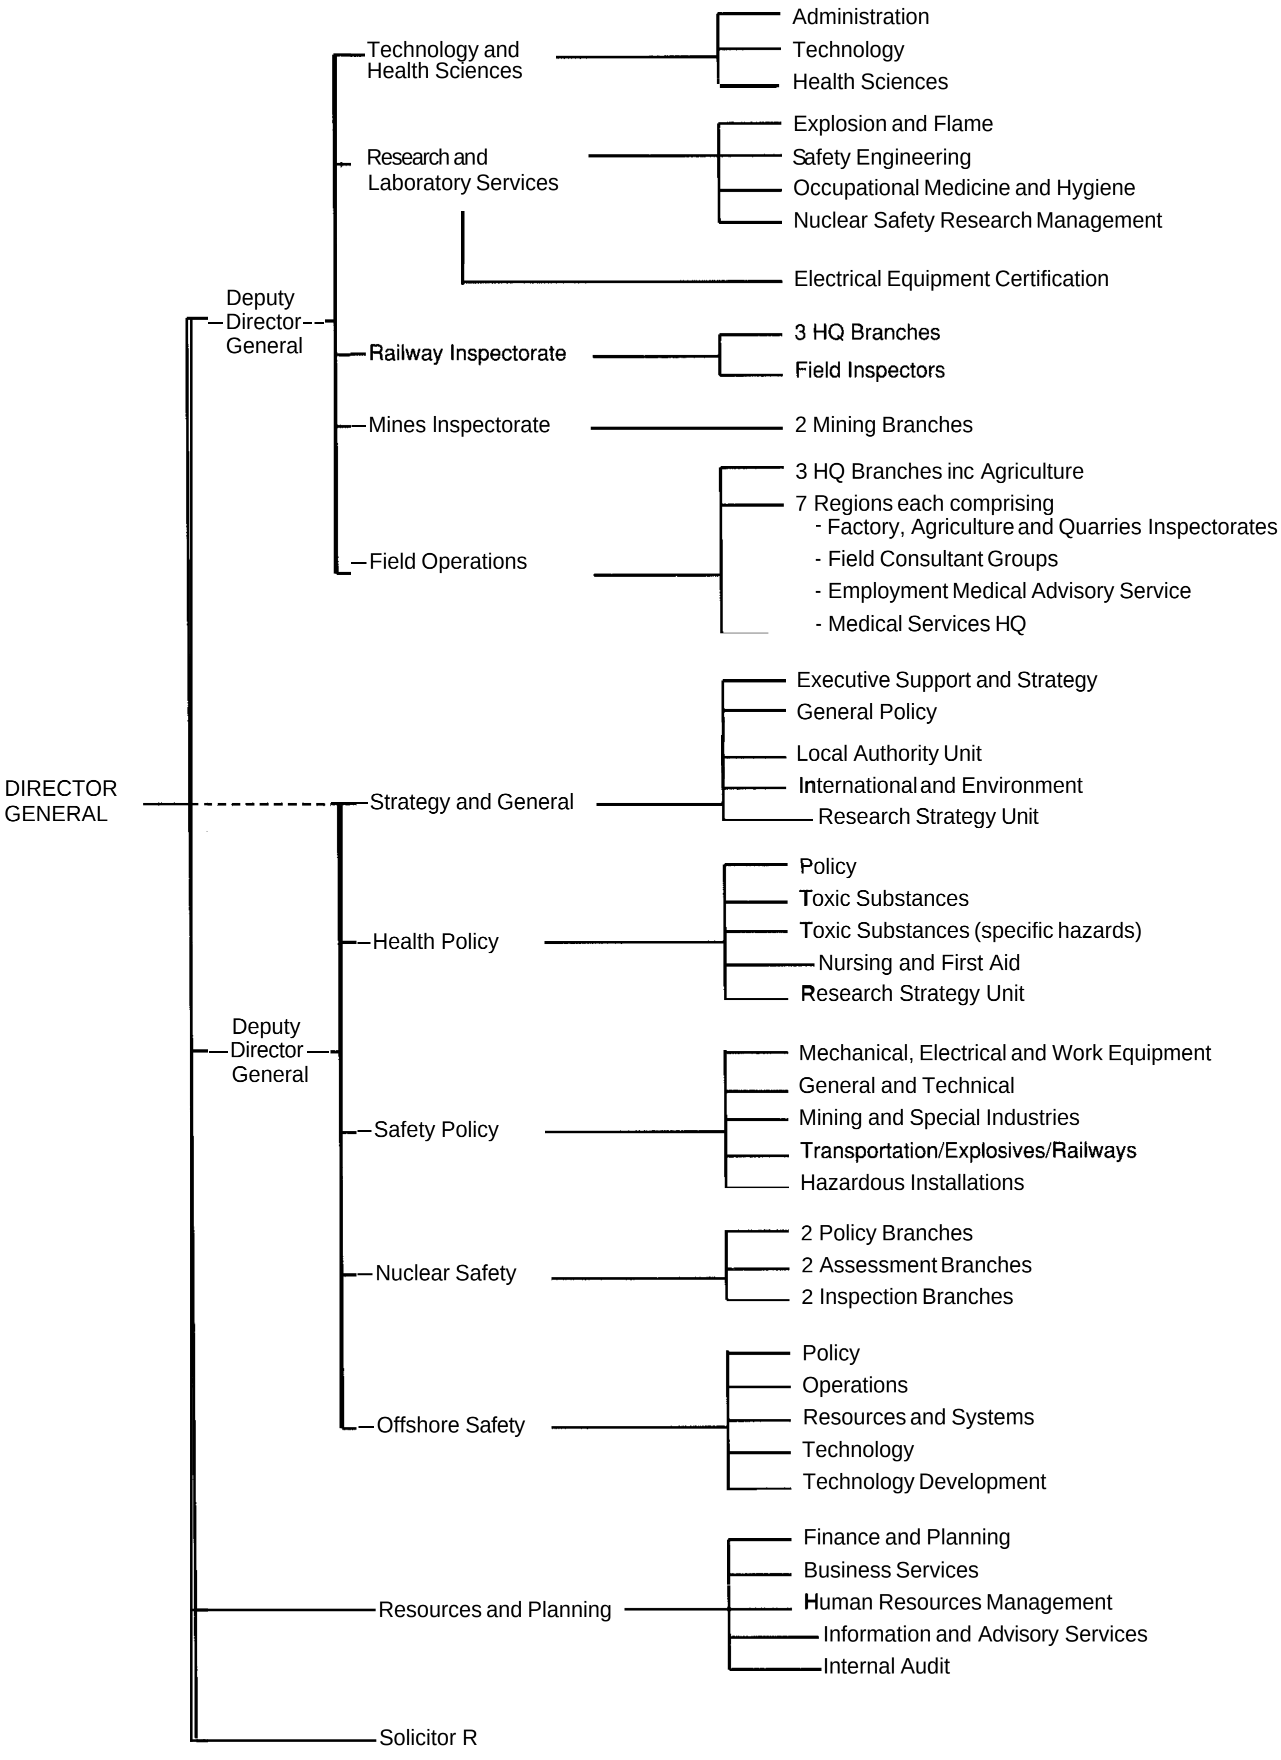

effective independence of the two main regulatory bodies from government departments. HSC is itself partly nominated by industry, but is debarred by statute from instructing HSE in its activity as an enforcing or a licensing body in any particular case. Unless directed in a specific matter by a Secretary of State, HSE is therefore able to take its decisions on whether to license or maintain the licence of a nuclear installation without interference from departments or from the power-generating or nuclear fuel industries, and in the light only of its view of the prospective or actual safety of the plant.

31 In the case of risks, such as the nuclear risk, where an important part of the harm can arise from the waste products which need to be disposed of, a very important role is played also by the Environmental (or Authorising) Departments (the Department of the Environment, the Ministry of Agriculture, Fisheries and Food, and the Offices for Scotland and Wales). Their inspectors visit nuclear installations to check on compliance with certificates of authorisation issued under the Radioactive Substances Act (1960) for disposal of wastes, to sample effluents and to monitor the environment and those parts of any plant which produce wastes. They also investigate relevant incidents and assess new plants, or help to do so, and they manage much relevant research.

General principles of regulation

32 In the UK, the legal responsibility for the safety of workers and the public is placed on whoever controls work activity - usually an employer. We can call this person an 'operator'. As we have seen, the primary principle which the HSW Act lays down and HSE enforces is that he must do whatever is **reasonably practicable** to reduce risk (ALARP, see paragraphs 25 and 26 above). Legally speaking, this means that unless the expense undertaken is in gross disproportion to the risk, the employer must undertake that expense. This principle also means that employers are entitled to take into account how much it is going to cost them to take a safety precaution, and that there is some point beyond which the regulator (HSE, with the Nuclear Installations Inspectorate (NII) acting for it on most nuclear matters) should not press them to go; but that they must err on the side of safety. In the case of risks where the public safety is seriously involved, employers will be pressed very hard indeed to say why they should not spend money to decrease even a low risk.

33 In applying these principles HSE as the enforcing body issues a large number of standards or codes and a great deal of guidance stating what in its view is 'reasonably practicable' for a very large number of industrial activities. The more important guidance documents are approved or published by HSC; and any of them is likely to be produced as evidence in legal proceedings against any employer who is believed by HSE not to have done all that is reasonably practicable to reduce risk.

34 For certain hazards, including cancer-producing materials such as radioactive materials and asbestos, and toxic substances such as lead, the safety regime is very stringent indeed. For such hazards, the approach is:

(a) to fix a level of personal exposure that can be regarded as just tolerable, but must not be exceeded; and

(b) **then** to say that each employer must do better by reducing the exposure and so the risk to the lowest level that is reasonably practicable.

The fixed standard may have to take account of what can be accurately measured and therefore enforced, but the level of safety which it guarantees may be improvable. What can then be done to improve performance is partly a matter of judgement in each case and partly a matter of ensuring that employers are following the best industrial practice; but they are legally required to do this, not just to stick at the level which is regarded as just tolerable.

35 Figure 3 roughly illustrates these requirements. Above a certain level, a risk is regarded as intolerable and cannot be justified in any ordinary circumstances. Below such levels, an activity is allowed to take place provided that the associated risks have been made as low as reasonably practicable. In pursuing any further safety improvements to demonstrate ALARP account can be taken of cost. It is in principle possible to apply formal cost-benefit techniques to

Figure 3 Levels of risk and ALARP

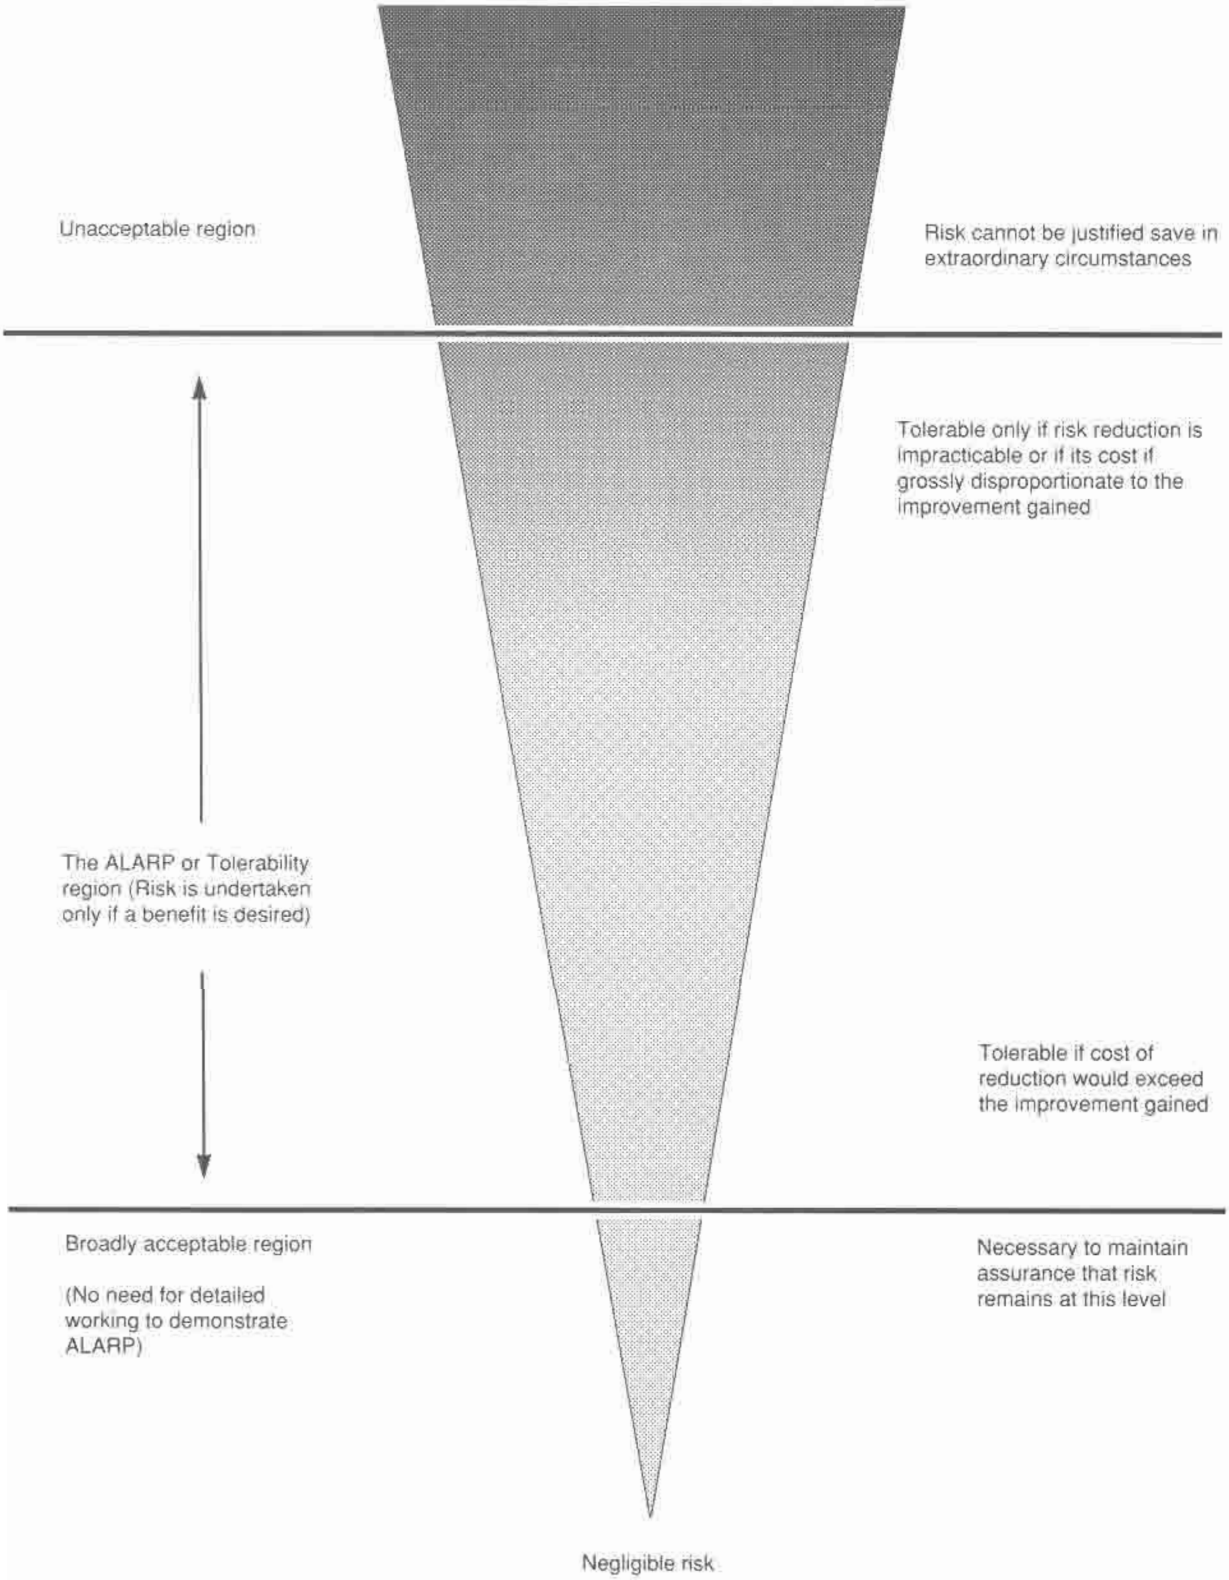

Figure 4 Levels of risk and ALARP

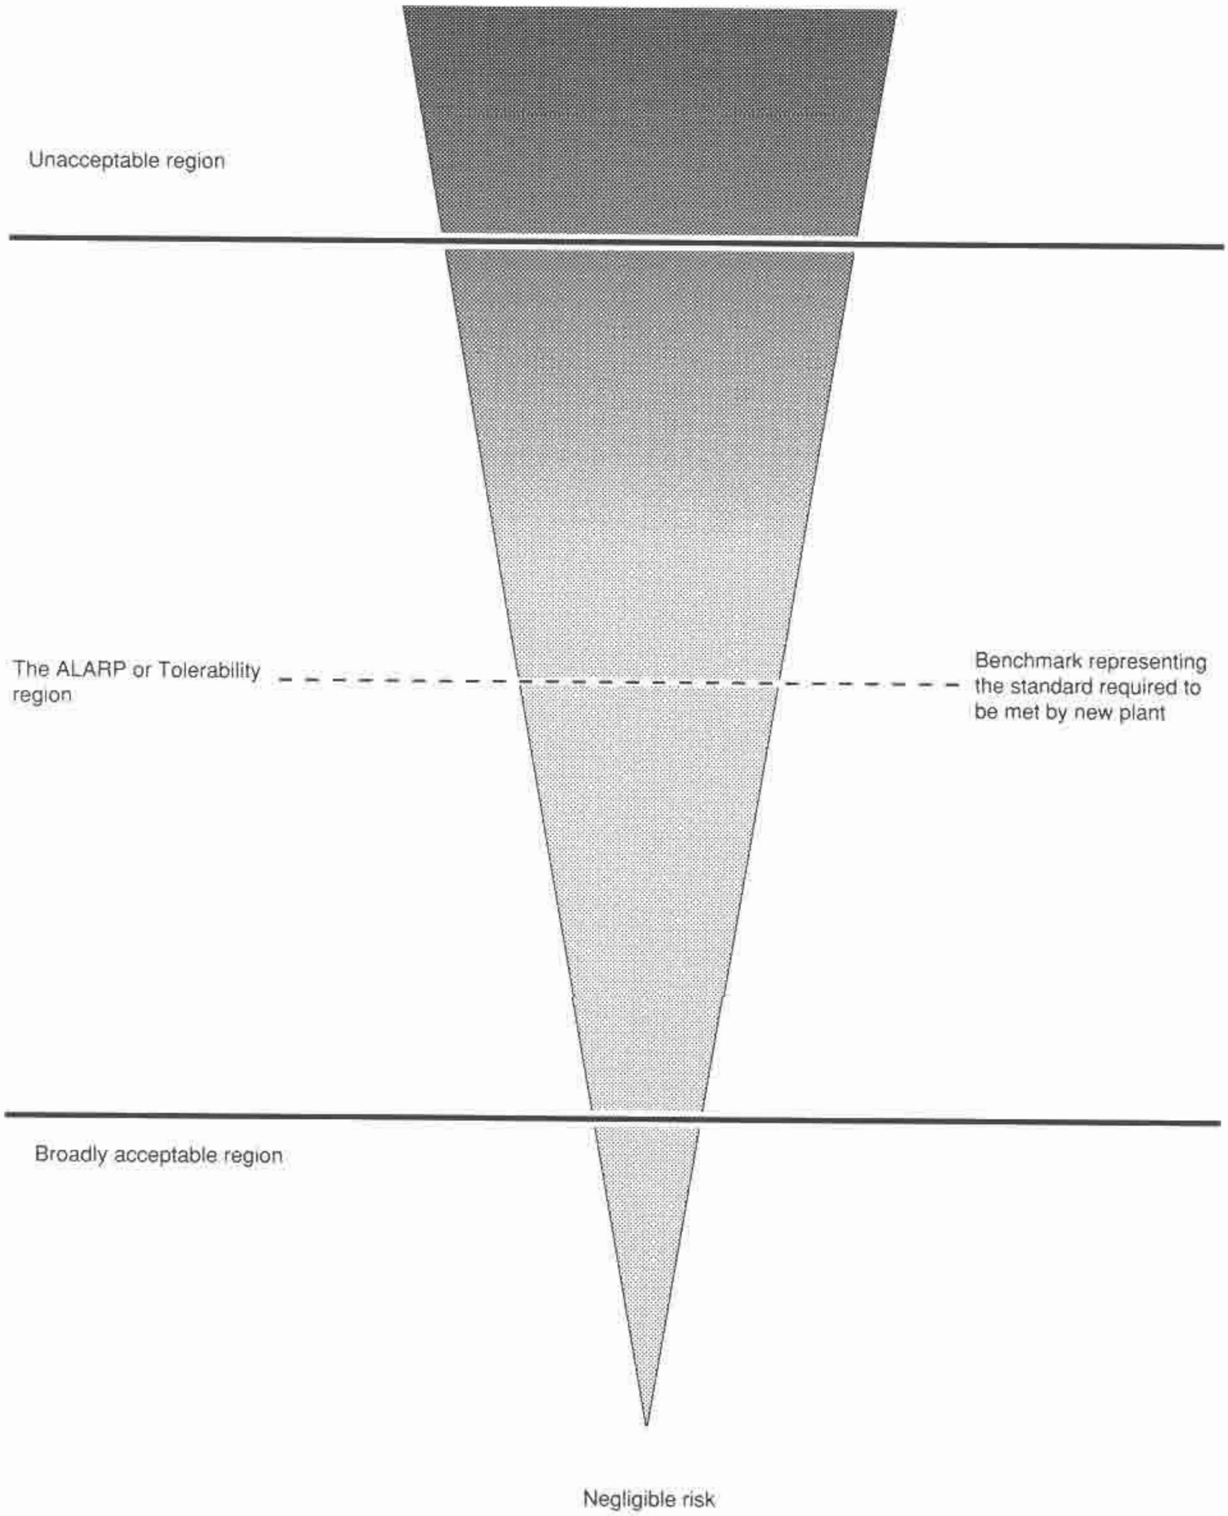

assist in making judgements of this kind. The practical issues are considered in Appendix 3.

36 In weighing the costs of extra safety measures the principle of reasonable practicability (ALARP) applies in such a way that the higher or more unacceptable a risk is, the more, proportionately, employers are expected to spend to reduce it. At the point just below the limit of tolerability (see Figure 3) they are, in fact, expected to spend up to the point where further expenditure would be grossly disproportionate to the risk; and as the risk will by definition be substantial, that implies a considerable effort even to achieve a marginal reduction. There may however come a point where in the existing state of technology even a marginal further reduction would be unjustifiably expensive and at that point the obligation to do better if reasonably practicable is discharged. These points are further pursued at Appendix 3.

37 Where the risks are less significant, the less, proportionately, it is worth spending to reduce them and at the lower end of the zone it may not be worth spending anything at all. Below this region the levels of risk are so insignificant that they need not claim attention, and the regulator need not ask employers to seek further improvement provided that they are satisfied that these low levels of risk will be attained in practice. Nevertheless employers might decide to spend even more, and some do.

38 The lower limit of the 'broadly acceptable' region in Figure 3 is set by the point at which the risk becomes truly negligible in comparison with other risks that the individual or society runs.

39 In practice, as technology improves it becomes possible to be confident that any new plant can be designed to meet a level of risk that is well below the level that would be regarded as intolerable by comparison with other risks, though it may still be in the region where further effort should be made if this is reasonably practicable. In this case it can be right to establish a new control level or 'benchmark' which new but not necessarily old plant, should generally meet or exceed. One consequence of Barnes' review at Hinkley Point[2] (paragraph 5) is to establish such a new control level for this risk from new nuclear plant. Figure 4 shows how this works. In some forms of risk control and regulation, a benchmark will take the form not of a risk level but of a written standard or description of design targets which represents an improvement on its predecessor. These are the ways whereby a regulator seeks to reduce risk levels as technology allows, always bearing in mind what is reasonably practicable.

40 In arriving at these standards, the regulatory authorities rely partly on their own experience and judgement, partly on international discussions and agreement, and partly on the best independent industrial, expert and scientific advice offered through a number of advisory committees. A considerable part is played by consultation with representatives of all parties including the people at risk. In this way, the best presently attainable regulatory standards are determined.

41 Such procedures accord with common sense. They recognise that not everything can be measured to the greatest degree of accuracy. They implicitly say that where a risk can be quantified, a definite standard of performance should be fixed. They recognise that measurement is not always possible, and that therefore expert assessment or judgement is called for. They compel employers to be always on the look-out to do better. In the case of the more serious risks they state flatly and definitely that there is a basic standard of achievement or a particular requirement that must be met at all costs, even if to do so would put the employer out of business. But once these mandatory requirements have been met they apply the idea of 'reasonableness' and do not insist on the impracticable or upon unnecessary expense.

Licensing of nuclear installations

42 In the case of nuclear installations, the main tool for applying these principles is the **site licence** which is granted only when NII is satisfied with the design safety of the plant. The conditions which are attached to each licence enable NII thereafter to control progress in meeting safety requirements at each stage of construction, commissioning and operation. The conditions which can be attached include:

(a) **consents,** which mean that the licensee cannot carry out a particular operation until NII has agreed it;

(b) **approvals,** whereby the licensee is required to describe the general arrangements for carrying out particular activities, and obtain NII's approval; and

(c) power to give **directions** whereby the licensee is directed to take any action which NII considers necessary.

43 The licensee remains ultimately responsible for the safety of the plant, except in so far as a direction might have been given. Licensees initiate and apply all the relevant safety standards and procedures and are responsible for ensuring that risks are reduced according to the ALARP principle. But the powers available to NII mean that there is a very powerful check on their actions and that if they wish to keep their licence they must carry out their duties. They must also produce a safety case for their plant which satisfies the NII assessors. HSE has published guidance for the assessors in the form of Safety Assessment Principles[3,4,5] which apply in detail the general theory of tolerability set out in this document, as well as setting down the NII's judgement of what constitutes the best modern practices in engineering and science.

44 The control exerted by NII upon a nuclear installation is applied prior to and during design; during construction; during testing and commissioning; at start up and during operation and dec'ommissioning. A licence is required before construction can commence; and thereafter, important steps in the construction, commissioning and operation of the plant are controlled by consents and approvals which are only granted after NII, in consultation with the Authorising Departments, is satisfied that the safety case is adequate.

45 Once the installation is operating its safety performance is monitored by the licensee's health and safety staff, and it is subject to intensive audits and inspections by the NII. All incidents of a significant character and any involving an unauthorised release of radioactivity however small have to be logged and reported and where necessary investigated. They are also announced quarterly by HSE.

46 There is usually one designated NII inspector for each major installation. For smaller installations (research reactors, fuel stores etc), several are grouped together under one inspector. Generally, site inspectors aim to spend at least 25% of their time on site, in visits of two to three days duration, but they also take part in special inspections, for example quality assurance audits to ensure that the relevant specifications are met, reactor start-up inspections after each biennial shut-down; and incident investigations. Some 30% of NII staff are engaged on site inspection duties; most of the rest are engaged in assessing the safety of new plants and reviewing the safety of existing plants as they grow older.

47 At power stations, each reactor is shut down periodically for inspection and maintenance and may not be started up again without NII's consent. If the licensee wishes to carry out any significant modification this will be assessed by NII in the same way as the original design. The licensee is required to subject every aspect of their reactor to scrutiny and investigation; all of their analyses are available to NII.

48 In addition to the day to day surveillance of safety NII requires the licensee to carry out periodical safety reviews if they wish to operate beyond an agreed period, relating to the original design calculations. The reviews have had three principal objectives: to demonstrate that the plants are operating as safely as originally intended; to identify any life-limiting features; and to compare their safety against modern standards. From these reviews, appropriate safety modifications are identified and implemented, after discussions between licensee and NII. Continued operation has to be justified to the satisfaction of the NII and the conclusions of NII's review process are made public.

49 The first two objectives are reasonably straightforward. The comparison with modern standards is more difficult. Safety standards change as the years go by partly because of improvements in methods, materials and technology but also, most importantly, because people's expectations increase. It is not always possible to demonstrate whether or not a reactor can meet a particular modern standard; in other cases it may be clear that the reactor cannot do that. In such cases the licensee has to justify further operation by particular arguments and to

make whatever improvements are reasonably practicable, taking into account the length of the period of further operation they wish to achieve. This may entail restrictions to operating limits such as altering temperatures or pressure, to increase the margins of safety. For any particular plant a stage will eventually be reached when the licensee will either be required to close their plant or will do so themselves because they are unwilling to incur the expense of improvements in safety needed for further operation.

50 The standards and techniques which are applied in nuclear regulation in the UK have developed partly as a result of extensive international discussion and review, within the International Atomic Energy Agency, the Nuclear Energy Agency of the Organisation for Economic Co-operation and Development (OECD) and elsewhere. NII participate regularly in these discussions and in specific bilateral discussions with regulatory authorities in a number of other countries. By these means the widest available experience is brought to bear on British regulatory practice, to ensure that the UK's reactors are operated as safely as reasonably practicable and to the highest standards agreed internationally.

BROAD PRINCIPLES OF RISK ASSESSMENT

51 In the case of major industrial hazards, which can be defined as 'any man-made industrial hazard which has the potential to cause large-scale injury and loss of life from a single brief event', it is nowadays the practice to insist that those who control the plant should carry out an assessment of the hazards it poses and submit this to HSE for examination. The Executive may then insist on additional precautions.

52 This practice has been applied for many years to risks from nuclear power stations; the way assessments are carried out in the nuclear case is described in the section on the risk of nuclear accidents.

53 In such assessments, a systematic approach has to be adopted. **First,** plant reliability and risk of plant failures have, so far as possible, to be measured or calculated in some way, so that the area left to judgement is circumscribed. Nevertheless the judgement of scientists and engineers will always be an indispensable feature of risk assessment processes, because it is not always possible to demonstrate how far an overall plant design or an individual item will perform in extreme circumstances nor is it possible to be sure that every contingency has been taken into account in the calculations. So it is always important to judge what reinforcement needs to be built in to cater for the unexpected. To protect against the consequences of the random failure of physical components, greater robustness can be specified for these components. Alternatively, identical or reinforcing back-up components might be provided (called **'redundancy').**

54 Redundancy will not, generally, provide protection against inherent design faults, which would simply repeat themselves in every affected component. One way of tackling this is to provide back-up via dissimilar components, ie ones that have been designed independently; this approach is called **'design diversity'.** For simple devices this can be very effective, since the versions would almost certainly fail independently of one another. When the task to be carried out is a complex one, however, it may only provide fairly modest reduction of risk. For example, experiments on diverse computer software suggest that there is a tendency for different designers, independently tackling the same problem, sometimes to make similar mistakes. If this happens, there will be a possibility for some of the faults to be present in all versions, thus creating a chance that these will fail simultaneously in certain circumstances.

55 **Second,** there is a rule of conservatism so that when figures are used they are generally on the cautious side. Finally, great attention is paid to two factors; the **quality of the plant itself** (including the management system for assuring it), at all stages of its design, construction, operation and maintenance; and the **operational procedures** which management apply.

Individual risk and societal risk

56 When putting a figure to a risk, it is always necessary to be clear as to whom or to what group of people the figure applies. It would for example be meaningless to say that the national average

risk of being killed while hang-gliding is one in twenty million. What is relevant is the risk to the people who actually go in for hang-gliding.

Individual risk

57 When we are considering the chance of a large scale industrial accident, the important question for each individual is, what is the risk to me and to my family? The way this is answered in risk estimation is to calculate the risk to **any individual** who lives within a particular distance from a plant; or who follows a particular pattern of life that might subject him or her to the consequences of an accident.

58 In making such an estimate the first question concerns how likely is an accident in the first place, and for this it is necessary to consider the likelihood of each important kind of plant failure. Data on the 'reliability' (ie failure rates) of plant and particular components are collected and computerised to assist in such studies. Account has to be taken of the reliability of human beings in the design, construction, testing, operation, maintenance, and modification of such plant and components. Then the results of particular failures, for example, how much toxic gas, or flammable substance, or radioactivity would be released, have to be considered. It is then necessary to calculate how such releases would be likely to affect a hypothetical person who was at some particular spot, taking account of the possibility that he or she might be in or out of doors; what dose would be received, and what harm would this do? Such calculations have also to take account of the behaviour of substances under different weather conditions. For example at coastal sites there is some probability that the wind would take the substance out to sea; or the substance might be in such a form that it would not be carried far, or might be deposited in larger concentrations in particular places. Population and sometimes transport patterns have to be considered. Finally, the chances of harm from all significant failure causes have to be summed up to give the overall level of risk from the installation.

59 Only when such complex studies and calculations have been made is it possible to predict the chance that any individual living within a particular radius, or behaving in a certain way, will be injured by an accident. Such calculations are referred to as 'individual risk' calculations, and enable us to say things like 'a person who lives within half a mile radius of such and such a plant has a chance of x per annum of being injured from an accident at the plant'.

60 According to the different forms of hazard for which calculations are made, there are various possible ways in which people might be injured. They might die more or less instantaneously from the effects of explosions, fire or toxic gas release. Radiation is a different case, since unless the dose is extremely large death within a few weeks or even years would be extremely unlikely. What could happen is that, depending on the size of the doses, a proportion of the people receiving them would develop a cancer some years after the exposure took place or possibly pass on a genetic abnormality to some of their descendants (a number of non-radioactive chemicals can cause the same sort of effects). So, by and large, the risk that an individual experiences from radiation exposure is to increase somewhat that person's chance later in life of developing a cancer, which might or might not be treatable.This needs to be borne in mind when comparing the effects of exposure to ionising radiation with the effects of conventional accidents, which are immediate.

Societal risk

61 The consequences of a major accident however go much wider than injury toparticular people, or to sets of people.

62 For example, in the case of the very serious nuclear accident at Chernobyl, not only were some people killed and others affected to a greater or lesser degree by fall-out, but there was also serious local disruption, and loss of plant and electricity production, as well as the fear that was generated both locally and internationally. Perhaps because each one of us has to die somehow, sometime, we tend not to take large numbers of individual accidental deaths so seriously as we do a single event killing a similarly large number of people. That is partly because these other consequences are frequently involved, so that a large scale accident raises questions of responsibility for safety and public accountability in a way that accidents to individuals do not.

63 To estimate such large events, we use the term 'societal risk', and in doing so, we have to consider not only the many different forms an accident could take, but the multiple consequences involved and their cost or severity. In principle at least, the different forms of harm that could follow from a very large accident could be added up and given a money value, provided of course that one is prepared to attach a value to the loss of human life. The things to be added would consist of:

(a) the price put on the loss of life adjusted for society's known extra aversion to multiple loss of life;

(b) the cost of coping with the emergency, the loss of plant destroyed or abandoned, of land rendered unproductive, and of opportunities for other investments foregone; and

(c) the costs, if they could be estimated, of shock and disruption to social and political life.

We can call these respectively the human, monetary and political **costs** of a major event.

64 Such a calculation would obviously beg many questions and require us to put values on sets of things that are difficult enough to estimate in themselves, and very difficult also to weigh with each other. But if we could do it, we would then have a sum or quantity that we could begin to weigh against the quantity of the **benefit** we were looking to receive from engaging in the investment - say in a nuclear programme - that could at an extreme involve society in these costs.

65 It seems likely that **no** benefit however large would cause us willingly to accept costs of the kind described above if they were certainly going to be incurred, unless the benefit took the form of avoiding some even greater misfortune. So in making this comparison we would of course have to take into account the **probability** of incurring the costs in question.

66 We can give the name **'detriment'** to the result of multiplying any cost by the probability of its being incurred. We would have to face many further complications in estimating such a quantity. An investment in a nuclear plant constructed, designed and operated to modern western standards is in fact very unlikely indeed to give rise to an accident on a very large scale. But it is rather more likely that it could give rise to some lesser but still quite disturbing accident and much more likely still that it might give rise to a minor accident that could have some health effects. Strictly the costs and the probabilities attaching to **each** of these orders of event would have to be calculated separately, and the detriments from each aggregated or set off in some way, in order to arrive at an estimate of **total detriment.**

67 **If** this could be done, a very useful and important step would have been taken in enabling society to take decisions about societally regulated risks (ie, the risks involving large numbers of people which Governments can choose on their behalf to take, avoid, or reduce: see paragraph 17). We would, not only, as stated, be able to weigh the total detriment from any risk against any perceived benefit, but we could weigh the net risks and benefits from doing one thing against the net risks and benefits from doing another, if they were alternatives. So, for example, we might weigh the total detriment from using fossil fuels against the total detriment from nuclear power. And even more than that, we could weigh the detriment from those activities that could give rise to **accidents** against the detriment from **continuing** forms of societally regulated harm - as for example arising from nitrate in water supplies - and so, both manage our social investments, and determine our priorities for reducing various risks or harms, better than we can at present.

68 Unfortunately, quite apart from the uncertainties already noted, we are far from having developed and agreed upon the accounting and valuing conventions that could enable us to perform such calculations. And so, in all estimates of probability of remote events, if we are to arrive at any judgement at all about societal risks, we have to find simpler means of weighing the factors involved. The most obvious way available is to confine ourselves to estimating what chance there is of a given number of people losing their lives from some 'typical' event, and comparing this roughly with other similar risks that are ordinarily accepted, or known well. We could at least say 'the chance of there being an accident at such and such a plant that would kill about y people is about x per annum'. The other consequences of such an

accident (such as economic cost) would then have to be fully borne in mind when assessing the seriousness of the event and comparing it either with other major risks or the benefits from running the risks in question.

69 HSE has already adopted such an approach in estimating risks for the purpose of giving its advice on land use planning in the vicinity of major chemical hazard sites. As Appendix 4 explains, in such cases individual risk calculations are more in terms of the chances of a hypothetical individual receiving a 'dangerous dose' of a chemical; and HSE's practice is to advise against homes being built in places where any individual's chance of receiving a dangerous dose was greater than 1 in 10^5 (1 in 100 000) per year. The chance of death from a dangerous dose would of course be lower than this.

70 Similarly, HSE will advise against major developments such as a housing estate being built where the individual risk of a dangerous dose is higher than 1 in 1 million per annum. In this case, the individual and the risk applying to them is being taken as a surrogate for the whole population and because of the difficulty of making 'societal risk' calculations, this is a procedure that must often be followed.

Summary on individual and societal risk

71 To summarise, there are two ways of estimating the overall risk from any hazardous industrial installation, namely by calculating:

(a) the risk to any particular individual, either a worker or a member of the public - this is **'individual risk'.** A member of the public can be defined either as anybody living at a defined radius from an establishment, or somebody following a particular pattern of life; or

(b) the risk to society as a whole - **'societal risk',** as represented, for example, by the chance of a large accident causing a defined number of deaths or injuries. More broadly, societal risk can be represented as a 'detriment', viz the product of the total amount of damage caused by a major accident and the probability of this happening during some defined period of time.

THE HARM FROM RADIATION

72 As is more fully explained in paragraphs 98 to 105, the main risk from nuclear installations concerns the potential for release of material which emits ionising radiations. There are some other man-made sources such as medical x-ray equipment, and there are natural sources. The following is a simplified account of what is known about the harm from radiation; for further accounts the reader should consult the further reading sections.

73 Ionising radiation is a naturally occurring phenomenon to which we are all exposed all the time: examples include the cosmic radiation entering the earth's atmosphere from space, and radiation from the natural radioactive materials in soils and rock. And some slightly radioactive materials, such as potassium, enter our bodies, usually via food or are present from birth. These are in addition to the man-made sources referred to above, and lead to the same kinds of harm.

74 All radioactive materials decay and eventually become stable, losing the harmful characteristics of their original radioactivity. However, the rate at which different forms of radioactive material decay varies enormously. So when radioactive materials enter the body, some of them will become harmless very quickly unless they are continuously replaced. Other materials may take so long to decay that they can remain active for periods much longer than the human lifetime. An important consideration for these long-lived materials is whether they become lodged in some of the body's organs, or are rapidly excreted by normal biological processes. The rate of decay of the material, the likelihood of its being excreted from the body, and the type and energy of the radiation that it emits are all taken into account when calculating how much radiation a person has received from radioactive materials in the body.

75 The amount of radiation received externally plus that from materials retained in the body is measured by a quantity called 'effective dose', often just called 'dose'. The effective dose is defined as the energy absorbed in body tissues, modified to allow for the different effect on the body of different kinds of radiation and for the different sensitivity of different organs of the body. The unit of effective dose is the sievert, but this is

inconveniently large for most purposes and effects on individuals are usually measured in millisieverts (mSv), ie one thousandths of a sievert.

Exposure to natural sources

76 Ionising radiations and radioactive materials can readily be detected and measured, even in small quantities, by sensitive instruments. The annual dose to people in this country from radioactivity in the ground, radioactivity in our bodies from birth or acquired since, and cosmic rays is about 1 mSv, with about a third coming from each of those sources. This value is an average. The doses to individuals may differ from the average by a factor of between 2 and 3, either higher or lower.

77 There is also an additional natural source, called radon, that is more variable. This is a radioactive gas that seeps out of the earth and can accumulate in our houses. The average annual dose in the United Kingdom from radon in houses is slightly more than 1 mSv making the average dose from all risks of natural radiation 2 mSv. In some areas, the dose is only about one third of this, but in others, such as parts of Cornwall, the annual dose from radon in homes is between 10 and 20 mSv. Still higher values occur locally and in some individual houses.

Exposure to man-made sources

78 Leaving aside the exposure of workers in certain occupations and of patients subject to radiation in diagnosis and treatment, the exposure of the public to ionising radiation from man-made sources is very slight.

79 The average annual dose to workers in nuclear power stations is about 1 mSv, with a few workers receiving doses greater than 5 mSv in a year. The limits now recommended by the International Commission on Radiological Protection (ICRP) are 20 mSv a year on average over a five year period with no more than 50 mSv in any one year. Generally the British nuclear industry works to a practical control level of 20 mSv in any year although a practical limit of 15 mSv per year is widely applied in the British nuclear industry. Some information about doses from medical treatment is given in table 3; this is the single highest man-made source of irradiation in our society.

80 A source which attracts a good deal of attention is radioactive waste. This currently gives rise to an annual dose, averaged over the whole population, of less than one thousandth of a millisievert (0.001 mSv), compared with more than 2 mSv from natural sources. Even the individuals receiving the highest doses from wastes are thought to get no more than 0.3 mSv per year.

81 Other man-made sources include the remains of the fall-out from early weapons tests, domestic smoke detectors, and traces of radioactive materials released by the burning of coal. Typically, the average annual dose per person from these miscellaneous sources amounts to about 0.01 mSv.

How radiation operates biologically

82 The body is made up of cells that are repeatedly replaced. Ionising radiation can damage these cells, mainly by causing errors in the DNA, which is the material that controls the function and replication of the cells. Some chemicals also damage the DNA. The natural process of replication is, itself, not free from errors and the cells have repair mechanisms that correct these various errors if they are not too severe.

83 The damage done to cells by ionising radiation, as with other causes, may prevent the survival or reproduction of the cell, but frequently the damage is successfully repaired by the cell itself. If that repair is not perfect, it may sometimes result in a cell that is modified, but can still reproduce itself.

84 **'Early effects' from high levels of radiation.** If enough cells in an organ or tissue are killed or prevented from reproducing there will be an observable damage to the organ. If the damage is slight, the injury will repair itself, but if the organ is vital and the damage bad enough the result will be death, perhaps in a matter of weeks.

85 **'Late effects' from lower levels of radiation.** A cell that is modified but not killed by radiation retains its ability to reproduce. The

resulting group of modified cells is usually recognised by the body as being 'foreign' and is eliminated by the body's defence mechanisms. But the body's defences can occasionally fail and the modified cells may, after long delay, develop into a cancer. If such a modification occurs in a germ cell whose function is to pass on hereditary information to future generations, effects could occur in the descendants of the exposed individual. (See paragraph 94 below.) All these effects occur long after the exposure and are called **'late effects'.**

86 The early effects of radiation will occur only if the dose is high and is received in a short time. For example, if the whole body is subjected to a dose greater than about 500 mSv over a few hours, blood changes will occur and the individual may feel feeble; if a person received 5000 mSv there is a high probability of dying in the following few weeks from severe depletion of the white blood cells as a result of the damage to bone marrow. A dose of 50 000 mSv is likely to cause severe gastrointestinal and central nervous system damage and death would follow quickly.

87 Tables 2 and 3 respectively give examples of some of the early effects of radiation, and examples of typical kinds of exposure in daily life.

Estimating the risks

88 Of the 'late' effects of radiation, the most important is cancer. Not all cancers are fatal, but since a high proportion are, it is usual to take the number of fatal cancers as the measure for the total harm, although the additional fact of the non-fatal cancers and possible hereditary defects has also to be considered. There is good practical evidence relating large radiation doses to cancer induction and a great deal of laboratory work that helps additionally to show that the same link can apply, though with a much lower chance, to the smaller doses that we all receive.

89 Conclusions as to the risk involved, that is to say how the dose received relates to the chance of getting cancer, are best obtained from studies on human populations. The more important studies include those on the survivors of the atomic bombs at Hiroshima and Nagasaki and on patients treated by doses of radiation.

The risk from low doses

90 Although there is little or no direct evidence of harmful effects at low doses, the assumption is made that there is no dose so low that it cannot cause harm (the no threshold assumption). It has been further assumed that, for doses well below the point at which 'early' effects become apparent, the chance of harm increases in direct proportion to the dose (the straight line dose-risk assumption). There are good grounds for believing that these assumptions are reasonable.

91 Studies on exposed **workers** have however so far yielded significant results only for groups with unusually high exposures, such as those who worked with radium in the early decades of the 20th century and those who inhaled radon and its daughters in uranium mining in the middle years of this century. More recent studies on other groups of workers with lower exposures, such as those in atomic energy installations in the US and the UK, are now beginning to provide useful information. At present, the precision of these new results is low and all that can be said now is that they are not inconsistent with other current estimates of the risk. However, the work is continuing and the precision will improve as the studies are extended to longer times and greater numbers of workers.

92 Much work has been done to link health effects to the variations in exposure to natural sources (see paragraph 77), but the results have been inconclusive. Other factors, such as industrial pollution and smoking, have much more influence and conceal any effects due to radiation.

93 Studies around nuclear reprocessing sites have sometimes shown slightly enhanced rates of cancer and small clusters of childhood leukaemia. Similar results have also been found near non-nuclear sites. One study in Cumbria suggested an association between childhood leukaemia and the employment of the fathers at the nuclear site at Sellafield and to the radiation dose to those fathers. The same study showed some association with employment in industries with no link to radiation, such as iron and steel, chemicals, and farming. At present, the cause of all those findings, including the observed clusters of leukaemia, is not established and further epidemiological and laboratory work is in progress.

Table 2

Some early effects of radiation doses delivered in a short time

Dose (mSv)	*Typical early effect*
50 000	Severe gastrointestinal and central nervous system damage; death rapidly ensues
10 000	Observable damage to exposed organ; death, probably within weeks, if whole body exposed
3000	50% chance of death if whole body exposed
1000	Some signs of radiation sickness if whole body exposed
500	Detectable changes in numbers of blood cells
200	Detectable chromosome changes in blood cells

Table 3

Typical exposures in daily life (millisieverts)

Dose (mSv)	
About 20	Annual doses received by individuals from radon in certain areas of the UK
5-15	Annual doses received by some maintenance workers in the nuclear industry
2	Typical x-ray examination of the lumbar spine
2	Typical annual dose from the natural background
1	Average annual dose received by nuclear workers, as a consequence of their work
0.1	25 hours in a jet aircraft at cruising height
0.05	Chest x-ray (single exposure)
0.001	Annual dose (averaged over whole population) from radioactive wastes.

94 The basis for believing that ionising radiation is capable of causing hereditary damage is that it can induce mutations in the cells from which sperm and eggs are derived. In fact, no evidence of an increase in hereditary defects caused by radiation has been found in irradiated human populations, such as the Japanese survivors, so estimates of the risk have to be made from animal experiments. These suggest that the risk of serious defects affecting future generations is well below that of fatal cancer to the present generation.

95 On the basis of all these studies, the risk of fatal cancer to a person is now reckoned to be 5 in 100 000 for every millisievert received uniformly over the whole of a life.[6]

96 This means that an extra radiation dose of 1 mSv to every one of a population of 100 000 people would result in five eventual deaths from cancer. We can reckon that each of us receives, on average, about 200 mSv in a lifetime from natural and man-made causes, and it follows that about 1000 people, out of any population of 100 000, die in the ordinary way from these causes of radiation.

97 We know that altogether, one person in four dies from cancer, ie 25 000 people in any population of 100 000. We can therefore say that roughly 1 in 25 of all cancer deaths result from natural or man-made radiation. Those who work in the nuclear industry and receive perhaps an extra **100 mSv** of man-made radiation in a lifetime will have their chance of dying from radiation increased by something like 50%,but it will remain only a small part of their total chance of dying from cancer. To put it another way, workers receiving a total of 100 mSv from their job will have their total chances of dying from cancer increased from about 25% to about 25.5%.

How radioactivity is produced and distributed in nuclear power generation

98 In a nuclear power plant (see Figure 5) a fission process in the reactor core releases heat that is then applied to produce steam which in turn drives turbines and produces electricity. The fission process produces ionising radiations. These are prevented from escaping from the core by thick concrete shielding, which absorbs practically all the radiation. However, because the fission process also releases neutrons, every material inside or immediately around a reactor core becomes radioactive. Moreover, in order to extract the heat and keep the reactor at the right temperature, a cooling medium (either a gas or a liquid) has to be pumped round a circuit through the core and out again to the boilers. When it emerges from the core, this coolant also contains radioactivity, so that in some stations the whole cooling circuit has to be surrounded by shielding.

99 Although the quantity of shielding is very great, power station workers receive small doses of radiation during the normal operation of the reactor and particularly when they have to carry out inspection or repairs to the coolant pipework and to pumps and boilers. In one or two older stations, where shielding is not so comprehensive anyone who remains for long periods on the site or near its perimeter receives a small dose above that they will also be getting from natural sources.

100 The waste liquid and gas which accumulates on a station is either routed to a waste treatment plant on the site, where it is further concentrated and the radioactivity retained, or it is discharged in small amounts to the air and to water. The quantities allowed to be so discharged are limited by the Authorising Departments (paragraph 31) who ensure that discharges are regularly monitored for compliance with the authorisations they have given and further that the licensee employs the best practicable means to limit the waste discharge to the lowest level reasonably achievable.

101 In considering what it is proper to authorise, the Authorising Departments take into account any reconcentration of the radioactive materials in the environment and their possible transfer to human foodchains.

102 At some stage, the spent fuel itself has to be withdrawn from the reactor core, and either reprocessed or stored. In the UK, most spent fuel is taken to the British Nuclear Fuels' nuclear chemical plant site at Sellafield for reprocessing. The passage of irradiated spent fuel from power stations to Sellafield by train in specially designed flasks is controlled by the Department of Transport. The reprocessing extracts the unused uranium and plutonium from the fuel so that they 5

can be re-used. The residues are mixtures of chemicals of varying radioactivity, including some that are highly radioactive and which are currently stored in specially designed containers on the Sellafield site.

103 The liquid wastes stored on the Sellafield site are gradually being converted from a liquid to a solid and more manageable state. Any site for disposal would be licensed and controlled as a nuclear installation. As with other nuclear installations, liquids of a very low radioactive content are discharged from Sellafield, following purification, to sea under the same arrangements as are described in paragraph 100. There is also a small discharge to air. The Authorising Departments (see paragraph 31) are responsible for regulating all these discharges.

104 The way in which power station accidents are prevented, and the levels of risk involved, are dealt with in the section on the risk of nuclear accidents, which also describes how risk assessment is applied.

105 In summary, therefore, there are three kinds of risk from any nuclear installation, each of which has to be separately dealt with:

(a) from radiation coming directly from the plant during its normal operation or maintenance (see the section on the risks in the normal operation of nuclear installations);

(b) from regular discharges of radioactive material to the environment (see the section on the risks in the normal operation of nuclear installations); and

(c) from accidents (see the section on the risk of nuclear accidents).

THE RISKS IN THE NORMAL OPERATION OF NUCLEAR INSTALLATIONS

106 As explained in the previous section, the normal operation of any plant will impose doses of radiation on workers and some members of the public. The difference between the exposures so resulting and those from an accident is that the former are actual, and continuously measurable, whereas the latter are a question of probabilities.

Figure 5 Reactor

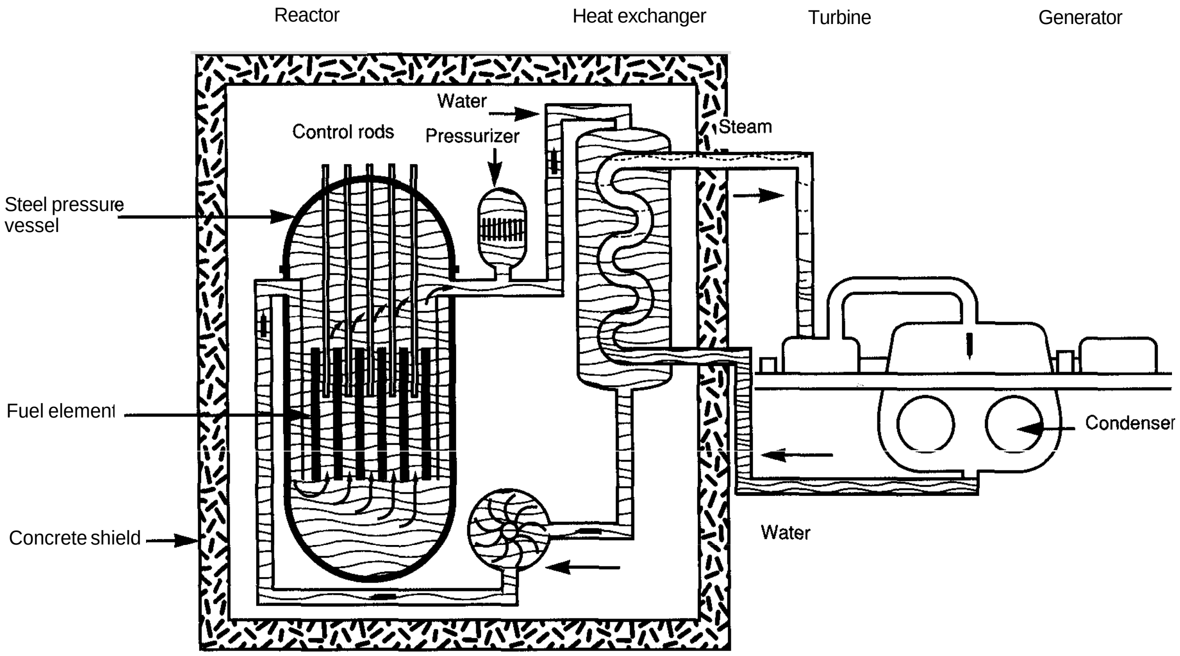

Note: This is a plan of a pressurised water reactor. Other kinds of reactor use different means (eg gases) to heat the water and make the steam; but the general principles are the same.

107 The dose limits for workers exposed to radiation and for members of the public are laid down in national[7] and international[8] provisions. In order to provide a strict measure of control, the maximum dose legally allowed is specified as an annual figure; in the case of workers the legal figure is 50 mSv in any one year and in the case of members of the public 5 mSv in any one year. The exposure of members of the public arises mainly out of discharges into air or water (as explained in paragraphs 100 to 103).

108 In fact, the Authorising Departments work to an average dose limit, over long periods of time, of 1 mSv. The Government has also made it clear that any waste disposal site must be so designed that the maximum risk to any member of the public should not exceed that equivalent to a dose of 0.1 mSv per year[9]. In practice therefore the legal limits are never nowadays approached either by members of the public or (see paragraph 109) by workers - unless there was an accident.

109 The legally specified limits of 50 and 5 mSv correspond to levels above which the risks have until recently been judged intolerable, on the very strict criteria applied to radiation risk. But as explained in paragraph 34, licensees are expected to ensure that the doses actually received are lower, to the limit of what is reasonably practicable (ALARP). Following revised estimates of the risk from radiation made by the International Commission on Radiological Protection, HSC acting upon the advice of the National Radiological Protection Board (NRPB) recently published revised guidance on dose limitation under which in effect the doses received by workers will be limited to 20 mSv per annum or less and which also requires an investigation to be carried out of the circumstances when a cumulative dose of 75 mSv or more is reached in any worker in any five consecutive years. This is recognised as an interim measure pending revision of the Ionising Radiations Regulations (1985)[7]. As explained in paragraph 79 the average dose for workers at nuclear installations is roughly 1 mSv per annum with a few workers in certain installations getting substantially more.

110 The doses which radiation workers receive are recorded on personal dosemeters which they wear. Employers are required to maintain records of these doses for inspection by the regulatory authorities and by the workers themselves. These records are used to confirm that safety controls are satisfactory and that the required limits are being complied with.

111 Workers are made aware of the risks they run, and have some choices in the matter. The public however cannot choose whether to be exposed or not; moreover there are specially susceptible groups of people. This is why the dose limits for the public are set at the lower levels described in paragraphs 107 and 108.

112 It is not possible to monitor the very small doses actually received by members of the public, these being indistinguishable from those produced anyway by natural radiation. The control is achieved as explained in paragraph 100 by strict limits on the amounts of different kinds of radioactivity which nuclear installations are allowed to discharge, by regular study of the ways in which radioactivity can get back into the food chain, and by protective measures such as restriction of particular foods, where environmental monitoring shows these to be required.

113 The significance of these doses and risks to individuals is further considered in the section on discussion of tolerable risk, particularly paragraphs 167 to 175.

THE RISK OF NUCLEAR ACCIDENTS

114 We have to rely on prediction and probability analysis to help us estimate the risk of accidental releases from nuclear power plants. The broad principles were discussed in the section on broad principles of risk assessment. It would be wrong to claim that such estimates are in anyway precise. Some components of the calculation are highly uncertain, some of the possibilities that have to be taken into account are very unlikely indeed, and some factors (including extremes of human aberration) are of their nature very hard to estimate. Nevertheless, as explained in paragraph 53, it is possible to build on the close study that has been given since the war to the reliability of plant within many industries to get a fair idea of what the chances are; and NII insists that this should be done as part of its examination of the safety of any nuclear plant that is to be built.

115 This section first explains how the chance of failure of the engineered parts of any nuclear plant is measured and reduced. It goes on to consider what account can be taken of contingencies such as natural disasters, external impacts and human error in operation.

Calculating the risk of plant failure

116 The hazard from radioactivity has long been recognised and the need to control its escape to very low levels was one of the first preoccupations of the designers of the first generation of stations. At that time there were no methods that could be used to quantify the risk of plant failure. The main safety precaution was therefore to ensure that all items of plant were exceedingly robust and that several layers of safety were built in where there was thought to be some chance of failure. Over the past few decades, however, techniques have gradually been developed for assessing the probability and consequences of failures and other events that could lead to accidental releases from nuclear and other forms of plant. These techniques are generally known as **Probabilistic Safety Analysis** (PSA); in other industries where there are major hazards the term more often used is **Quantified Risk Assessment** (QRA).

117 Figure 5 illustrates a 'typical' nuclear reactor. The principal hazard arises from the fission process in the core, and the most important aim in safe design is to ensure that the reactor can be shut down (ie the reaction stopped) quickly and the core cooled reliably. The reactor must 'trip' (ie shut down automatically) if operating limits are exceeded from any cause whatever, including an error by its operator or because of plant failure. It is also important to ensure that all radioactive material, eg in the fuel or in the coolant liquid or gas, remains contained within the structure.

118 A Probabilistic Safety Analysis (PSA) of a design begins with a careful identification of so called 'initiating events', that is, the things that could fail or go wrong and lead either directly or through a succession of other events, possibly including human error to a release of radioactivity. All the events that can be imagined and their possible consequences are plotted in the form of logical sequences called 'fault' or 'event' trees. Many thousands of these are plotted and considered for every nuclear plant design.

119 Systematic attention is then given to the means of preventing any of the possible sequences developing to the point where an important failure would cause a release. The availability of extensive data on the reliability or integrity of particular items of plant enables the chance of failures to be broadly estimated with a fair degree of confidence; as also the chance of failure of any safety precaution built in.

120 This approach enables the assessors and the designers themselves to ask a series of 'what if?' questions and also to consider when they should be satisfied with the answers, that is, when the chance of any of the fault sequences which could result in an accident has been driven low enough by provisions in the design. Paragraphs 122 to 134 describe some of the rules that are applied, and discuss the question 'how low is 'low enough'?'. Paragraphs 164 to 189 continue this discussion, and suggests risk limits that are tolerable.

121 PSA also helps the designers to produce a plant that is balanced from the safety point of view. The aim must be and is to reduce all important risks to some level of acceptability. Finally, it becomes possible to put an overall figure to the whole set, so that licensees can demonstrate that they have met whatever safety goals have been set for the plant as a whole. Though the figures are far from exact, they can give an idea of the magnitude of risk from engineering and from most types of human failure. Above all, the PSA process has the benefit that it ensures a systematic process of examination of the design and its risks, to which judgement and common sense as well as numerical calculations can be applied.

122 In assessing the design of any licensed nuclear installation in the United Kingdom the NII follows the guidance set out in its Safety Assessment Principle[3,4,5] (see paragraphs 6 and 43) which apply, in some detail, the general philosophy of tolerability set out in this document. The Safety Assessment Principles take into account the notion of a hypothetical person so situated as to be at greatest risk. (See also paragraphs 176 and 177).

123 The Safety Assessment Principles seek to ensure that any plant is so designed that a release of radioactivity that could give this hypothetical person an effective dose of 100 mSv or more

would have a very low chance of occurrence. A dose of 100 mSv is equivalent to an additional risk of fatal cancer of about 5 in 1000 using the current risk factors (see paragraph 95).

124 A release of the size that would produce this effect on a person at the critical point would diminish rapidly in significance at any more distant point. Nevertheless, emergency plans exist which would be activated in the event of any incident which might lead to consequences off-site (see paragraph 62).

125 **Design basis accidents:** Certain accidents are known as **'design basis'** to signify that the characteristics of the engineered safety systems of the plant will cater for them so that they do not produce unacceptable consequences. The implication is that any release bigger than that involved in a design basis accident could only occur as the result of the sequential failure of several levels of safety protection, or of some major and very unlikely event, such as the failure of the very strong vessel surrounding the reactor core. Such larger releases are called 'beyond design basis' accidents. They could range in size from those bigger than the design basis to very severe accidents.

126 **Large releases:** Even in the extreme, however, a nuclear power reactor could not explode in the manner of an atomic bomb because this requires a critical mass of fissile material to be forced together and held there while the nuclear reaction builds up to its full explosive force. For this to happen, very sophisticated and deliberate design provision has to be made. But even if, hypothetically, such a situation could develop in a nuclear reactor the very rapid heating in the first few moments of a reaction would disperse the fissile material and terminate the reaction. Thus the process would be essentially self limiting.

127 Though for these reasons a nuclear explosion cannot result from a power reactor, the Chernobyl accident showed that very large releases involving a high proportion of the fissile material and fission products in the core are possible. In recent years the possible size and content of potential releases has been carefully considered and a range of releases has been defined for purposes of calculating risks and devising counter measures. For purposes of the Safety Assessment Principles5 two particular types of release have been identified as a basis for design assessment; these are specified in more detail than the types of release considered in the original 'Tolerability' document. For one, known as a 'large release' (though it covers levels of release very much smaller than Chernobyl), the objective of the assessor is for the chance of the event to be reduced, so far as reasonably practicable (paragraphs 25 to 50), to not more than once in 10 million reactor years. For the other, a still smaller but 'beyond design basis' accident which could give a dose of 1000 mSv or more to the hypothetical person at greatest risk (paragraph 122), the objective of the assessor is for the chance of the event to be reduced, so far as reasonably practicable, to not more than once in 1 million reactor years. The working rule that is adopted for this purpose is that no single sequence of possible faults should be calculated to have a chance of more than 1 in 10 million per annum of causing the second type of release.

128 **Methods and results:** The following is a highly simplified illustration of how this is done and checked. Let us suppose that the chances of failure of the main pumps providing feedwater to the boilers (or steam generators), which could cause the reactor to overheat, is after a study of failure of similar systems elsewhere shown to be equivalent to once in 100 years of operation (1 in 10^2 per year). Such a failure if it actually occurred should be sensed by the instrumentation and cause the reactor to shut down. From experience and from experiments and analysis, it is known that shutdown systems have a very high reliability indeed but a chance of failure is nevertheless assigned to them, at about 1 in 100 000 per demand for shutdown (1 in 10^5). So the overall chance of a main feedwater system failing and the reactor subsequently failing to shut down is calculated on this basis to be about 1 in 10 million (1 in 10^7 per annum (1 in 10^2 x 1 in 10^5)), which is judged sufficiently remote to require no further precautions unless there is any possibility that some event affecting the main feedwater system could also affect the chance of failure of the shut-down system.

129 The assessor has then to consider for example a less serious sequence in which the main feedwater system fails, the reactor shuts down correctly, but the emergency feedwater

system does not remove the residual heat. Let us say that this chance is found to be greater than 1 in 10^7 per annum. The assessor is likely then to insist that not one but two such emergency feedwater systems must be installed so that if one system fails the heat will be removed by the second. This applies the principle of **redundancy** [7] (see paragraph 53).

130 However assessors may next consider that **both** of the emergency feedwater pumps they have demanded could be put out of action by a common cause - let us say the failure of the electrical supplies. They may on this account demand that there must be two emergency feedwater systems operating in a different manner, perhaps one using electrical drives and the other steam turbines. This applies the principle of **diversity** (see paragraph 54); it helps to avoid what is known as 'common mode failure'.

131 The use of fault and event 'trees' (see paragraph 118) ensures that such possibilities are considered in a systematic manner, so that at each point the objective is to ensure the provision of sufficient layers of protection to reduce the chances of releases to those specified in the Safety Assessment Principles.

132 It must not be supposed that precise numbers can be put to all the events that might take place. Though there is an increasing availability of plant reliability data, there are many gaps where engineering judgement has to be applied to produce a number; and judgement, too, is often needed to fortify numbers derived from reliability data. Moreover, there are some sources of uncertainty that it is difficult to take account of in a fully systematic way. These, and how they are dealt with, are described in paragraphs 140 to 152.

133 These calculations, and the systematic approach made possible by Probabilistic Safety Analysis, give a good measure of assurance that, provided that a plant is well and reliably constructed and maintained, and satisfactorily operated, the chance of failure leading to a release significantly in excess of the 'design basis' should be very low indeed (see paragraphs 127 and 162).

134 As has already been explained, these numbers cannot be precise which is why it is necessary to use such terms as "of the order of". They may become still less precise when the various chances attaching to different possible failure sequences are added up to give an overall risk of serious failure attaching to a plant as a whole. The huge number of such calculations made for any plant gives considerable assurance that the various doors to important failures are securely blocked. However, no-one can be completely certain that nothing has been overlooked, or that forms of deliberate intervention or interference, or unexpected common modes of failure cannot occur - even though in any conceivable circumstances, the multiple safety systems in a nuclear plant have a very high chance of succeeding in their purpose from whatever unexpected angle they may be challenged. When therefore the overall chance of failure is said, after all these calculations and judgements to be, say 1 in a million (1 in 10^6) it would be wrong to take this in a literal sense. What is being said is that the chance is two dimensions lower than the sort of chances we can easily judge from common knowledge and experience - such as one's chance of being killed this year in a motor accident (1 in 10^4).

Safety-critical computing systems in nuclear plant

135 Modern plant is becoming increasingly complex: new technology allows us to have finer control and so run processes more efficiently; it also provides novel means of monitoring for increased safety. An important factor in these changes has been the increased use of computer systems, both to control plant and to help ensure its safe operation. Computers can monitor enormous amounts of data and thus allow plant designers to do things that would not be practicable with conventional engineering hardware. In nuclear power reactors, computers have been used for several years to help the operators control the plant by presenting them with information in ways that are easily understandable: they are now also being used to control the plant automatically and shut it down safely in an emergency. In principle, and often sometimes in practice, computer systems can provide us with means of control that are more efficient and safer than would otherwise be possible.

136 However, the complexity of both the plant and the computer systems brings problems as well

as benefits: the greater the complexity, the harder it is to be sure that there are no hidden design faults which can reveal themselves as failures during operation. This problem is particularly acute for computer software, which may be unreliable if its designers have not fully understood all aspects of the plant being controlled or if mistakes have been made in writing the program. Some software faults are minor and are relatively easy to detect, rather like spelling mistakes. Fundamental design faults are potentially much more serious and the more complex the software, the more likely they are to arise and the more difficult they are to detect.

137 Good design practices can assist. Safety-critical software needs to have the greatest simplicity compatible with the system performing its tasks. The temptation for a designer, released from hardware limitations, to build programs of ever greater complexity has to be resisted. In addition, separation between computer programs performing safety functions, such as emergency shut-down, and those providing normal everyday control leads to a simpler safety system, aids an understanding of how it works, and minimises the chance of a single fault causing both loss of plant control and failure of the emergency shut-down.

138 Modern software practices include formal mathematical techniques, which sometimes allow a rigorous proof that a program will behave as specified. However, such a specification must itself be a formal mathematical document, and there is always a chance that this does not adequately capture all the engineered provisions for the safe functioning of the overall plant. Another recent approach has been to try to make software systems 'fault-tolerant' by using design diversity. This technique appears to bring useful, but not dramatic, benefits.

139 It is generally not practicable to test a computer program exhaustively for all its different input signals, since the number of different combinations of these is usually astronomical. Fairly modest levels of software reliability can be demonstrated, using a sample of these inputs that is statistically representative of operational use, but these fall short of the levels that are currently demonstrated for comparatively simpler, conventionally engineered hardware systems.

140 For all these reasons, there are at present important limits on the extent to which computer software is relied upon for safety-critical functions in nuclear reactors. At present, a probability of failure of about 1 in 10 000 per demand is the best that can justifiably be claimed. It should be emphasised, however, that these limitations do not preclude the use of computers in safety roles. In the case of a reactor protection system, for example, a primary computer-based system of modest claimed reliability can be backed up by a simpler, conventionally engineered ('hardwired') secondary system to provide the necessary confidence that shutdown will occur when needed. Such systems are in fact insisted on in the design of modern UK nuclear plant, and strongly recommended in other industries also.

Natural events

141 A second source of the uncertainties referred to above is the chance of some external event such as an earthquake which, if severe enough, would affect all systems in different degrees and so upset the calculations described above. NII require that a modern plant be so designed as to be able to withstand safely all earthquakes except those of a severity whose chance of occurrence is judged to be less than 1 in 10 000 per annum. with such provision the plant would generally be expected to withstand safely even larger events (paragraph 142). Checking this prediction for British conditions is difficult, since Britain is singularly free from significant earthquakes.

142 Even greater uncertainty attends calculations of the chance or consequences of some substantially bigger earthquake than this or some even less likely natural event such as a massive inrush of seas or the impact of a large meteorite. Such events if they occurred would of course create very widespread dangers quite aside from their possible effect on nuclear plants. It seems therefore unreasonable to require a plant to demonstrate ability to cope with every natural event we might imagine; but rather to insist on margins of strength in all major components so that it is likely to do better than any estimates suggest.

The human factor

143 In trying to estimate the total risk, the possibility of **human error** must also be taken into account. It may enter, not merely into the use of

the plant once it exists, but also into the original decisions taken when that plant is planned and introduced. Human error in the **design** of the plant is one of the things that the systematic approach described above is intended to reduce; with its emphasis on tracing out possible fault trees and their analysis by PSA. Human error in **construction of** and **modification to** plant can likewise be, and is, minimised by the application of modern quality assurance and reliability systems. Moreover, some errors of this type are already taken into account in the risk figures used by PSA, since as explained these are derived from the failure rates of other operating systems constructed and modified in similar ways.

144 Human errors in **operation** or **maintenance** of a plant, though also taken into account in the risk figures, require a different method of control. Some human intervention is required even in a highly automatic system, because although people may make mistakes, they also have a capacity unparalleled by machines, to operate skilfully and make correct judgements in unprecedented situations. This unique human capacity has to be preserved, while deliberately making harmful intervention as difficult as possible.

145 To do this, the exact point at which people are able to intervene in the automatic system, and what actions they could possibly take has to be carefully planned. The goal is to make sure that any possible error has a low probability of causing serious hazard. As in the case of physical components, this is done by analysing the possible sequences of events that might result in a release. From this analysis, one can locate the points at which human action can occur in the chain; and the type of decision the person needs to make at such points.

146 The probability of an error can then be estimated from several sources. On the one hand, it may be known what error rates are found for people performing such tasks elsewhere. It may be possible to carry out deliberate experiments, or to use the results of those done elsewhere for other purposes. Finally, it may be possible to extrapolate by starting from what is known about slightly different situations, and then correcting for special causes of ease or difficulty in this particular decision. Analysis of this kind emphasises the importance of the instrumentation available to the operators, of their controls, and of their working environment, since the probability of error does depend on the extent to which they have clear and unambiguous information on which to make their judgements.

147 In many possible chains of events, human beings do not intervene; the engineered safety systems provide a defence in depth, designed to provide an appropriate response without further human intervention. Typically, several such protection systems must fail to respond adequately before the safety of a reactor can be seriously jeopardised. They are also normally simpler than the system they protect, because they only have to prevent or stop a process rather than control it. Where a human action can prejudice only a single safety system, therefore, the impact on overall risk from that action may be slight, and the accuracy of estimation is then likely to be sufficiently high.

148 The actions of most concern are those that affect several protection systems at once, since this undermines the 'defence in depth'. In the accident at Three Mile Island, operators misinterpreted and reacted incorrectly to the signals they were receiving and deliberately suspended their most essential safety system, the emergency core cooling arrangement. In the disaster at Chernobyl, successive safety devices were removed to produce a free run for an experimental test. Such events have provoked concern as to whether the same kind of thing could happen in Britain. Typically, these actions involving intervention in multiple safety systems mean that the operational staff have seriously misunderstood the state of the plant, and have undermined its safety, often with the best of intent. Such inappropriate plans of action may involve deliberate, although well intentioned, violations of essential safety rules and procedures.

149 It is possible to apply Probabilistic Safety Analysis (PSA) to human errors in ways resembling its application to mechanical systems. The class of actions considered above can be represented in PSA currently in two ways. First, PSA studies need to allow for the dependencies between safety systems, that is to say situations where the overall reliability is significantly less than if systems were totally independent. Past experience of incidents is used to indicate factors that reduce or increase dependence. The overall

reliability is then represented by a 'dependence' or 'common cause failure' model.

150 Second, PSA seeks to address human error directly. For errors of skill, slips and lapses, there are well established models and data that can give reasonable numerical estimates of human error. For 'knowledge-based' actions, which include the errors of misunderstanding mentioned above, the situation is more difficult. Since such incidents are less common, our database of past experience is insufficient for a reliable estimate specific to a particular plant design. However, we do have a reasonable understanding of factors that help to avoid or mitigate these potentially serious errors. Such defences include diverse supervisory staff, computer-based diagnostic tools, well-designed symptom-based procedures and alarms, built-in time delays, good instructional systems, providing feedback on the effects of actions taken. They also provide against interference with any system while thought is given to the cause of any difficulty.

151 PSA can therefore evaluate qualitatively the strength of these defences, so as to say which of two situations is safer; even though we cannot predict quantitatively the exact frequency with which these knowledge-based errors may occur. We can also judge an upper bound on the contribution these errors make to the risk from the plant, because by definition errors for which there are few data are very infrequent events. In time, our ability to model this class of errors will improve both with increased data from past experience, and with more validation of the models that are currently being proposed.

152 The total of these approaches to human error analysis should be able to demonstrate, and in some recent cases has demonstrated, that a well-designed plant can limit the effect of operational human errors to being of the same order as more conventional engineering failures. Since the main residual uncertainty involves procedural and knowledge-based error, the NII regards the scrutiny of operating procedures and contingency planning of all kinds as being as much a part of its business as attention to design, or to the actual operation's record and log. The need to specify procedures and follow them is a part of the licensing conditions of any station. The procedures and any incidents throwing light on them are regularly reviewed.

The safety culture

153 Good procedures must be supported by a culture which disposes every individual in the organisation towards safe working practices. The lead has to come from the top, with the Chief Executive assuming active personal responsibility for the safety of the organisation's employees and for the impact of its operations on the safety of the general public. There needs to be an effective system of safety management which the most senior management in the organisation clearly support and are involved in. Managers! safety responsibilities at all levels need to be correctly allocated and clearly defined. They should be allocated to line managers with safety specialists acting as advisers. Similarly, responsibility for ensuring that operations are safe cannot be sub-contracted to consultants.

154 As well as properly allocated and defined management responsibilities, an effective system of safety management must include well considered and articulated safety policies; organisational arrangements which ensure control and promote cooperation, communication and competence; planning systems which identify objectives and targets, establish priorities and set performance standards; performance measurement which actively monitors the achievement of plans and the extent of compliance with standards, and reactively monitors accidents, incidents and ill health; and effective systems for auditing and performance review. These essential elements of a successful safety management system are described in detail in a recent HSE publication. The effect of such a system should be that top management's commitment to high safety standards permeates every level of the organisation.

155 As with operating procedures and contingency planning, HSE regards assessment of an organisation's safety management systems, safety culture and safety performance as critically important and it takes them very much into account both in the Safety Assessment Principles on which NII bases its licensing assessments, and in programmes of inspection. Evidence from a number of industries suggests that effective safety management is associated not only with reduced frequency of incidents but also with commercial success.

Safety in operation

156 The operation of a nuclear reactor does not present unusual difficulties by comparison with that of other major industrial plants and, apart from its shut-down system, which is designed to operate very quickly indeed, it responds quite slowly to operational signals. Manual controls do not need to be activated quickly: the operator has time to think.

157 The control systems are designed in such a way that any attempt to continue to operate the reactor outside its safe limits of pressure, temperature, etc results in the plant shutting down. These systems operate independently and cannot readily be defeated by the operator in the event of any of the plant failures foreseen and provided for in the design process. It would be extremely difficult to defeat the protective systems of a British plant in the way that occurred at Chernobyl, because of the very large number of triggers that cause reactors to close down quickly and automatically. It would hardly be possible for a determined attempt to defeat the systems to escape the attention of competent management. That, however only serves to emphasise the importance of careful selection and training, and of thorough attention to safety attitudes.

158 Because of the nature of a reactor, a sequence of faults which could potentially involve an uncontrolled release would be unlikely to happen very quickly. The final part of a sequence might indeed happen rapidly and perhaps catastrophically but the precursors would be likely to develop slowly, as at Chernobyl, where the systematic and deliberate removal of safety systems and subsequent operation of an undefended reactor occurred over a period of many hours. There is therefore time for control. The automatic systems built into modern designs (paragraphs 144 to 145) do not require human intervention for the first half hour after shut-down. This avoids any necessity for rushed decisions by the operators which might be based on faulty diagnosis and affect operating conditions or prevent safe shutdown, as happened in the Three Mile island incident.

159 If, however, some series of faults continued to develop after an initial period of automatic operation it would neither be possible nor desirable to continue not involving the operator. At this point, as pointed out in paragraph 144, the adaptability of the human beings who understand a plant's capability and who have had time to think carefully about the signals that have been reaching them is an asset. There has also been time to key into a prearranged formal system of technical advice or to obtain the appropriate level of managerial clearance to take non-standard actions. The training and retraining of operators on simulators to deal with emergency situations is required by NII and is a standard part of operational procedure.

160 This combination of safeguards cannot of course make human error impossible. What it is designed to do is to assure the competence and motivation of those who manage and operate nuclear power plants, and to restrict the scope for serious error to situations that are covered by precautions such as automatic shutdown, or to extremes of wilfulness that could not escape detection. It is intended to maximise the amount of thinking time available to the operator, and the skill and knowledge that can be applied to any situation.

161 Having considered all these factors there remains the difficulty of assigning some degree of probability to a release of radioactivity caused either directly by the human factor or by some combination of plant fault and error. We discussed earlier, in paragraphs 143 to 155, the current position with regard to human factors not otherwise quantified, which naturally operate for both good and ill, and the positive impact that proper training, management and regulatory procedures can have. We have to bear in mind too that the safety systems provided for the failures that have been foreseen will also cover situations that have been overlooked, whether these arise from plant failure or human error. Our conclusion is that it is possible to design the system so that the influence of the factors referred to above will be to increase the probability of large releases by less than a factor of ten. We also consider that the conservatisms built into the design of modern nuclear power stations should ensure that other "unquantifiable" hazards, such as very remote natural events, should add very little to the risk derived from comprehensive Probabilistic Safety Analyses (PSAs).

162 The original 'Tolerability' document examined the chance of plant failure leading to what was

called an 'uncontrolled release'; and concluded that the chance of such an event was of the order of 1 in 1 million per annum. This, as the document explained, had to be adjusted to take account of the then unquantifiable factors and on that basis, was estimated to be greater than 1 in a million but less than 1 in 100 000. Modern forms of PSA take many of these factors into account and we are now able to say that the chance of a large release, as described in paragraph 127, is significantly lower than 1 in 100 000 per annum.

Mitigation and emergency procedures

163 No matter how improbable a serious accident may be, it is essential to have a **safety net** for the protection of the public from radiation and for the evacuation of people if this becomes necessary. It is also necessary for the regulatory authorities to have well rehearsed procedures for assessment of the outcome of any accident that might develop and for decisions for people's safety, protection, and information to be taken in good time. It is a mandatory requirement of the site licence of all nuclear installations that the licensee should make suitable approved emergency arrangements including liaison with local authorities such as police, hospitals and fire services and for immediate notification of Government and regulatory agencies. The licensee and the other authorities rehearse and demonstrate these to NII in emergency exercises. The HSE publication, *Emergency Procedures for UK Civil Nuclear Installations*[11] describes these arrangements.

DISCUSSION OF TOLERABLE RISK

164 Paragraphs 106 to 113 considered the risks in the **normal operation** of a power station, which are principally to workers but which also include those associated with authorised disposal of wastes of low radioactivity. Paragraphs 114 to 163 considered the risk of an **accident,** either of the 'design basis' kind or of one leading to a 'large release'. It now remains to discuss what these risks actually mean when they are translated into risks to individuals (paragraph 57) and to society (paragraph 63) and how much they amount to when compared with the other risks we commonly accept.

165 It is necessary once again to emphasise the considerable uncertainties that attend all calculations of risk and in particular to what was said in paragraphs 114 to 163 about estimates of remote risk. Where figures are used as they are in the following paragraphs, they are unavoidably broad estimates.

166 Before turning to the figures, we must also remember what was said about risk in paragraphs 10 to 24. It is quite reasonable not to want to accept a particular kind of risk, so long as we have first ascertained what is known about it. In doing so we have to consider it in due proportion to other risks, and particularly to the other kinds of risk we are usually in fact choosing when we shun one particular kind.

Levels of individual risk

167 As explained more fully in the fourth section, the effect of any exposure to ionising radiations is mainly to increase the statistical probability of contracting cancer later in life. For workers exposed over a number of years to low levels of radiation from the day to day generation of nuclear power the associated risk of radiation-related death is essentially zero over the first 10 to 20 years and builds up to a peak some 10 to 20 years after they cease to be exposed. In reality therefore the risk to radiation workers additional to that from natural sources will generally be expected to express itself in their later years, if at all. In order to permit comparison with conventional risks it is necessary to average the total radiation risk over the number of years of exposure.

168 **Risks in normal operations:** It was explained in paragraph 79 that the average levels of dose received by **workers** exposed to radiation at nuclear installations vary between about 1 mSv and 5 mSv per annum; at power stations the level has usually been less than 2 mSv. Applying the latest risk factors, the average risk of death associated with an annual dose at these levels would be between 1 in 20 000 and 1 in 4000 per annum with a risk of 1 in 10 000 or better at power stations. The upper figure is broadly comparable with the risks borne on average by the workforce in such heavy risk industries as metal manufacturing and mineral extraction, and the lower figure with 11

the average for manufacturing industry. Further comparisons are set out in Appendix 2. The bias from the fact that the radiation risk is deferred is partly balanced by the fact that radiation workers will also bear some of the conventional risks in addition.

169 The level of risk borne by the very small number of workers whose dose is near to the level of **15** mSv recommended by the National Radiological Protection Board as not to be regularly exceeded would probably approximate to that of many workers in the riskier groups in risky industries; such as that of workers in the offshore oil industry, faceworkers in mining, or roofworkers in the construction industry. The level of these risks is difficult to estimate precisely because of gaps in the statistics, but we can say that broadly, a risk of death around 1 in **1000** per annum is the most that is ordinarily accepted by substantial groups of workers in any industry in the UK, with that level being exceeded only by fishermen and relatively small sub-groups such as helicopter pilots, divers and demolition workers. It seems therefore reasonable to adopt a risk of death of around 1 in **1000** as the dividing line between what is just about tolerable as a risk to be accepted by any substantial category for any large part of a working life, and what is unacceptable for any but fairly exceptional groups.

170 In his report on the Hinkley Point Inquiry, Barnes[2] suggested a limit of 1 in **5000** per annum as the average risk to workers on a plant. We have considered how far this should be required as a regulatory benchmark (see paragraph **39)** and have concluded that it is more appropriate as a managerial check rather than a standard of tolerability since it is entirely dependent on the number of workers included in arriving at the average, eg whether temporary or contract workers or non-radiation workers are included or not.

171 For reasons explained in paragraph 111, the maximum tolerable risk levels for workers are very much greater than those which apply, by law or in fact, to any **members of the public,** even those who live near enough to nuclear installations to face any meaningful risk from the authorised emissions during normal operation. The precautions taken to protect them are explained in paragraphs **100** to **103,** and the legal dose limits applying are discussed in paragraphs **107** and **108.** As explained there, people living near to plants will actually receive on average much less than the legal limits.

172 If the maximum tolerable risk for any worker is set at around 1 in **1000** per annum (see paragraph **169),** it seems clear that the maximum level that we should be prepared to tolerate for any individual member of the public from any single large scale hazardous plant, nuclear or other, could not be less than ten times lower, ie, 1 in **10 000** (1 in 10^4**)**. Such a level would as it happens equate to the average annual risk of dying in a traffic accident, and can be compared with everyone's general chance of contracting fatal cancer, which is an average of 1 in **300** per annum.

173 Barnes[2] proposed that the limit of acceptability of risk to individual members of the public from such plant should be set at 1 in 100000 (1 in 10^5) per annum. We consider that Barnes was, in effect, saying that, in order for him to find the Hinkley Point 'C' proposal acceptable, the maximum risk to any member of the public from its operation should be 1 in **10^5** per annum. This is not the same as saying that for every industrial plant in the UK the maximum tolerable risk to any individual member of the public should be less than 1 in 10^5 per annum. **We propose to maintain our existing position that a risk of 1 in 10^4 per annum to any member of the public is the maximum that should be tolerated from any large industrial plant in any industry with, of course, the ALARP principle applying to ensure that the risk from most plant is in fact lower or much lower. But, in accordance with Barnes' findings, we propose to adopt a risk of 1 in 10^5 per annum as the benchmark for new nuclear power stations in the UK,** recognising that this is, in the case of a new station, broadly achievable and measurable.

174 In practice, the measures taken for nuclear installations mean that the risk borne on average by members of the public in the vicinity of a plant from its **normal operation** will generally be no more than 1 in 1 million (1 in 10^6) per annum. To obtain the overall risk in terms of Barnes' proposed formula (paragraph **173)** the risks from possible **accidental events** have to be added (paragraphs **176** to **177** and Figure **6).**

175 Having considered what might be regarded as levels of risk that are **just tolerable** or can be used as benchmarks we must now consider what might be a **broadly acceptable** risk to an individual dying from some particular cause, ie, what is the level of risk below which, so long as precautions are maintained, it would not be reasonable to consider further improvements to standards if these involved a cost. **This level might be taken to be 1 in a million (1 in 10^6) per annum** bearing in mind the very small addition this would involve to the ordinary risks of life (paragraphs 22,23 and Appendix 2). An annual risk of 1 in a million is of course not altogether negligible; it is broadly the same as that of being electrocuted at home (and is about a hundred times less than the annual average risk of dying in a traffic accident). But it is a level of risk which, provided there is a benefit to be gained, and proper precautions are taken, does not worry us or cause us to alter our ordinary behaviour in any way.

176 **Accidental risks:** For a nuclear plant designed just to meet the NII's original Safety Assessment Principles[3] the risk of a fatal cancer to the hypothetical individual at greatest risk from possible releases of all sizes added together, has been estimated previously to be in the region of 1 in 1 million per annum[11]. But, as with the case where we were considering the probability of a large release some further allowance has to be made for those elements that are of their nature very difficult to predict and which give an added uncertainty and therefore an added risk.

177 On the other hand the risk calculations discussed in paragraph 176 took no account either of people's real pattern of behaviour, which do not involve constant exposure to the risks, or of the mitigation actions (paragraph 163) that would help to protect them. Taking these real factors into account, the risk to any individual in the UK of dying from cancer caused by a nuclear accident even if he or she lives quite near to a station is a great deal lower than that of the hypothetical most at risk person; and provided that the plant conforms to NII's Principles as revised, should be rather less than 1 in 1 million per year. If however we were to **add up** the risks to the range of people living near a plant **both** from ordinary operation (paragraph 174) and from an accident (paragraph 176), we might conclude that **most people in the vicinity are at or near the 1 in 1 million level and well below the benchmark of 1 in 100 000 (1 in 10^5) per annum. Some people might be near to the benchmark, while a handful could be a little above that level.** The risk from such a plant to the average person living elsewhere in the UK would be very much below these levels. Figure 6 illustrates the actual levels of risk by comparison with the level that might be regarded as just tolerable.

Levels of societal risk

178 For people living nearby any hazardous plant, nuclear or otherwise, the principal consideration is the risk to themselves and their families, ie 'individual risk' as just described. In connection however with large accidents with potential effects going well beyond those on human life, society needs to base its judgements on some measure of the whole risk and harm, as explained in paragraph 63. It will be recalled that such a measure is called 'societal risk', ie the risk borne by society as a whole in relation to the totality of the potential harm.

179 As with individual risk, judgement has to be assisted by comparison with other societally regulated risks and harms. In the nuclear case we have to take into account the possibility of a release where a sequence of safety systems has broken down (paragraph 127). If ever it becomes possible to assign a 'total detriment' (paragraph 66), we could do sums on the basis of calculating the risk and harm of a series of possible accidents including one with an exceedingly remote chance, of the size, say, of Chernobyl, and others of greater chance of occurrence but of much smaller size. As matters are, the only alternative procedure is to select an accident of some considerable size, treat it as a point of reference, and compare it with other major events against which precautions are taken. We then have to make allowance for the possibility of much larger and exceedingly improbable events and much smaller ones that are more likely.

180 A large number of studies have been done in recent years on the probability of very large **non-nuclear** accidents of various kinds, as an aid to determining what it is worth paying to reduce the risks. Of these, the largest and the most thorough, which is nowadays taken as a model and standard of comparison in risk studies world wide, is the

survey carried out by HSE some years ago of the potential of the industrial installations at Canvey Island on the Thames for causing a major accident affecting the surrounding population. This led to improvements to safety which were calculated to have reduced the risks to a chance of about 1 in 5000 per annum of a major accident capable of causing more than 1500 casualties.

181 Again, when the Thames Barrier was built, the specification insisted that the chances of its being overtopped by a freak tide should be less than 1 in 1000 per annum. (This as it happens is also the predicted annual chance of an aircraft crash in the UK killing 500 or more people, ie of unprecedented size). Both a serious accident at Canvey Island and the overtopping of the Thames Barrier under extreme tidal conditions would have catastrophic effects well beyond killing numerous people. Like hypothetical major nuclear accidents they therefore represent standards of societal harm and risk whose proper measure should theoretically be in terms of 'total detriment'.

182 We might deduce from these rough comparisons with Canvey and the Thames Barrier that where we have little choice but to accept a major societal risk, we require its chance of occurrence to be reduced to less than 1 in 1000 and, if possible, less than 1 in 5000 per annum. A 1 in 1000 chance of a major event occurring per year can perhaps be taken as the maximum calculated major 'societal' risk we are prepared to tolerate; and we aim to do better. Society might very reasonably demand a lower order of risk than this where, as with nuclear, we have some choice whether to accept it or not.

183 In considering how to compare a major nuclear risk with the kinds of risk discussed above, a number of factors enter in. First, **any one of the plants** that are part of a programme of reactors could be the source of such an accident, so it is not just the risk attaching to one plant that matters, but the greater risk attaching to the whole family of plants. Second, we have to bear in mind that a very large nuclear accident, say approaching the size of Chernobyl, would have long term economic effects, eg in rendering land and buildings sterile, and this has to be taken into account in any estimation, however rough and ready, of total detriment.

184 Third, we have to consider the proposition that people feel greater aversion to death from radiation than from other causes, and that a major nuclear accident could have long term health effects. Against this, as previously explained, deaths from most non-nuclear accidents are immediate; most people whose death could be attributable to some nuclear event would in fact suffer no immediate harm, but would die earlier than they otherwise would, from the eventual development of a cancer. To put it another way, among the group of people affected by the radiation, most would suffer no harm, rather more would die from cancer than otherwise would, a few would be likely to suffer early effects and a degree of harm would probably be passed on to future generations. It is hard for the human mind to compare such an outcome with, say, the chance of large numbers of people dying from burns in the event of an accident at Canvey Island or being drowned or killed by other consequences if the Thames Barrier were overtopped; but making the attempt helps us to grasp the size and nature of the risk.

185 To achieve such a comparison, we first have to specify a nuclear accident whose effects might, taking everything into account, be thought to be of the same order as Canvey and then to consider the risks of its happening. On the whole, it seems right to specify an accident which could lead to the immediate or eventual deaths of 100 to 1000 people. By specifying such a range, we are recognising explicity that any particular accident for which we can establish a scenario could have greater or less effects according to such wholly arbitrary factors as the weather. Also the chances of an accident leading to the deaths of a hundred people do not seem, so far as we can calculate, greatly different from those of one leading to the deaths of a thousand people, whereas the chances of an accident of a much greater magnitude are calculated to be substantially lower.

186 The chances of a nuclear accident with the effects specified above at any one plant are, so far as can be calculated, very much lower than the chances attaching to an accident at Canvey Island. So far as can be calculated, and taking as much account of the human factor as is practicable in modern forms of risk calculation, a programme of between 20 and 50 modern nuclear reactors would have a similar chance of causing death to some

hundreds of people as the installations at Canvey Island.

187 There are about 35 nuclear power reactors of varying designs in current operation in the UK at the moment, some of them approaching the end of their working lives. Their owners have sought to calculate the risk of an accident involving a large release (as discussed in paragraph 127) as part of the periodic safety review procedures required by NII (paragraph 48). Uncertainty attaches to these calculations, in addition to the uncertainties already discussed, for example, in paragraphs 114 and 161, since, although the reactors' designs are very robust, their design and construction was not governed by modern systems of quality assurance or quantitative risk estimation. Nor in any calculation we might make can we include risk estimates from the numerous French nuclear reactors, some of which are in principle capable of creating a nuclear accident which could affect southern Britain. The risk flowing from these and from existing reactors in the UK is in effect one which exists and is accepted, subject to the close consideration currently being given to the continued safety of the Magnox plants (paragraph 49).

188 It is in principle possible to suggest a 'tolerability limit' for a hypothetical programme of modern nuclear reactors in the UK, in terms of the chance of occurrence of a major accident, bearing in mind that such a major accident would have consequences well beyond the vicinity of the plant. This was the approach adopted in the 1988 version of this document, which concluded that an overall risk of one considerable accident per 10 000 years from any one of a programme of modern reactors would be just tolerable. We also pointed out that (on risk calculations that took little account of the human factor) a programme of 20 modern reactors might have a chance roughly ten times less than Canvey Island producing an accident of comparable size.

189 There clearly is some maximum to the number of reactors that could be tolerated in any future programme. But it continues to be difficult or impossible to specify what this should be in relation solely to current estimates of risks. In the first place we do not know what risks will attach to future designs of nuclear power stations. Furthermore, in practice, decisions on new nuclear power stations would have to be taken one at a time, and different proposals may well be made by competing utilities. A decision as to whether or not the aggregated risk is tolerable would have to be taken on the basis of a very broad estimate, taking into account whatever benefit society chooses to assign to nuclear power, as for example any economic benefits or disbenefits arising from alternatives such as the large scale use of fossil fuel. It seems right therefore to bear these considerations in mind but nevertheless to apply a limit to the risk of a major event at any one plant that could have 'societal' consequences.

CONCLUSION

190 The intention of this document has been to make clear what the present safeguards are and how they are exercised, following public discussion based on expert assessment of the original discussion document on this subject published in 1988. It is the best estimate of the position that HSE as a regulatory body can make, having taken the best expert opinion, including that of its Nuclear Installations Inspectorate. It is not however for the regulatory authorities but for Parliament and the public to weigh the benefits of nuclear power with the risks we have outlined.

Figure 6

Tolerable and actual levels of risk to workers and the public

Where appropriate the specified risk ranges assume that the risk factors will be increased as recently suggested by NRPB

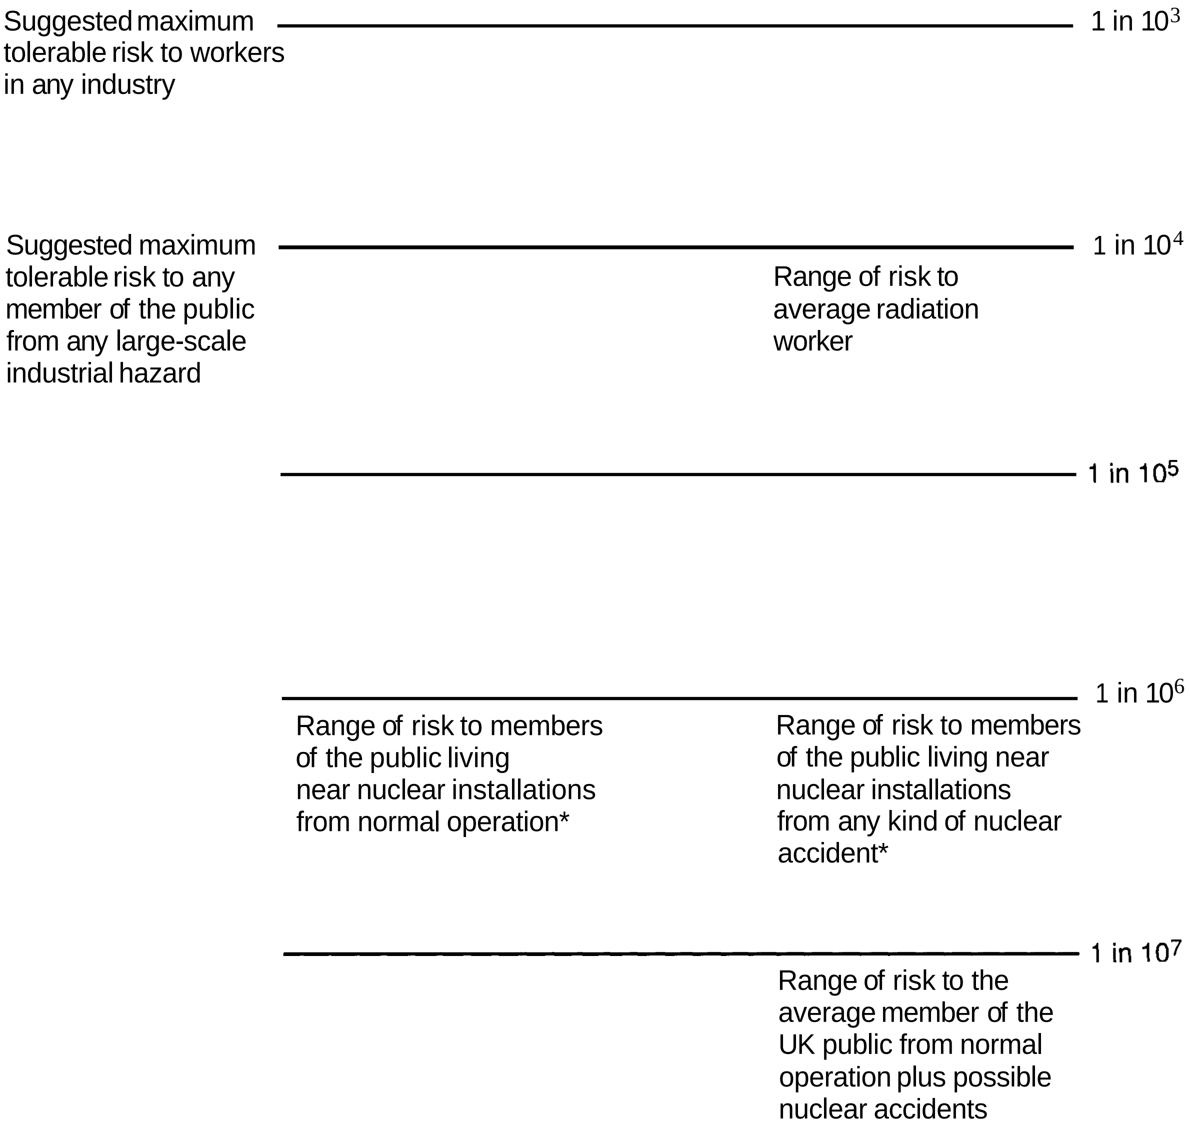

* It is very difficult to assign a probability to the risk borne by people who live close to a plant from its normal operation, since any doses which may be received by individuals are not only very small but are unascertainable; for instance only a very few people living close to a few plants are regularly exposed. The estimate gives only a broad idea of the risks borne by the whole range of people living close enough to be affected, on pessimistic assumptions.

APPENDIX 1 The specification of risk

1 The word 'risk' is defined in paragraph 11 of the main text as 'the chance that something adverse will happen'. More strictly this means 'the probability that a specified undesirable event will occur in a specified period or as the result of a specified situation'. In this usage, both the probability and the event, and perhaps also its severity, have to be specified. Thus risk can never be reduced to a single quantity; it must always contain at least two separate components. In the context of the risks from a nuclear power station, the risks of greatest interest are those associated with radiation.

The risk to individuals

2 The consequences of an exposure of an individual to radiation have been described in paragraphs 72 to 105 of the main text. These consequences can be grouped into two classes - early effects and late effects. The early effects will occur if the radiation dose is large, much larger than ever results from normal operations, and, as far as the public is involved, larger than those resulting from all but the most serious of accidents. The late effects, of which cancer is the most important, are quite different. No dose, however large, is certain to cause cancer, but any dose results in some additional probability of cancer, the probability rising as the dose increases.

3 Despite the range of possible outcomes, it is possible to use our simple definition of risk by specifying for an individual the probability of each outcome separately. However, this may be unnecessarily complicated and it is common to simplify the picture by considering death, usually fatal cancer, as being an adequate representation of all the possible outcomes. It must be remembered that the risk being considered is that of a death attributable to the radiation exposure under discussion. This use of the probability of an attributable death makes it easy to compare radiation risks with other, more familiar, ones, but it fails to consider time of death.

4 The early deaths caused by radiation can legitimately be compared with deaths due to mechanical injury or acute poisoning, because the deaths occur quite soon after the event that causes them. For the late deaths, the risk is spread very unevenly over several decades and the attributable death is unlikely to occur until ages in the region of 60 to 80 years.

5 If the radiation exposure is spread over a lifetime, or occurs in mid-life, the typical period of life lost due to an attributable death is about 15 years. For early effects, or early deaths from other causes, such as accidents, the period lost is typically about 35 years. Ignoring this difference over-emphasises the importance of the deaths attributable to radiation.

6 There are at least two ways of specifying the risk to an individual in a way which corrects this bias. One way is to give both the probability of the attributable death and an indication of the length of life lost, leaving the reader to make the necessary judgement. The other is to apply a weighting factor to the probability of death in order to allow a direct comparison between the risks of early and late deaths.

7 If the length of life lost is regarded as the most important difference between early and late effects, the weighting factor for late deaths is simply the ratio of the time lost, 15/35, or about 0.43. Alternatively, conventional financial discounting techniques could be applied to the 'cost' of one death compared with another some years later. The discounting period would then be 35 minus 15 years, ie 20 years. With a discount rate in real values of 3% per year, the weighting factor for the late deaths would be about 0.55. At 5% the weighting factor would be about 0.45. The two approaches give similar results, but the first approach is simpler and probably less open to argument.

8 In some situations, part at least of the exposure to radiation will be delayed, sometimes for hundreds or thousands of years. The use of discounting procedures for the costs of effects on health over long periods raises ethical problems. In the regulation and management of nuclear industries, the policies that would follow from discounting the costs of future health effects have been regarded as inappropriate. From the earliest days of the development of civil nuclear power, it has been the policy to provide protection of individuals in the foreseeable future to at least the same standard as is applied now.

Societal risk

9 When considering the risks associated with a plant or an operation, the risk to an individual is not an adequate measure of the total risks: the number of individuals at risk is also important. The presentation of the combined risk to a number of people, sometimes called a societal risk, is very complex. The individuals may be widely dispersed geographically and the risk may extend over many generations. Both routine releases of radioactive materials and accidents come in a wide range of magnitude. The resultant exposures may be to many or to few individuals, each of whom will be at their own level of risk. A single measure of individual risk may be insufficient to express fully the significance of such risks.

10 Ideally, what is needed is a new quantity to represent a combination of the likelihood and the severity of the whole range of adverse outcomes for society that may follow both routine situations and accidents. This quantity has sometimes been called 'detriment'. It can be represented by a table of specified outcomes, each with its own probability. The table may be supplemented by, but preferably not replaced by, an aggregated total detriment. This total is obtained by taking some combination of the severity and the probability of each outcome, weighting each combination to reflect the importance given to that type of outcome, and adding these weighted contributions to give the aggregated total. It is the complete tabulation of detriment that is what we mean by the phrase 'societal risk'.

11 For a routine discharge of radioactive waste, it has become conventional to add together all the resulting doses to all the exposed people, now and in the future. This so called collective dose is a measure of the total detriment from all the late effects attributable to the discharge. The magnitude of the individual doses is still of some interest, so the collective dose is often subdivided into blocks to separate the components due to broad bands of individual dose. A similar subdivision is sometimes used when the exposures stretch far into the future, so that the estimates of parts of the collective dose are subject to considerable uncertainties.

12 A similar technique can be used to assess the severity of accidents once they have occurred. For forward planning, it is also necessary to make allowance for the improbability of accidents. For this purpose, it is necessary to specify both the probability of an accident and the magnitude of its consequences.

13 In this report, the historical risk of accidents has been described by the frequency of occurrence of accidents of a given kind causing more than a specified number of deaths. For example, the worldwide frequency of chemical accidents causing 100 or more deaths is about 0.25 per year, ie an average of about one in four years. For most accidents, nearly all the deaths are early deaths, whereas most of the deaths from a nuclear accident, even from a serious nuclear accident, will be late deaths. These will occur against the background of a much larger number of similar natural deaths in a large population, most of whom will have received only small doses and will thus be subject only to small additional risk. The result will be a small, quite undetectable, increase in cancer incidence. One simplified approach is to define a serious accident as one that is likely to result in about 1000 late deaths and then to estimate the likelihood of such an accident, either at a single plant or anywhere in the country. This is a limited, but still useful, measure of the societal risk of severe accidents. The use of this approach is discussed in paragraphs 164 to 189 of the main text.

APPENDIX 2 Comparisons of risk

1 This appendix compares the size and frequency of various risks we run.

2 Different kinds of risk have to be compared in different ways. Some kinds of risk, such as being killed by lightning or in a road accident or by some other violent cause, are borne by large numbers of people or even by all of us all the time, so it is reasonable to give the chance per million per annum, even though some of us would have a better chance than others. See Table B1.

3 However some kinds of risk need to be compared in a way that takes account of the extent to which the risk is being run. For example, to compare the risk of death from travelling by air, road or rail we need to express it as a proportion of the number of kilometres or the number of journeys travelled. For these kinds of comparison, see Table **82.**

4 Tables 83 and B4 list a number of disasters that have happened in Great Britain and abroad, giving the numbers believed to have been killed. It is not an exhaustive list. Some of the figures are subject to a great deal of uncertainty, particularly where the accident occurred in a developing country, or where many of the deaths did not happen immediately, for example with accidents that reduce life expectancy such as those connected with toxic substances or radioactivity.

5 Table 85 shows the estimated annual chance of certain major events occurring in Great Britain. Estimates of this kind can sometimes be based on direct or historical experience - we know for example how many major fires occur each year and we can expect the same trend to continue more or less. Sometimes, however, the estimates represent no more than a complex set of expert judgements based on a variety of factors such as the known rate of failure of engineering components. Some others, such as the estimated chance of an aircraft crash in Britain killing 500 people, represent a scaling down of world experience. All of them are subject to large margins of error, and those that depend on engineering judgement may be overstated, because of the caution and pessimism which it is customary to build into such estimates.

6 All the tables have to be looked at with certain other points in mind:

(a) not all of us bear the same risk, even, for example, of being killed by lightning. It depends on how much we expose ourselves to the risk; and sometimes there is variation according to age, residence, profession etc;

(b) the tables compare only the chances of **death.** But each kind of hazard carries an additional risk of injury or ill health to people exposed to it, and the extent of this 'tail' will differ from hazard to hazard. For example, no one is known to have been killed by the release of dioxins at Seveso in Italy in 1976; but the health of a large number of people. may have been affected;

(c) in the case of exposure to certain hazards, such as radiation, early death may be unlikely but those exposed may die sooner than they otherwise would. It is only within limits reasonable to compare the number of deaths so resulting with, say, the number dying immediately in an aircraft crash. Strictly speaking in such cases we should compare the actual loss of life expectations. So great caution must be used in comparing the various hazards and events listed in the tables;

(d) where events are infrequent, estimates of the risk may rest upon small numbers of incidents and be affected by some particular recent event;

(e) in comparing different kinds of industrial hazard, account must be taken of the paucity of statistics of industrial ill health. This arises from the difficulty of connecting a death, eg from cancer, with particular conditions encountered at work, and from the fact that the importance of such causes and connections may only be identified long after the event: eg asbestos, which has led to the early deaths of many thousands of people who in earlier years encountered it at work.

7 Notwithstanding these important reservations, the tables give some idea of how the different risks we run compare with each other in size and probability.

Table B1 Some risks of death expressed as annual experiences

	Risk as annual experiences	*Risk as annual experiences per million*	*Basis of risk*
Dying from all causes			
Average over entire population	1 in 87	11 490	UK 1989
Men aged 55-64	1 in 65	15 280	UK 1989
Women aged 55-64	1 in 110	9 060	UK 1989
Men aged 35-44	1 in 578	1 730	UK 1989
Women aged 35-44	1 in 873	1 145	UK 1989
BOYS 5-14	1 in 4400	225	UK 1989
Girls 5-14	1 in 6250	160	UK 1989
Dying from cancer			
average over entire population)	1 in 374	2 880	GB 1989
Death by all violent causes (accidents, homicides, suicides, others) (averaged over population)	1 in 2700	365	GB 1989
Death by accidents (all)	1 in 4200	240	GB 1989
Death by road accidents			
(averaged over population)	1 in 10 204	98	GB 1989
Death by gas incident (fire, explosion or carbon monoxide poisoning, averaged over population)	1 in 1 100 000	0.9	GB 1986 to 1990 average
Death by lightning	1 in 10m	0.1	UK (average over several years)
Death by industrial accident to employees			
Deep sea fishermen on vessels registered in UK	1 in 750	1 340	UK 1990
Extraction of mineral oil and gas	1 in 990	1 011	GB 1986/7 to 1990/1 provisional average
Extraction of minerals and ores	1 in 3 900	254	
Coal extraction	1 in 7 100	141	
Construction	1 in 10 200	98	
Agriculture	1 in 13 500	74	
All manufacturing industry, including:	1 in 53 000	19	
Metal manufacturing industry	1 in 17 000	60	
Instrument engineering industry	1 in 1 million	1	GB one death only in 10 years to April 1992
All service industries	1 in 150 000	6.6	GB 1986/7 to 1990/1 provisonal average

Sources

Annual abstracts of statistics 1991 HMSO 1991 ISBN 0 11 620446 X

Health and Safety Commission *Annual Report 1990/91* HMSO 1991 ISBN 0 11 885726 6

Department of Transport Marine Accident Investigation Branch Annual Report 1990 HMSO 1991
ISBN 0 11 55104 9

Table B2 Some risk of death expressed as consequence of an activity

	Risk expressed a consequences of activity	*Risk expressed as consequence of each million units of activity*	*Basis of risk*
Pregnancy and associated conditions	1 in 13 000 per birth live	77 per million births	UK 1989
Surgical anaesthesia	1 in 25 000 per case	40 per million cases	England and Wales 1970-73
Vaccination	1 in 1m per case	1 per million cases	England and Wales 1967-76
Rock climbing*	1 in 25 000 per hour	0.04 per thousand participant hours	UK 1961
Canoeing	1 in 100 000 per hour	0.01 per thousand participant hours	UK 1960-62
Hang-gliding	1 in 670 per year	1.5 per thousand participant years. (without allowance for time spent in UK activity)	1977-79
Driving by car (drivers and passengers)	1 in 200 per million km	0.005 per million km	GB 1989
Flying, UK scheduled airlines, passengers	1 in 5 000 per million km	0.0002 per million km	UK 1985-89 average.
Rail travel, passengers	1 in 645 per million km	0.00155 per million km	GB 1986-89 average

Sources

Transport and pregnancy risks: annual abstract of statistics HMSO 1991.

Medical and sports risks: quoted from section 4.6 of the report of the Royal Society study group on risk assessment London 1983

Note

* No longer applicable to general rock climbing, in view of technological advances, but still applicable to solo rock climbing.

Table B3 Some man-made disasters that have happened in Great Britain or on UK planes and boats due to accidents

Place	*Date*	*Numbers killed (sometimes approx.)*	*Comments*
Titanic	1912	1513	Collision of liner with iceberg. Insufficient lifeboats.
Senghenydd Colliery	1913	440	Electric circuit caused coal dust and methane gas explosion.
Silvertown. London	1917	73 contemporary (300*)	TNT factory explosion in wartime conditions. Contemporary official figures were quoted in later literature.
Gresford Colliery	1934	265	Methane gas explosion in colliery.
Bolton Wanderers	1946	33	Overcrowding at football ground.
Harrow train crash	1952	112	Collision between express train and stationary commuter train. A third train then collided.
London smog	1952	400+	Protracted smog trapping coal fire etc fumes. Most of those who died already suffered from respiratory or cardiac diseases.
Windscale	1957	0 immediate. Up to approx. 100 long term (estimated)	Fire in graphite moderated reactor.
Aberfan	1966	144	Collapse of coal-mine waste tip on neighbouring school. 116 of those killed were children.
Glasgow Rangers	1971	66	Spectators crushed as crowd surged back into ground attracted by late goal.
Heathrow Airport	1972	118	Crash of Trident on take-off.
Flixborough	1974	28	Explosion following escape of gas from chemical plant modified from original design. Occurred weekend. Had it been during normal working hours, many more casualties would have occurred.
M6	1985	13	Multiple vehicle crash in fog.
Manchester Airport	1985	56	Aircraft fire on take-off.
Bradford	1985	56	Fire in crowded football stadium.

Sumburgh, Shetland	1986	45	Helicopter crashed into sea.
Herald of Free Enterprise	1987	184	Zeebrugge-Dover ferry capsized. Put to sea with vehicle access doors left open.
Kings Cross Station	1987	31	Fire in underground railway station.
Clapham Junction	1988	34	Signal failure due to defective wiring resulted in multiple train crash.
Piper Alpha	1988	167	Explosion destroyed offshore oil platform.
Kegworth	1989	47	British Midlands airliner crashed on M1 at Kegworth while attempting emergency landing at East Midlands Airport.
Marchioness	1989	51	Thames River pleasure boat sank after collision with dredger 'Bow Belle'.
Hillsborough Stadium	1989	95	Failure of crowd control arrangements of football stadium.

Notes

* The higher figure is quoted in Nash J R *Darkest hours* 1979

Numbers quoted are sometimes approximate as many of the estimates of casualties are in the form 'around x deaths or more'. Alternative estimates of casualty numbers from different and often secondary sources in brackets.

'Immediate deaths' include those who died shortly afterwards from injuries, or the immediate effects of radiation, in contrast to those dying much later from long term illnesses.

TABLE B4 Some man-made disasters that have happened abroad

Place	*Date*	*Number killed (sometimes appropriate)*	*Comments*
Johnston, Pennsylvnnia, USA	1889	2000†	Dam burst.
Halifax, Canada	1917	2000(1650)	Munitions ship explosion in harbour.
Oppau, Germany	1921	400+(561†)	Explosion of 4500 tons ammonium nitrate.
Zarnesti, Rumania	1939	60	Escape of 24 tons chlorine.
Honkeiko, China	1942	1572†	Coal dust explosion in colliery.
Cleveland Ohio, USA	1944	128 (131†)	Liquefied natural gas tank ruptured. Flowing boiling liquid engulfed workers and nearby urban area.
Bombay, India	1944	1000	Explosion on ship in harbour, carrying munitions.
Texas City, USA	1947	500+ (576†)	Explosion of 3000 tons ammonium nitrate in cargo of ship. Oil tanks ignited. 330 homes damaged.
Ludwigshafen, Germany	1948	209	Release of flammable vapour from overfilled rail tanker. Explosion caused devastation in surrounding chemical plant and collapse of buildings.
Frejus,	1959	400	Dam foundation failed.
Vaiont, Italy	1963	2000 (2600-3000‡)	Hillside collapsed into reservoir. Dam overtopped.
Potchefstroom, S Africa	1973	18	Ammonia tank burst; sudden failure.
Seveso, Italy	1976	0 immediate Long term effects suspected.	Release of dioxide produced in runaway chemical reaction. Suburb, evacuated, long term health effects being monitored. Major decontamination task lasting several years.
San Carlos, Spain	1978	200+ (215†)	Propylene fire in holiday camp.
Three Mile Island, Pennsylvania, USA	1979	0 immediate One long term death (original estimate)	Minor reactor fault, followed by a series of misdiagnoses by staff led to loss of coolant accident and major core damage. Most of radioactivity successfully contained. Reactor written off.

Mexico City	1984	500+ (452[†])	Liquefied petroleum gas explosion at refinery
Bhopal, India	1984	2000 (2500-5000[†])	Accidental release of methyl isocyanate in pesticides factory. Long term effects feared for many more. Over 170000 people received treatment.
Stava, Italy	1985	250	Dam failure.
Chernobyl, USSR	1986	31 immediate 30,000 long term deaths worldwide over the next 70 years or so (United Nations estimate)	Prompt criticality in nuclear reactor carrying out improperly authorised experiment.
Phillipines	1987	4386	Capsize of ferry Dona Paz.
Islamabad, Pakistan	1988	10+	Fire in ammunition dump on city outskirts. Local residential areas showered with rockets, shells and shrapnel. Over 100 injured.
Eastern Coast, Canada	1988	29	13 year old petrol tanker exploded. All crew lost.
Deesa, India	1988	22	Industrial acid tanker overturned and spilled load. Many of the victims burned as they rushed to salvage tanker contents mistaken for fuel oil.
Arzomas, USSR	1988	73	3 box crates of industrial explosives detonated as train approached station. 230 injured and 150 houses destroyed.
Shanghai, China	1988	25	LPG leak at oil refinery ignited by sparks from nearby construction shed causing explosion and fire.
Bombay India	1988	35	Naphtha overflowing from storage tank with faulty level gauge ignited. Resulting fire engulfed nearby genzene/toluene tanks.
Asha-Ufa, USSR	1989	500+	Leaking LPG pipeline. Gas cloud ignited by spark from passing train.
Garom Chasma, Pakistan	1989	40	Explosion in ammunition depot showered nearby homes with rockets and shrapnel.

Pasadena, USA	1989	23	Vapour cloud release from polyethylene plant ignited destroying plant.
Henan Province, China	1989	27	Explosion at illegal fireworks factory in populated area. 45 houses destroyed, 176 others damaged.
Nagothane, India	1990	31	Explosion at gas cracking plant resulting from leaking pipe.
Bangkok, Thailand	1990	63	Reckless driving led to overturning of LPG tanker in busy city traffic and release of LPG which ignited.
Maphrao, Thailand	1991	171	Truck carrying dynamite overturned. Villagers looting truck and wreckage killed in blast thought to be caused by cigarette.
Red Sea Coast, Egypt	1991	464	Capsize of ferry *Salem Express.*
Borneo	1991	17	Tanker carrying ammonia collided with general cargo ship and sank.
Strasbourg, France	1992	87	Airbus A320 airliner crashed into mountain.
Eastern Turkey	1992	400+	Methane explosion in coal mine.

Notes

* Quoted in Rowe W D *An anatomy of risk,* Wiley New York 1979.

† Quoted in Hohenemser C *Penultimate risks* 53 Clark University Worcester Mass 1989.

‡ Quoted in Inhaber H *Risk of production* AEGB-I119/Rev 1 Atomic Energy Control Board Ottawa 1978.

Numbers quoted are sometimes approximate as many of the estimates of casualties are in the form 'around x deaths or more'. Alternative estimates of casualty numbers from different and often secondary sources in brackets.

'Immediate deaths' include those who died shortly afterwards from injuries, or the immediate effects of radiation, in contrast to those dying much later from long term illnesses.

Table B5 **Major events occurring or estimated per year in Great Britain**

	Event	*Approx chance per annum*	*Basis*
(a)	A fire killing 10 or more people	1	Experience
(b)	A railway accident killing or seriously injuring 100 or more people	1 in 15 or 20	Experience of last 40 years
(c)	An aircraft accident killing 500 people	1 in 1000	Very limited world experience, scaled
(d)	A tidal surge too large for Thames Barrier to control	1 in 1000	Greater London Council design specification
(e)	Event at Canvey Island complex	1 in 5000	Expert estimates of risks causing 1500 + deaths or serious injuries following improvements. 'Conservative' ie likely to over-estimate, rather than under-estimate
(f)	Similar event causing 18000+ deaths or serious injuries	1 in 100 000	
(g)	Aeroplane crashing into any one of London's many football stadia whilst empty	1 in a million	Based on pattern of actual crashes in Home Counties
(h)	Aeroplane crashing into full football stadium	1 in a hundred million	

Sources

(a) Griffiths R F and Fryer L S *The incidence of multiple fatality accidents in the UK* UKAEA Safety and Reliability Directorate Report SRD110 1978

(b) Railway Inspectorate.

(c) Proof of evidence by Mr J Locke to the Sizewell Inquiry, 1984. The Range of Risks from a PWR at Sizewell - An Overview NII/S/92/Saf.

(d) Horner R W The Thames barrier project *Geographical Journal* 145(2) 242-253 (1979).

(e) and (f) Health and Safety Executive *Canvey: a second report,* HMSO 1981.

(g) and (h) Quoted in chapter 6 of the *Royal Society study group on risk assessment,* 1983. More recent figures given by D W Phillips lead to a similar estimate (SRD/HSE/R435. *Criteria for the rapid assessment of the aircraft crash rate onto major hazards according to their location* 1987).

Notes

These figures depend critically on the period of observation, if only one or two events are involved, or on the nature of expert estimates. For any particular hazard the societal risk is best shown as a graph - the observed or predicted chance of an event killing N or more people. Some of the figures in the table are selected from such graphs.

On average there is approximately one fire killing ten or more people each year in this country. Each of these deaths is included in the statistics in Table B5. But because such multiple deaths cause particular concern, it becomes useful to present such figures separately as a 'societal risk', local or national. The next figure, for railway accidents killing or seriously injuring 100 or more people, is based on historical rather than immediate experience. The GB frequency is about one every 20 years taken over an extended time span, although there has not been an accident of this size during the past 30 years. The third figure, for the likelihood of an aircraft accident killing 500 people, is more tenuous. It is derived by scaling, to UK traffic levels, the very limited world experience. There have been very few such accidents anywhere. Worldwide experience may reflect different safety standards from airlines operating in the UK.

The tidal surge too large for the Thames Barrier to control does not of course come from direct experience. It was the design specification by the GLC after considering past experience of flood frequencies and expert predictions. Engineers designing the barrier used their professional judgement and experience to meet this specification.

The Canvey Island predictions represent a complex set of expert judgements, based on experience of chemical accidents and component failures at home and abroad and appropriate judgements of local conditions. They rest in part upon evidence of the effect of dangerous substances upon people - and this will be subject to some variation between (say) the old and very young, and others. The figure is specific to the region around Canvey. The total GB national societal risk from petrochemical complexes would be the sum of various separate and different estimates. The 'aeroplane crash into football stadium' figures are derived from the observed frequency of crashes in this country of planes of all sizes, and the area of land covered by a typical football stadium.

Apart from deaths in fires, these figures are expert estimates based on limited historical accident or event data. Hence they are not hard and fast figures as in a balance sheet. As with all forecasts they contain uncertainties which vary in size depending upon the quality and quantity of the accident data and the nature and extent of the expert assessments and judgements involved. Some of these figures may err on the side of caution, for safety's sake, by a factor of up to ten. They are useful guides to policy making, but their limitations must be made clear to the policy maker.

APPENDIX 3 The application of cost benefit analysis to nuclear safety assessment

1 The Health and Safety at Work Act and other safety legislation requires that safety measures should be taken 'so far as is reasonably practicable'; legally speaking, this means that the risk of harm has to be balanced against the cost of preventive measures, and the latter need not be taken if the cost is grossly disproportionate to the risks. Some risks, of course, are so large that preventive measures would be taken whatever the cost, or the risk not run.

2 In the case of nuclear installations, the potential effects of accidents could be widespread and long lasting, and thus very large sums of money must be spent if necessary to make the chance of an accident very remote, and to reduce harm from radiation in normal operation to low levels.

3 Nevertheless it is necessary for those concerned with the assessment and regulation of the design of nuclear plant to bear the costs of their requirements in mind, and once they are satisfied that risks have been reduced to very low levels, not to insist on further measures whose cost would be out of proportion to the remaining risk.

4 One technique that may be applied in deciding whether a particular requirement is necessary is 'Cost-Benefit Analysis' (CBA). The Layfield Report[1] recommended that assessment levels for radiation doses in normal operations and for the risk of accidents should be founded, so far as practicable, on cost benefit analysis and that HSE should develop the application of CBA to assessment of nuclear safety measures. The following paragraphs set out in simple terms what is involved, and some associated difficulties.

Applying CBA : General issues

5 CBA involves the **identification** and **quantification,** in common (and hence usually monetary) units, of all the desirable and undesirable consequences of a particular measure to society as a whole. A CBA seeks in principle to ascertain whether the benefits of a particular measure are sufficiently great to enable those who gain as a result to be able to compensate those made worse off and still be better off than before. CBA was designed for appraisal of public sector projects. It is not, however, the concern of the safety regulator whether the overall benefits of, say, a nuclear power station outweigh the costs but with whether the additional benefits of making the plant safer justify the additional costs involved. Thus, if, as the safety regulator, we are to make use of CBA, we must be able to isolate the effects of safety related changes.

6 It will usually be the case that those parts of society bearing the costs of a particular measure are not the same as those who receive the benefits. CBA is not concerned with the distribution effects of a project (ie with who gains and who loses) but these effects may be very important to the decision maker. The French power utility (EdF) ensures that the local community by a nuclear power station are effectively compensated for bearing the risk by providing reduced price electricity and subsidised local amenities.

7 The use of CBA in this area necessarily involves attaching monetary values, or at least approximate values, to human life, health and all the other possible consequences of a major nuclear accident. Furthermore, since some of the consequences for human health and the environment could affect future as well as present generations, the application of conventional 'discounting' procedures to obtain equivalent present values is problematic.

8 Finally the process of deciding whether a benefit is sufficient to justify a cost must incorporate an assessment of what represents 'gross disproportion' within the meaning of the law (referred to in paragraph 32 of the main text).

Isolating the costs of safety measures

9 It is very often possible, though never very easy, for a regulatory body to estimate roughly the **cost** of a particular **'extra'** requirement. There are nevertheless at least two difficulties. First, the regulatory body will always be partly dependent on the industry concerned for its estimate of the cost,

which may be very expensive or difficult to check; and the estimate may well be exaggerated. However, regulators can turn to specialist independent advice.

10 Second, there may be not one but many ways of meeting a requirement, some of which may depend on the ability and willingness of the industry to make rearrangements. So isolating the 'additional' element whose cost is to be determined may be a very lengthy and complex business, involving disputable issues as well as questions well outside the regulator's interest or specialist knowledge. Indeed, in the medium or longer term the cost of a particular measure is often untraceable in the flow of new investment. Equally, particular items of plant may serve both safety and commercial purposes at the same time, since there is always a demand for reliability and robustness for purely commercial purposes. It is, in principle, possible in such cases to estimate such commercial benefits and subtract them from the overall estimated additional cost to obtain the net cost of the safety gain. A more intractable problem is that to assess the impact of improving a particular design requirement it is necessary to compare the costs of a plant conforming to the required standards with that for a plant that would satisfy existing standards. This presumes that the information required for both designs is available at the appraisal stage. In practice it may only be possible to provide such information after designing and, possibly only, after (at least partially) constructing the plant. This problem often arises with conventional use of CBA in project appraisal and, while it cannot be avoided, the lessons from previous post project evaluations can help minimise the problem.

11 The cost benefit technique is easier to apply where a particular item serves only a safety purpose, and is clearly additional to the rest. Even where a safety measure can be isolated in this way, the structure of the cost may itself be complex and hard to determine accurately. It would include for example elements for the design, testing and construction of the equipment; its installation and its operation including the auxiliary costs of maintenance, training, instruction etc. These factors are rarely calculated under the same budget heads, so an exhaustive examination would be necessary to achieve accuracy. Nevertheless, the regulator needs to be aware of the broad magnitude of expenditure involved and the very exercise of cost examination can be of value both to the designer and operator.

Quantifying and valuing safety benefits

12 The use of CBA to assess nuclear safety measures necessarily involves quantifying the change in potential **detriment** associated with the safety measures in question. There have been considerable advances in the valuation of human safety and environmental detriments in recent years, although there is understandable scepticism about attaching monetary values to human life and health and to irreversible environmental consequences.

13 No amount of money could compensate someone for the loss of their life. The traditional approach to valuing life which equated a life to the value of a person's future of stream of economic output plus a notional sum for the 'pain, grief and suffering' felt by those affected by a given person's death is now recognised as an inappropriate method. However people do accept small additional risks of death or other harm to themselves in return for financial or other benefits. Thus a value can be inferred for a small additional increment of risk. From this, statistical life values can be derived - ie an expected loss of life from among a large population at risk can be valued in advance even though actual identified deaths that may result cannot.

14 There are two techniques which can and have been applied to elicit monetary values for 'statistical life'. These are:

(a) **seeing what people spend** in other situations to reduce the risk they face, or accept by means of financial compensation for small measures in risk, controlling for other factors. This method is called the 'revealed preference' approach;

(b) **asking people** for the amount they would spend/accept for a reduction/increase in their risk of loss of life in some hypothetical situation. This is called the 'stated preference' or 'expressed preference' method.

15 There have now been numerous results obtained from both methods, both in Great Britain and abroad. These results vary quite markedly and there are problems of interpretation of research findings associated with both methods. The Department of Transport in 1987 commissioned an extensive review of such studies[13]. This literature survey was followed by the issue of a consultation paper[14] proposing a value for a statistical life for use in the Department's appraisal of road schemes of £500 000 (in 1987 prices), to be updated in line with increases in per capital national income.

16 The Department of Transport's proposed value for a life was pitched at the very bottom of the range of values for a life suggested by individual studies. This reflected the Department's concern not to make too large a shift from previous values which had enjoyed governmental and public acceptance at the time.

17 The Department of Transport's consultation exercise secured widespread endorsement for its proposed value of life for application in road transport appraisal. This value now stands at £660 000 for a life. The literature on willingness to pay for changes in risk suggests that the value to be applied to a given risk reduction increases with the level of risk and varies also with the nature of risk. The various factors involved have been identified by research both in the USA and elsewhere into people's perceptions of and attitudes to risk (eg by Slovic, Fischoff and Lichtenstein) and by studies of risk decision (eg in HSE's QRA and its input into decision making[16]). Amongst these are - whether the risk is voluntary or involuntary; the degree of benefit to those at risk from the activity; perception of the level of personal control of the risk; whether the risk is new or familiar; whether it is purely individual or also has a societal dimension; and whether it involves painful or 'dread' forms of death. It would thus appear reasonable to take the Department of Transport's £660 000 value for preventing a road accident fatality as providing the **minimum value** for the loss of a life in an individual accident situation, with higher values applying in some other risk situations. Further work will however be necessary to establish appropriate quantitative weightings for other risk situations and this is presently being explored with other interested departments, including the Department of Transport.

18 Variations of the same basic 'willingness to pay' methodology have been applied to obtain monetary values for various other ill health effects short of immediate death. A variety of 'relative utility loss indices' have been developed, both in the UK and abroad, which use market survey research among selected groups to map the relative disutility of different states of injury and disability in comparison with normal health. These can be used with an established value of life to obtain a money value for any given injury or ill health effect. This approach has certain limitations. It relies on the judgement of experts to analyse injury states to isolate the various dimensions used to measure loss of utility (eg the degree of physical immobility and the level of distress). It is also concerned with actual levels of harm rather than with the risk of different forms of harm. One variant, applied in a recent research project for the Department of Transport, involved the use of questions seeking the chance of successful recovery leading to normal health or failure leading to death that accident victims would be prepared to take as the price of treatment of an otherwise permanent level of injury. This introduces an element of risk, but again in the context of responding to actual harm suffered rather than in the context of a risk of harm.

19 The Department of Transport expect to publish the results of this research and a related project reviewing the relative utility loss approach in the context of arriving at monetary values for a serious road accident injury later in 1992.

Quantifying and valuing other detriments

20 There has been a considerable amount of recent work applying the same willingness to pay principles to obtain monetary values for environmental consequences. Most of these studies have been concerned with valuing the loss of localised environmental amenities. While the values obtained from such studies have little direct application to valuing the environmental risk associated with nuclear power, they provide evidence that a workable methodology exists for deriving appropriate values for certain environmental effects, although whether the methodology can be applied to very major, irreversible environmental losses has yet to be fully tested.

21 Many of the other potential nuclear hazards, such as damage and loss to plant and property and loss of land can be handled through more conventional valuation techniques since there are established market prices for these. There are however a set of potential social and political effects, eg loss of confidence in the Government, demands for the ending of all applications of nuclear technology, which are difficult, if not impossible, to fully identify let alone quantify and translate into monetary values.

Multi-attribute analysis

22 While cost benefit analysis attempts to express all the effects of a decision in monetary terms, 'multi-attribute utility analysis' is a technique designed to encompass factors, such as aversion to low probability accidents and socio-political aspects, which are difficult to quantify in monetary terms. The essence of the technique is to use a scoring scheme for the relevant factors based on the judgement of a group of informed people. The weightings that are produced are thus necessarily subjective.

23 A way of overcoming this is to use a 'multi-criteria outranking' technique where each option is compared to every other option to see whether that option outranks (or is preferred to) other options.

Comparing costs and benefits over time

24 It will be seen from this discussion that the valuation of benefit to be balanced against the cost of any measure is a difficult and uncertain business. Moreover, to compare the costs and benefits which accrue at different times in the future it is necessary to convert them all to present values through 'discounting'. The application of a discount rate reflecting commercial risk, the opportunity cost of capital and general social time preference (ie for jam today rather than a little more jam tomorrow) is an accepted, and largely uncontroversial practice in conventional project appraisal. Its use in nuclear safety appraisal, however, raises some objections. Even if the principle of discounting is accepted, the rate of discount applied to the appraisal of a nuclear power plant may not be the most appropriate rate for appraisal of additional safety measures.

25 The application of even a pure time preference discount rate of 3 to 4% suggests that the benefit of averting a death delayed by some 20 years is worth only about half that of averting an immediate death. There are some who object to this, believing that the benefit of saving a life in the future should have the same value as saving a life today. However, comparing the average years of loss of life per attributable death for someone exposed to an annual occupation risk of immediate death of 1 in 1000 with someone exposed to an equivalent 1 in 1000 risk of developing a fatal cancer, one notes that the former results in more than twice the number of years of life being lost. The relative value of a delayed death against an immediate death produced by discounting seems not unreasonable in the context of comparative risk to those alive today. After all, we all face a certainty of dying, we can only change the date and kind of our death.

26 The application of discounting to very long term effects upon future generations as yet unborn is more questionable. There is a case on equity grounds that one should not value such effects any less than effects upon present members of society. However, while some consequences of a serious nuclear accident would be extremely long lived and this problem also applies to the costs of handling nuclear waste, the life of most of the costs and benefits of nuclear safety measures will be limited to the life of the plant, perhaps 40 years.

Issues of uncertainty

27 The scale of uncertainties involved, in particular those relating to the risk estimates and to the valuation of potential detriments, will usually be so considerable that it is not possible to say with any degree of precision whether the expected costs of a proposed design improvement will exceed the expected safety benefits or not.

28 The use of sensitivity testing can, however, examine the degree of confidence in the assumptions that need to be made if the benefits are to exceed costs and thus help provide at least a qualitative estimate of the likelihood of this.

The ALARP decision rule

29 The process of determining whether a benefit

is sufficient to justify a cost depends on a judgement as to what constitutes 'gross disproportion'. This in turn depends on the prior level of risk. Where this is above the 'broadly acceptable level', 'gross disproportion' essentially takes the form of a multiplier applied to the value of the health and safety benefits and increasing with the level of risk. Precise values for this multiplier have never been defined by the courts and neither the regulator nor the regulated have sought this; both recognise the drawbacks associated with trying to regulate by means of (arbitrary) numbers. Where there are smooth continuous safety cost functions this framework does not provide sufficient information to decide at what point the additional costs become 'grossly disproportionate' to the extra health and safety benefits. However, in most cases there will be discontinuities in the marginal safety cost function or points where rapidly diminishing marginal returns set in. At such points it will usually be fairly easy to decide, by comparing the marginal costs and benefits of further safety improvements, that any extra expenditure would be excessive relative to the increment in health and safety benefits.

Summary for nuclear safety assessment

Application of cost benefit analysis techniques

30 Since 1982 it has been the policy of the Health and Safety Commission that all new regulatory controls (ie both new regulations and Approved Codes of Practice) should be supported by an assessment of the additional costs and benefits involved. These vary flexibly from case to case and this experience demonstrates that the general CBA approach is helpful in providing yardsticks and helping steer a balance between improvements to health and safety and commercial considerations.

31 Cost Benefit Analysis is very far from being a precise calculation. While HSE has continued to develop the application of CBA methodology to health and safety issues, full quantification is often impractical.

32 Full cost benefit analysis is extremely demanding of information and by conflating all effects into a single measure can both give the impression of spurious precision and present both decisionmakers and members of the public with an impenetrable 'black box'.

Radiation protection measures

33 In the limited case of proposals for additional measures of radiation protection at an operating installation, even accepting the limitations referred to above, the costs and effects of such measures can usually be fairly securely established and there is an accepted scale of 'harm' resulting from the doses received and which would be reduced (see paragraphs 90 to 97 of the main text). The NRPB has given practical advice on these matters[16].

Design assessment

34 However, the general application of the technique to indications given by NII in the course of the assessment of the design of an installation involves a number of very serious difficulties. There are two main difficulties. The first is making sufficiently accurate assessments of the probabilities and consequences associated with accidents, particularly of the low probability/high consequence events. The second is that as the process of assessment is conducted, these indications are not formulated as instructions to take a particular design course; they are frequently put forward in the form of problems that need to be solved, or of refusals to accept as fully satisfactory particular solutions to such problems; and the designer then takes these into account in further thinking. Thus neither the assessor nor the designer can readily disentangle the effects of the safety indications given by the assessor from the influences of the other parameters within which the designer is working; and this, from the point of view of cost assessment, involves a particularly difficult form of the problem referred to in paragraph 10 of this appendix. Because of all these difficulties the NII has so far found only a limited use for quantified CBA in aiding its decision making about what reasonably practicable measures should be incorporated into nuclear power stations to reduce the probability, or to mitigate the consequences, of accidents.

35 Where a requirement is objected to as exceeding the assessment principles applied by NII, or where it can readily be disentangled from other design considerations, the question of the cost and of the advantages would always be jointly considered, and clarified by whatever numerical estimates can readily be brought to bear.

Conclusions

36 HSE will continue to develop and apply cost benefit analysis techniques in a flexible and pragmatic way, quantifying and valuing those effects where the necessary information can be obtained without disproportionate effort and delay, rather than attempting to develop and apply a standard CBA 'rulebook' mechanistically across the board. The ultimate aim in all cases should continue to be to inform and clarify how a final judgement on any safety measure should be made, recognising that a variety of additional factors will usually be involved which cannot be reduced to a single measure.

APPENDIX 4 Consideration of societal risk for certain non-nuclear hazards

1 This appendix summarises:

- the approaches used by HSE in risk criteria for land use planning near major hazards;
- the analysis in the report on the transport of major hazards substances by the Advisory Committee on Dangerous Substances; and
- proposals for offshore safety criteria.

Particular attention is paid to the treatment of societal risk in these cases, to illustrate ways in which this matter has been tackled in practice.

Risk criteria for land use planning

2 HSE has stated criteria which are used when it gives advice on building developments near industrial major hazards such as LPG bulk storage, chlorine storage and large petro-chemical process plant[18]. This drew on the concepts put forward for nuclear installations in *Tolerability of risk from nuclear power stations* (HSE 1988). It also consolidated ideas developed for major hazards over a 10 year period within HSE and incorporated the thinking of the Advisory Committee on Major Hazards which in turn quoted a Royal Society Study Group on Risk Assessment[20].

3 The land use planning criteria document relates to developments which have not yet started and are therefore relatively inexpensive to stop. The individual risk criteria used to define the upper and lower bounds of tolerability relate to the probability of receiving a dangerous dose, or worse, of an effect, such as heat, over-pressure or toxic gas. The 'dangerous dose' has the potential to cause death in susceptible people and severe distress or injury to a majority of the remainder of a typical cross-section of the national population. The upper and lower bounds of such individual risk for land use planning purposes are 10 in a million per year and 1 in a million per year respectively, with an additional lower bound of 1 in 3 million per year for developments with higher proportions of highly susceptible people.

4 In using these criteria, HSE indicates negligible risk for development proposals such as housing where people might well be present most of the time below the lower bound (1 in a million per year). Similarly HSE automatically indicates substantial risk for proposals with a substantial number of people (25 or more) above the upper limit since there would probably be one or more highly susceptible people in such a number.

5 To use the criteria, a risk assessment is done to show the zones of land where people would be subject to individual risks at the levels indicated, if they lived in houses in the zone locations. Then the actual development proposals are identified relative to the zones, to see what HSE advice should be.

6 Developments other than housing may produce quite trivial levels of individual risk, but substantial societal risk. Consider a supermarket situated in a zone where the risk to a hypothetical house resident woula be 1 in 100 000 per year. The individual risk to any particular shopper would be very small. However, the supermarket will be full of people so a major accident would be a disaster. HSE therefore, assumes that where the risk to different types of development is mainly societal such developments may be equated in their significance to various sizes of housing development. The equivalencies used are:

7 Within any type of development those which are very large, (shops over 500 m^2 floor area), or contain populations which are vulnerable, (handicapped, outdoor) or where evacuation may be difficult (multi-storey buildings) are given special consideration. In general HSE would advise against special cases where the housing risk was greater than 1 in a million per year and in many cases would look for such risk levels below 1 in 3 million per year. Special cases might also include large theme parks, tower blocks and hospitals.

8 It should be noted that:

(a) the harm is not the risk of death, but the risk of a 'dangerous dose or worse'. This avoids conceptual and precision problems due to variations in individual sensitivity, but it adds an extra dimension of judgement in comparing this measure with risks of death from other causes.

Table D1 Population equivalences

Housing	*Retail*	*Leisure (Pub Restaurant etc)*	*Hotel etc*
10 houses	100 people	100 people	25 people
30 houses	300 people	300 people	75 people

(Note: 'People' at peak periods)

(b) in developing its approach here HSE assumes, implicitly, an 'average' degree of planning benefit for a project, following the rule that permission should be granted unless there are sound and clear-cut reasons against it. Where the local authority is particularly strongly in favour or otherwise, HSE is willing to discuss the implications for the particular case.

(c) there is some degree of flexibility in the criteria - a middle zone - where the size of the development (and other factors) can be considered. This acknowledges the significance of societal risk. It is very similar to the approach applied by Barnes in the report of the Hinkley Point Inquiry.

9 HSE's approach here does not explicitly consider national societal risk. Each development on its own would add very little to the total national risk (HSE has assessed in detail over 5,000 cases in the last 15 years). However, taken all together there would eventually be a significant addition to national risk if such developments were not controlled.

10 Calculating societal risk values before and after a development may throw extra light on the significance of a development. For example, with a new supermarket it may help to see what difference it makes to local societal risk. There are no societal risk criteria for such cases, but the extra information would be fed into the judgement process. (NB this would be in addition to the normal approach, which is to consider the risk zone for housing on the same site and judge accordingly).

11 From time to time, HSE calculates the risks for an existing installation vis-a-vis an estimated reduction by measures at the installation. This uses the results in a relative rather than absolute mode. It has been very helpful on occasions, for example, to show that the best measure was to protect the operator rather than expensive changes to the hardware.

The major hazards transport study

12 The major hazards transport study[14] is a landmark in the development of quantified risk assessment and criteria for industrial major hazards. As part of the study, the ACDS sub-committee adopted criteria against which to set the results of QRA of hazardous substances on road and rail and in sea ports. The objective was to derive benchmarks for tolerability for existing risks to existing populations against which to test the need and scope for improvements in safety. A brief summary is attempted here and illustrated at Figure D1:

(a) The starting point was the framework set out in the original Tolerability document using the second Canvey Island risk assessment report to determine the point at which societal risks were deemed to be just tolerable. This set an upper limit of societal risk to a local community at 1 in 5000 per year of 500 fatalities. The societal risk criteria chosen applied no specific adjustment for aversion to high numbers of fatalities (a slope for the F/N curve of -1 was applied).

(b) Limit lines based on the starting point would be used as initial benchmarks for judgements at other sea- ports. (The 'Local intolerability line' on Figure D1).

(c) The criteria refer only to the risk of fatality - they make no allowance for such as serious injury, ill health, damage to property or the environment, all of which may be factors in any particular decision on the tolerability of major hazard risks and the reasonable practicability of risk reduction.

(d) In addition the risk per tonne of dangerous substances handled at Canvey was multiplied by the tonnage of dangerous substances handled in Britain as a whole to obtain a criterion which was referred to as the 'national scrutiny level'. Similarly a 'local scrutiny level' may be obtained for a particular locality by multiplying the 'risk per tonne' from Canvey by the tonnage of dangerous substances handled at the locality of interest.

(e) Clearly for Canvey the 'local scrutiny level' is identical to the 'local intolerability level' but for other localities the relationship between the two would depend upon the tonnage of dangerous substances handled compared to Canvey. For localities handling greater quantities of dangerous substances than at Canvey the 'scrutiny level' would be an intolerable risk since it would be at a higher risk than the 'intolerability level'. In such situations only the 'intolerability level' would be a relevant criterion.

(f) The 'national scrutiny level' is at a higher risk level than the 'local intolerability level' or the 'local scrutiny level' at any one port since it relates to the total movement of dangerous substances in Britain as a whole. The sub-committee was unable to decide whether or not the 'national scrutiny level' or any other level constituted a 'national intolerability level'.

(g) If the actual assessed risks (local or national) approach the relevant 'scrutiny' level then although these risks may not be intolerable they will, pro rata with the tonnage of dangerous substances moved, be comparable to the risks at Canvey. Such a situation would merit special study (scrutiny) to ensure all reasonably practicable measures were being taken to reduce the risks.

(h) A 'negligible' level is set with frequency three orders of magnitude below the Canvey limit line and with the same slope.

(i) For road and rail, one test was to use the Canvey limit lines for any particular major route or national trade for a particular substance. This would ensure that no population near such a route was exposed to higher societal risk than the Canvey community.

(j) The road and rail risks were also (separately) aggregated to produce national societal risks. These were compared with the national scrutiny level for ports. In addition, they were compared with the Canvey limit line. Both of these comparisons were applied to test the likely significance of the risk and practicability of safety precautions, noting the rather tenuous links with the initial benchmark.

13 It will be seen that a degree of 'reading-across' (that is, cross-comparison of different hazards) has been involved here. This was felt to be justified by the similarities between the Canvey hazards and those elsewhere in the study. The Canvey work had received a great deal of scrutiny, from detailed examination in public inquiries to Ministerial and Parliamentary debate.

14 The report[22] also accepted that building developments near transport sites where individual risk was relevant might be judged using HSE land planning criteria, but it stated a preference for the societal risk approach outlined above, where possible.

Proposals for criteria for offshore risks

15 It may be of interest to note that HSE is proposing a numerical risk criterion for Offshore Installation safety cases[23]. This follows the recommendation in the Lord Cullen report on Piper Alpha, that a QRA should be done of the risk that the 'temporary safe refuge' (TSR) on an installation might be breached. (The TSR is a means of protecting people against a major fire or explosion). Lord Cullen recommended that all risk criteria should be set by installation operators, subject to scrutiny by HSE, but initially HSE itself should propose a criterion for the TSR.

16 HSE suggests that the frequency of breach should be less than 1 in 1000 per year. This would not necessarily result in the death of people in the TSR, but for the purposes of comparison it is conservative to assume that it will do so. Then for a platform with 100 people on board, this point is shown in Figure D 1 for illustration. Note that this frequency was proposed as a benchmark for its particular purpose (HSC co-ordinated consultations on draft offshore installations safety regulations). Its location on the line shown should not be taken to imply a more detailed or sophisticated usage.

Summary and conclusions

17 HSE's risk criteria for land use planning near major hazards rely heavily on individual risk plus judgement to allow for societal risk aspects.

18 The major hazards transport study develops the concepts of local and national societal risk into workable tools to aid decision-making for the context of the study. There may well be scope for extension of this to other types of hazard.

19 HSE's proposed tolerability limit for offshore installation temporary safe refuges may be expressed in similar units, but this is not to imply too close a parallel.

20 In all this work, there are degrees of 'reading across', and the use of simplified measures for complex multi-dimensional hazards, and assumptions about society's requirements. However, there is little sign of serious challenge (except sometimes from industry on cost, rather than policy, grounds).

Figure D1 Risk criteria developed for major hazards of transport study

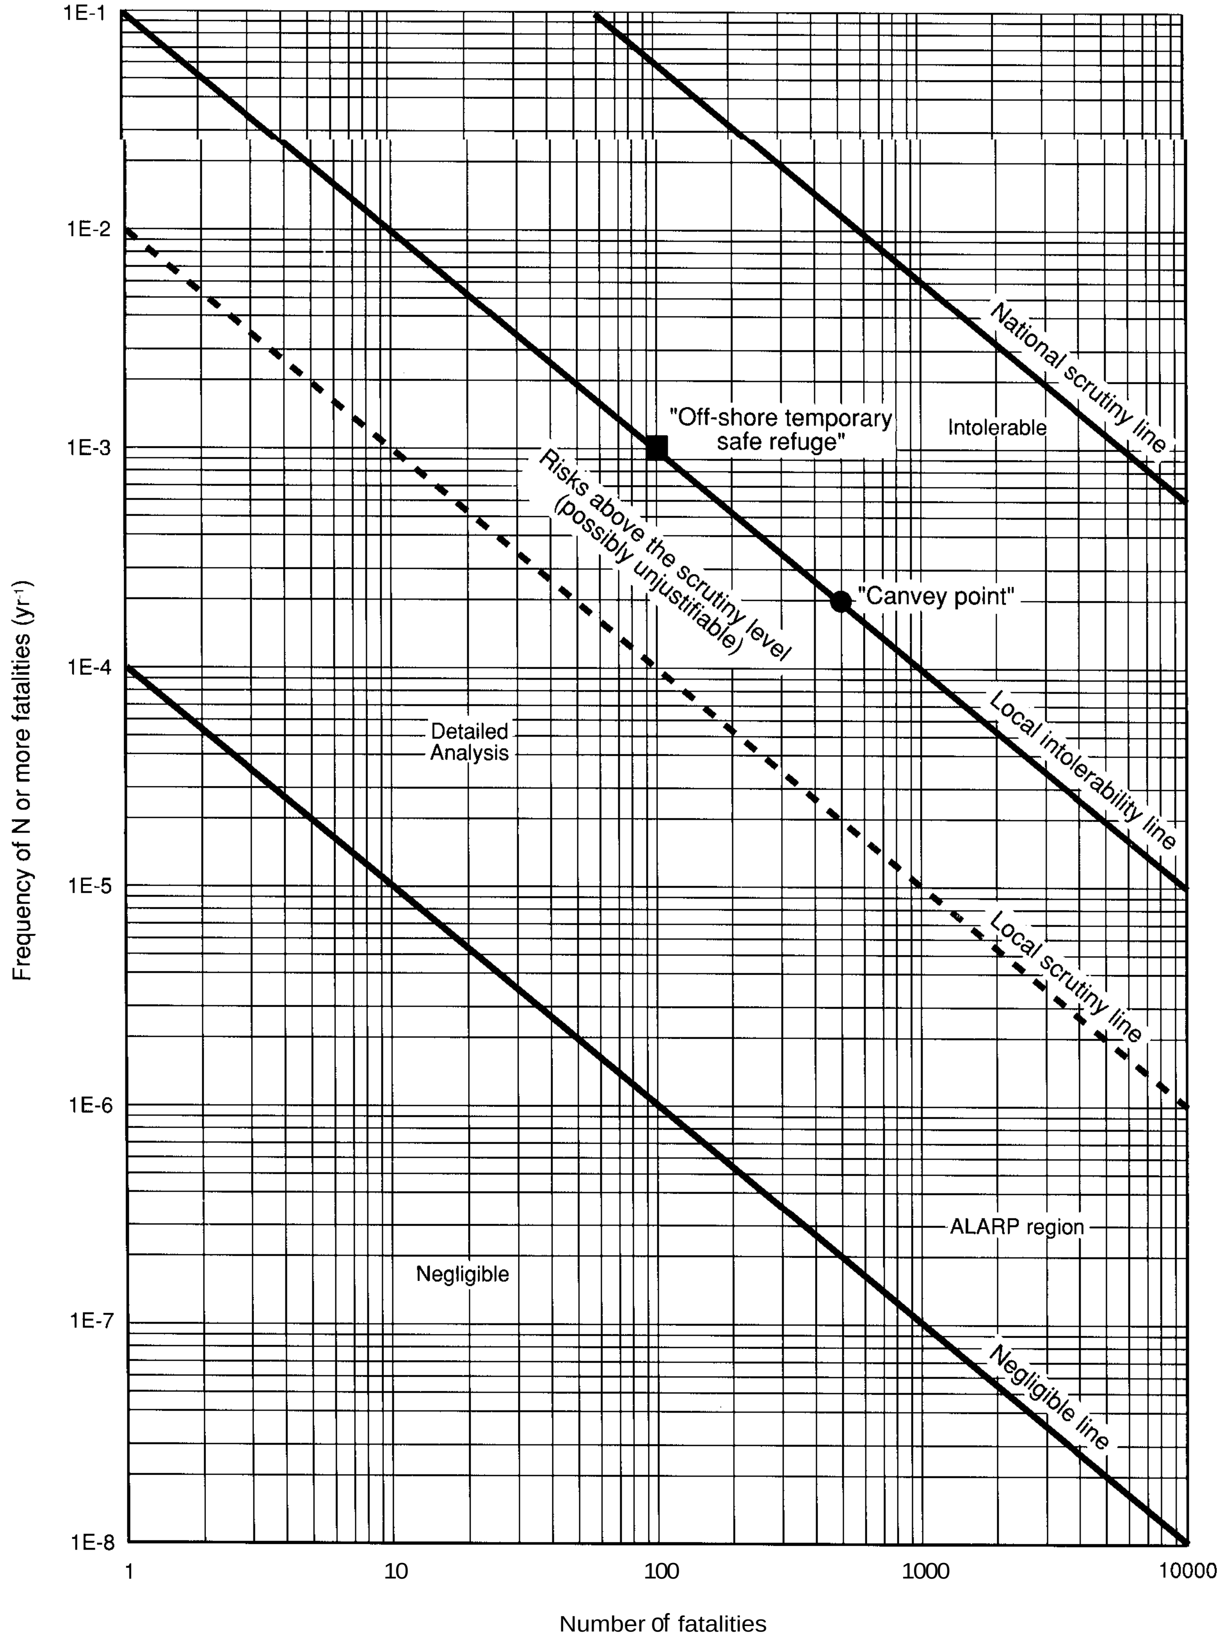

REFERENCES

1 *Sizewell B public inquiry* Report by Sir Frank Layfield HMSO 1987 ISBN 0 11 411576 1

2 *The Hinkley Point public inquiries. A report by Michael Barnes QC* HMSO 1990 ISBN 0 11 412955 X

3 HSE *Safety assessment principles for nuclear reactors* HMSO 1982 ISBN 0 11 883235 2

4 HSE *Safety assessment principles for nuclear chemical plants* HMSO 1983 ISBN 0 7176 0153 6

5 HSE *Safety assessment principles for nuclear plants* revised version to be published late 1992

6 Recommendations of the International Commission on Radiological Protection 1990 (ICRP 60) *Annals of the ICRP* ISBN 0 08041144 4

7 *Ionising Radiations Regulations 1985* SI 1985 No 1333 HMSO 1985 ISBN 0 11 057333 1

8 Council for the European Communities *Council Directive 80/836 as amended by 84/467 Euratom 1984*

9 DOE *Disposal facilities on land for low and intermediate level radioactive waste: principles for the protection* of *the human environment* HMSO 1984 ISBN 0 11 751775 5

10 *Successful health and safety management* HS(G)65 HMSO 1991 ISBN 0 11 885988 9

11 HSE *Arrangements for responding to nuclear emergencies* HMSO 1990 ISBN 0 11 885525 5

12 Kelly GN, Hemming CR and Harbinson SR *Procedures to relate the NII safety assessment principles for nuclear reactors to risk* National Radiological Protection Board R180 HMSO 1985 ISBN 085 951246 7

13 Dalvi MQ *The value of life and safety: a search for a consensus estimate* Department of Transport report 1988

14 Department of Transport *Valuation of road accident fatalities* 1988

15 Slovic P, Feschoff B and Lichtenstein S Facts and fears: understanding perceived risk in *Societal risk assessment: how safe is safe enough?* (ed R Schwing and W Albers) Plenum Press NY 1980

16 HSE *Quantified risk assessment: its input to decision making* HMSO 1989 ISBN 0 11 885499 2

17 National Radiological Protection Board *Cost-benefit advice in the optimisation of radiological pratection* Advice on Standards of Protection Volume 9 HMSO I986 ISBN 085 951 2649

18 HSE *Risk criteria for land-use planning in the vicinity of major industrial hazards* HMSO 1989 ISBN 0 11 885491 7

19 HSC *The control of major hazards* third report of the HSC **Advisory** Committee on Major Hazards, HMSO 1984 ISBN 0 11 883753 2

20 Royal Society *Risk assessment: a study group report* 1983

21 HSC *Major hazard aspects of the transport of dangerous substances* HMSO 1991 ISBN 0 11 885699 5

FURTHER READING

Risk criteria for land use planning in the vicinity of major industrial hazards HMSO ISBN 0 11 885491 7

Major hazard aspects of the transport of dangerous substances HMSO 1991 ISBN 0 11 885699 5

Risk assessment: a study group report The Royal Society 1983 ISBN 0 854032088

The assessment and acceptability of risks

Lord Zuckerman The risks of a no-risk society *The Year Book of World Affairs* 1980 7 Stevens & Sons London 1980 **34** ISBN 0 420457305

Lowrance W W, *Of acceptable risk: science and determination of safety* William Kaufmann: Los Altos California 1976 ISBN 0 913232 31 9
Weinberg A M Science and its limits: the regulators dilemma *Issues in Science and Technology* 1985 11 (1) 59-72

Douglas M and Widawsky A How do we know the risks we face? Why risk selection is a social process *Risk Analysis* 1982 2(2) 45-52

Slovic P Informing and educating the public about risk *Risk Analysis* 1986 6(4) 403-415

Starr C Risk management, assessment and acceptability *Risk Analysis* 1985 5(2) 57-102

Risk assessment: a study group report The Royal Society 1983 ISBN 0 85403 208 8

Dawson J *Living with risk: the British Medical Association guide* John Wiley & Sons 1987 ISBN 0 47191598 X

Cooper M G(ed) *Risk: man-made hazards to man* Oxford University Press 1985 ISBN 0 19 854155 4

Wilson R and Crouch E A C Risk assessment and comparisons: an introduction *Science* 236 267-270
Kleindorfer P R and Kunreuther H C (eds) *Insuring and managing hazardous risks. From Seveso to Bhopal and beyond* Springer Verlag Berlin 1987 ISBN 3 54017732 9

Rimington and Pape *Nuclear power performance and safety. Volume 3; Safety and international co-operation.* International Atomic Energy Agency Vienna 1988 IAEA-CN-48/35. Available from HMSO ISBN 9 20050281

Interpretation of major accidents to the environment for the purpose of the CIMAH Regulations DOE 1991

HSE *A guide to the Control of Industrial Major Accident Hazards Regulations 1984* HS(R)21 (Rev) HMSO 1991 ISBN 0 11 885579 4

Dept of Env Circular 20190. (Welsh Office Circular 34/90). EC Directive on protection of groundwater against pollution caused by certain dangerous substances (80168lEEC)

A guide to hazard and operability studies Chemical Industrial Association 1987

RCEP *Twelfth report best practicable environmental option* HMSO CM3lO 1988 ISBN 0 101031025

RCEP *Fourteenth Report GENHAZ. A system for the critical appraisal of genetically modified organisms to the environment* CM1557 HMSO 1991 ISBN 0101 155727

Emergency exposure indices for industrial chemicals Ecetox Technical Report No 43. Brussels European Chemical Ecology and Technology Centre 1991

Preuss WP and Ehrlich AM *The EPA's risk assessment guidelines* 1990

The use and value of quantitative risk assessment (QRA) and its development within the BP Group GSC Report No 04-90-0005 1990

The development of risk acceptability criteria for the BP Group CSS Report No 04-90-0014 1990

Stringer D A *Hazard assessment of chemical contaminants of soil* Ectox Technical Report No 40 Brussels European Chemical Ecology and Technology Centre 1990

Anderson E *Scientific trends in risk assessment research* Tox Ind Health 1989 5 777-790

Risks from radiation

Pochin E Nuclear radiation: risks and benefits (monographs on science, technology and society) Clarendon Press 1983 ISBN 0 19858329 X

National Radiological Protection Board Living with radiation HMSO 1986 ISBN 0 85951 255 X

National Radiological Protection Board Interim guidance on the implications of recent revisions of risk estimates IRCP 1987 Como statement HMSO 1987 ISBN 0 85951 2932

Dunster H J The biological effects of radiation and the relative risks of conventional energy sources Proceedings of the Royal College of Physicians Edinburgh 19 (2) April 1989.

Nuclear power programme

Royal Commission on Environmental Pollution, Nuclear power and the environment. Sixth Report Cmmd 6618 HMSO 1976 ISBN 0 101 661 800

Pentreath R J Nuclear power, man and the environment Taylor & Francis 1980 ISBN 0851 098401

Fremlin J H Power production what are the risks? Oxford University Press 1987 ISBN 0 19286078 X

Trades Union Congress Nuclear energy: general council report London 1987 ISBN 1850 061 31 9

Recent technical reports

Gittus J N Probabilistic safety analysis for the pressurised water reactor Science and Public Affairs 2 1 21

Training of Technical Staff for Nuclear Power Stations Mechanical Engineering Publications Papers presented at a seminar held at the Institution of Mechanical Engineers on 23 September 1986 Bury St Edmunds 1986 ISBN 0 85298 61 30