HSE
Health & Safety Executive

Good practice and pitfalls in risk assessment

Prepared by **the Health & Safety Laboratory**
for the Health and Safety Executive 2003

RESEARCH REPORT 151

HSE
Health & Safety Executive

Good practice and pitfalls in risk assessment

Dr Sandra Gadd, Dr Deborah Keeley,
Dr Helen Balmforth
Health & Safety Laboratory
Broad Lane
Sheffield
S3 7HQ

This report presents findings from stage 2 of a project conducted by the Health & Safety Laboratory on behalf of the Health & Safety Executive (HSE). The aim of this stage of the project was to identify and highlight common pitfalls in industry risk assessment methodologies and their application, and to briefly summarise how different HSE divisions evaluate risk assessments. The report sets out examples of good practice in relation to risk assessment, and contrasts these with pitfalls encountered by HSE, illustrated by case study examples. The project was jointly funded by HSE's Field Operations and Hazardous Installations Directorates.

This report and the work it describes were funded by the Health and Safety Executive (HSE). Its contents, including any opinions and/or conclusions expressed, are those of the authors alone and do not necessarily reflect HSE policy.

HSE BOOKS

First published 2003

ISBN 0 7176 2732 2

ACKNOWLEDGEMENTS

The authors would like to thank all those in HSE who kindly responded to requests for help in identifying examples of pitfalls encountered by industry in doing risk assessments and provided relevant information. The assistance provided is greatly appreciated. The authors would also like to thank the following individuals who provided helpful comments on drafts of the report: David Bosworth, SPD; Phil Brighton, NSD; David Carter, HID; Kim Fox, RI; Peter Harper, HID; Mary Trainor, HSL; Shane Turner, HSL; Geoff Vaughan, NSD; Ian Whewell, OSD; and Nicola Worsell, HSL.

CONTENTS

1 INTRODUCTION 1

1.1 RISK ASSESSMENT TERMINOLOGY 1

1.2 THE APPLICATION OF RISK ASSESSMENT 2

1.3 REVIEW OF PUBLISHED CRITIQUES OF RISK ASSESSMENT METHODOLOGIES 2

1.4 STRUCTURE OF REPORT 3

2 BACKGROUND - THE LEGISLATIVE FRAMEWORK 4

3 GOOD PRACTICE IN THE USE OF RISK ASSESSMENT, AND ILLUSTRATION OF COMMON PITFALLS 6

3.1 DEFINING THE SCOPE OF THE ASSESSMENT 7

3.1.1 Risk 'of what?' 8

3.1.2 Risk 'to what?' 9

3.1.3 Risk 'from what?' 9

3.2 DECIDING ON AN APPROPRIATE APPROACH AND DETERMINING THE LEVEL OF DETAIL NEEDED 10

3.2.1 Determining the level of detail needed in the assessment 11

3.2.1.1 Qualitative risk assessment 13

3.2.1.2 Semi-quantitative risk assessment 14

3.2.1.3 Quantitative risk assessment 14

3.2.2 Deciding whether it is appropriate to consider individual or societal risk, or both 15

3.2.3 Approaches to estimating individual risk 16

3.2.3.1 Estimating the risk to particular groups of persons 17

3.2.3.2 Estimating the risk in cases of 'transient exposure' - avoiding 'salami slicing' 18

3.2.3.3 Estimating the risk from 'risk hotspots' 18

3.3 DECIDING WHO SHOULD BE INVOLVED 19

3.4 IDENTIFYING THE HAZARDS 21

3.5 IDENTIFYING THE POSSIBLE CONSEQUENCES 24

3.6 ESTIMATING THE LIKELIHOOD OF THE POSSIBLE CONSEQUENCES 26

3.7 ESTIMATING THE RISK 30

3.7.1 Risk indices 31

3.7.2 Individual risk measures 32

3.7.3 Societal risk measures 33
3.8 EVALUATING THE RISK 33
3.9 RECORDING THE FINDINGS 40
3.10 ENSURING TRANSPARENCY THROUGHOUT 41
3.11 ENSURING APPROPRIATE CONSIDERATION OF HUMAN FACTORS 41
3.12 HANDLING UNCERTAINTY 43
3.13 ACTING UPON THE FINDINGS 45
3.14 REVIEWING THE ASSESSMENT 47
4 BRIEF REVIEW OF GUIDANCE FOR HSE INSPECTORS ON RISK ASSESSMENT 48
5 SUMMARY OF MAIN FINDINGS 51
6 REFERENCES 52

EXECUTIVE SUMMARY

This report presents findings from stage 2 of a project conducted by the Health & Safety Laboratory on behalf of the Health & Safety Executive (HSE). The aim of this stage of the project was to identify and highlight common pitfalls in industry risk assessment methodologies and their application, and to briefly summarise how different HSE divisions evaluate risk assessments. The report sets out examples of good practice in relation to risk assessment, and contrasts these with pitfalls encountered by HSE, illustrated by case study examples. The project was jointly funded by HSE's Field Operations and Hazardous Installations Directorates.

Objectives

The objectives of this stage of the project were to:

- Carry out a review of published critiques of both general and specific risk assessment methodologies;
- Identify examples of inadequate industry risk assessments that illustrate common pitfalls in the application of risk assessment;
- Carry out a brief review of HSE guidance for Inspectors on assessing the adequacy of risk assessments.

The scope of this work was restricted to consideration of the assessment and control of risks to people's health and safety arising from work activities, e.g. the assessment of risks to the environment or of health and safety risks arising from non work related activities are not within scope. The methodologies considered and pitfalls identified will therefore be of relevance across the entire spectrum of industry within the United Kingdom, primarily to work activities under the remit of HSE enforcement, but will also be applicable to many areas of Local Authority enforcement.

Main Findings

A comprehensive literature search was carried out to attempt to identify any published critiques of both general and specific risk assessment methodologies, or any references that include descriptions of risk assessment pitfalls. While there exists a large body of published material on the general topic of risk assessment and its application, very few references were found that include material of relevance to the critical review of methodologies or information on risk assessment pitfalls. A brief review of the few identified references was carried out.

A large number of examples of inadequate industry risk assessments that illustrate common pitfalls in the application of risk assessment were identified by collating experience from HSE operational divisions, and from HSL's own experience of carrying out support and research work for HSE. An outline of good practice in the use of risk assessment is presented in the report, and common industry pitfalls are illustrated throughout this section of the report by the inclusion of twenty six case study examples.

The identified pitfalls were as follows:

- Carrying out a risk assessment to attempt to justify a decision that has already been made;
- Using a generic assessment when a site-specific assessment is needed;
- Carrying out a detailed quantified risk assessment without first considering whether any relevant good practice was applicable, or when relevant good practice exists;

- Carrying out a risk assessment using inappropriate good practice;
- Making decisions on the basis of individual risk estimates when societal risk is the appropriate measure;
- Only considering the risk from one activity;
- Dividing the time spent on the hazardous activity between several individuals - the 'salami slicing' approach to risk estimation;
- Not involving a team of people in the assessment or not including employees with practical knowledge of the process/activity being assessed;
- Ineffective use of consultants;
- Failure to identify all hazards associated with a particular activity;
- Failure to fully consider all possible outcomes;
- Inappropriate use of data;
- Inappropriate definition of a representative sample of events;
- Inappropriate use of risk criteria;
- No consideration of ALARP or further measures that could be taken;
- Inappropriate use of cost benefit analysis;
- Using 'Reverse ALARP' arguments (i.e. using cost benefit analysis to attempt to argue that it is acceptable to reduce existing safety standards);
- Not doing anything with the results of the assessment;
- Not linking hazards with risk controls.

A review of HSE guidance for Inspectors was carried out, the most relevant being specific guidance for Inspectors on risk assessment relating to the assessment of safety cases in the nuclear, offshore, railway and chemical sectors. A high level review of this guidance was carried out.

Recommendations

It is hoped that this report will provide useful guidance for Inspectors involved in the assessment of industry risk assessments on the appropriateness of the adopted approaches, and also to practitioners in industry involved in the process of carrying out workplace risk assessments of how to avoid common pitfalls.

1 INTRODUCTION

This report presents findings from stage 2 of a project conducted by the Health & Safety Laboratory (HSL) on behalf of the Health & Safety Executive (HSE). The aim of the first stage of the project was to gain an insight into the extent to which risk assessment has been adopted by industry. This was achieved by carrying out a scoping study looking at the extent of adoption of risk assessment within the polymers and plastics sector of the polymers and fibres industry. The findings from stage 1 of the research were presented in Wright *et al.* (2002).

The aim of this second stage of the project is to identify and highlight common pitfalls in industry risk assessment methodologies and their application, and to briefly summarise how different HSE divisions evaluate risk assessments. The project was jointly funded by HSE's Field Operations and Hazardous Installations Directorates.

The objectives for this stage of the project are to:

- Carry out a review of published critiques of both general and specific risk assessment methodologies;
- Identify examples of inadequate industry risk assessments that illustrate common pitfalls in the application of risk assessment;
- Carry out a brief review of HSE guidance for Inspectors on assessing the adequacy of risk assessments.

This work is concerned with the assessment and control of risks to people's health and safety arising from work activities (for example, the assessment of risks to the environment or of health and safety risks arising from non work related activities are not within scope). The methodologies considered and pitfalls identified will therefore be of relevance across the entire spectrum of industry within the United Kingdom, primarily to work activities under the remit of HSE enforcement, but will also be applicable to many areas of Local Authority enforcement. While the assessment of risks to both health and safety are within the scope of the work, the sections of the report that describe risk assessment methodologies and approaches in detail concentrate primarily on the assessment of risks to safety, as the approaches needed for the assessment of risks to health are often quite different. Only limited discussions of the differences in approach that would be appropriate when considering risks to health have been included in the text. The report sets out examples of good practice in relation to risk assessment, and contrasts these with pitfalls encountered by HSE, illustrated by case study examples.

1.1 RISK ASSESSMENT TERMINOLOGY

There are numerous definitions of risk assessment, and of hazard and risk, both terms that are central to the process, and so it is useful to begin by defining what is meant by these terms within this report.

A *hazard* is any physical situation or object that has the potential to cause harm to people, and *risk* is the likelihood of a specific undesired event occurring within a specified period. Risk is therefore a function of both the likelihood and consequence of a specific hazard being realised.

Risk assessment is the process of estimating the likelihood of occurrence of specific undesirable events (the realisation of identified hazards), and the severity of the harm or damage caused, together with a value judgement concerning the significance of the results. It therefore has two distinct elements: risk estimation and risk evaluation.

1.2 THE APPLICATION OF RISK ASSESSMENT

Risk assessment can be thought of as a tool to aid decision making. The process of carrying out a risk assessment will result in an understanding of the level and significance of workplace risks that should inform decisions relating to the implementation of appropriate risk control and reduction measures.

Risk assessment is an integral part of successful health and safety management. Employers are required to have arrangements in place to cover their management of health and safety (MHSWR, 1999). The effective management of health and safety will depend, amongst other things, on a suitable and sufficient risk assessment being carried out and the findings being used effectively. The arrangements for management of health and safety can be integrated into the management system for all other aspects of the organisation's activities, and will need to reflect the complexity of the organisation's activities and working environment (MHSWR, 1999). The key elements of effective management systems as outlined in HSE (1997) are: policy, organisation, planning and implementation, measuring performance and review.

Risk assessment is an essential component of the 'planning and implementation' element of effective health and safety management. Risk assessment methods can be used to decide on priorities and to set objectives for eliminating hazards and controlling and reducing risks (HSE, 1997).

1.3 REVIEW OF PUBLISHED CRITIQUES OF RISK ASSESSMENT METHODOLOGIES

A comprehensive literature search was carried out to attempt to identify any published critiques of both general and specific risk assessment methodologies, or any references that included descriptions of risk assessment pitfalls. The OSH-ROM literature database (which covers HSELINE, RILOSH, CISDOC, NIOSHTIC and NIOSHTIC-2) was searched, for references including the keywords 'risk assessment' or 'probabilistic safety assessment' in their title, with one or more of the following keywords appearing in the abstract: review; problem; best practice; good practice; comparison.

There exists a large body of published material on the topic of risk assessment and its application. It is recognised that some of this material may include the sought for critiques of methodologies or descriptions of risk assessment pitfalls, without explicitly identifying this in the title or abstract, and would therefore not have been identified by this literature search. However, the volume of general risk assessment literature is so large that a more wide ranging literature search was beyond the scope of this work. Only the following relevant references were identified from the literature search that was carried out: Beale (2001), Simpson (1996) and Stowe (2001), and hence only a brief review of these is included.

Simpson (1996) outlines common problems and misunderstandings raised by risk assessment practitioners, and suggests a modification to the traditional approach to risk assessment to make it more practicable. Further references to this article are made where appropriate throughout the report. Beale (2001) provides an industry perspective on the selection of appropriate

approaches to risk assessment in different situations, selecting five hypothetical scenarios to illustrate how risk assessment might be used appropriately for managing each scenario. Beale (2001) also identifies five common pitfalls that can undermine the usefulness of the risk assessment approach. Stowe (2001) provides guidance to employers (aimed specifically at small and medium sized enterprises) on how to assess and manage workplace risks using practical 'no-frills' techniques that are also cost effective. Practical cases of failure to manage risks successfully are illustrated by a case study approach, along with details of losses incurred and the control measures which, if they had been in place, would have prevented the loss in the first place.

1.4 STRUCTURE OF REPORT

Section 2 of the report includes as relevant background a description of the regulatory requirements for risk assessment. In section 3 of the report, an outline of good practice in the use of risk assessment is given, throughout which common pitfalls are illustrated through identified case study examples. The case study examples of inadequate industry risk assessments were identified by collating experience from HSE operational divisions, and from HSL's own experience from carrying out support and research work for HSE.

Section 4 of the report includes a high level review of HSE internal guidance relevant to assessment of the adequacy of risk assessments, and a brief summary of the main findings of the report is given in section 5.

2 BACKGROUND - THE LEGISLATIVE FRAMEWORK

Sections 2(1) and 3(1) of the Health and Safety at Work etc. Act (HSW, 1974) place a general duty on every employer to conduct their undertaking in such a way as to ensure, so far as is reasonably practicable (SFAIRP), that both employees and persons not in their employment who may be affected thereby are not exposed to risks to their health or safety. HSW (1974) therefore includes an implicit requirement for risk assessment. A number of other regulations, both general and industry sector specific, include explicit requirements for risk assessment. Industry may also undertake risk assessment for other purposes, e.g. insurance.

The Management of Health and Safety at Work Regulations 1999 (MHSWR, 1999) require all employers to assess the risks to employees and any others who may be affected by their undertaking, to enable them to identify measures necessary to comply with their duties under health and safety law. The aim of the risk assessment is to identify the significant risks to health and safety to any person arising out of, or in connection with any work activity. It should identify how the risks arise, and how they impact on those affected. This information is needed so that decisions can be made about how to manage the risks in an informed, rational and structured manner, and the action taken proportionate.

Risk assessments are required by MHSWR (1999) to be 'suitable and sufficient'. This means they should: identify the significant risks arising out of the work activity, and consider all those who may be affected; be appropriate to the nature of the work and be such that they remain valid for a reasonable period of time. The level of detail in a risk assessment should be proportionate to the level of the intrinsic hazards. Once the risks have been assessed and taken into account, insignificant risks can usually be ignored as can risks arising from routine activities associated with everyday life, unless the activity compounds or significantly alters those risks. Employers are expected to take reasonable steps to help themselves identify risks, e.g. by looking at appropriate sources of information, such as relevant legislation, appropriate guidance, supplier manuals and manufacturers' instructions and reading trade press. They should also look at and use relevant examples of good practice from within their industry. The risk assessment should include only what an employer could reasonably be expected to know; they would not be expected to anticipate risks that were not foreseeable.

In addition to the general requirement on all employers to carry out risk assessments imposed by MHSWR (1999) a number of industry sector specific regulations also include requirements for risk assessment. Examples of such industry sector specific legislation include:

- Control of Substances Hazardous to Health Regulations 1999, requiring employers to carry out assessments of the risks from exposure to substances hazardous to health (COSHH, 1999);
- Construction (Design and Management) Regulations 1994, covering construction activities (CDM, 2001);
- Offshore (Safety Case) Regulations 1992, requiring operators of offshore installations to prepare comprehensive safety cases (including detailed risk assessments) detailing their arrangements for health and safety (OSC, 1998);
- Control of Major Accident Hazard Regulations 1999, requiring operators of certain major hazard sites to prepare formal Safety Reports (which implicitly require them to carry out detailed risk assessments) to demonstrate that all measures necessary have been taken to control the off-site risks (COMAH, 1999);

- Railways (Safety Case) Regulations 2000, requiring railway operators to prepare comprehensive safety cases (including detailed risk assessments) detailing their arrangements for ensuring the health and safety of the public (HSE, 2001b);
- Nuclear site licensing arrangements. Under the nuclear site licensing regime, safety cases are required for the design, construction, commissioning, modification and decommissioning of individual plants. The Nuclear Installations Inspectorate (NII) expectations on the scope and detail of the associated risk assessment have been published (HSE, 1992b).

3 GOOD PRACTICE IN THE USE OF RISK ASSESSMENT, AND ILLUSTRATION OF COMMON PITFALLS

In this section, an outline of good practice in the use of risk assessment is given. An outline of good practice at each stage of the risk assessment process is presented, accompanied by case study examples illustrating commonly seen pitfalls. The section begins with a brief overview of why and how risk assessments are carried out.

The purpose of carrying out a risk assessment is to determine whether the level of risk arising from workplace activities is acceptable, or whether more needs to be done to control or reduce the risk. Risk assessment involves both an estimation of the magnitude of the risk (i.e. how big is it?) and an evaluation of the significance of the risk (i.e. is it acceptable?). The process of risk assessment should be carried out in a rational, logical and structured manner.

The findings from a risk assessment can be used to inform decisions as to whether any existing precautions or control measures are adequate, or whether additional prevention or control measures are needed. Risk assessment can also be used to perform a systematic comparison of different risk control/reduction options so that the optimal decision can be made. It is therefore not appropriate to carry out a risk assessment to attempt to justify a decision that has already been made (see Case Study 1).

The process of using risk assessment to inform decisions relating to the control and reduction of workplace risks can be conveniently divided into 3 stages:

- Preparing for the assessment;
- Carrying out the assessment;
- Post-assessment activities.

In practice the distinctions between the 3 stages outlined are not clear cut, but this division provides a useful framework for outlining the guiding principles and factors that should be considered throughout the process.

In preparing for the assessment, the following factors should be considered:

- What is the appropriate scope for the assessment?;
- What is an appropriate approach, and what level of detail is needed?;
- Who is going to be involved?

In carrying out the assessment, the basic steps to be followed are:

- Identify the hazards;
- Identify the possible consequences;
- Estimate the likelihood of the possible consequences;
- Estimate the risk;
- Evaluate the risk;
- Record the findings.

The order presented for the identification of possible consequences and estimation of their likelihood is not intended to be prescriptive; these steps can be completed in either order, or simultaneously, but both are needed in order to estimate the risk.

In carrying out the risk assessment it is also very important to:

- Ensure transparency throughout;
- Ensure appropriate and adequate consideration of human factors*;*
- Ensure adequate handling of uncertainty.

Once the assessment has been completed, it is essential that appropriate action is then taken. It is essential that:

- The findings are acted upon;
- There is a system for regular review of the assessment.

Case Study 1: 'Carrying out a risk assessment to attempt to justify a decision that has already been made'

Projectiles were being thrown from the doors and windows of trains crossing a viaduct, aimed at a car showroom. 28 projectiles were known to have been thrown in 30 months including a headrest, seat and stones. HSE viewed this as an obvious and serious risk of death or injury to members of the public and served an Improvement Notice on the rail operator to erect a fence at an approximate cost of £100k. The duty holders subsequently employed consultants to carry out a Quantified Risk Assessment (QRA), which included a cost benefit analysis (CBA), to attempt to show that the cost involved was grossly disproportionate to the improvement in risk.

Pitfall - The QRA was used to attempt to justify a decision that had already been made, i.e. doing nothing, rather than doing a risk assessment in advance to inform the decision making process.

Other problems with the Risk Assessment (RA) - As there is a clear risk of death or serious injury, QRA is not appropriate in this case, common sense and good industry practice should have been used instead. In addition, the QRA itself was flawed as it was not a cautious best estimate and inappropriate assumptions were used in the CBA (see Section 3.12).

3.1 DEFINING THE SCOPE OF THE ASSESSMENT

The scope of any risk assessment should be clearly defined at the outset. All those involved in the assessment should be agreed on what is being assessed, i.e. risk 'of what?' 'to what?' 'from what?'.

3.1.1 Risk 'of what?'

The first stage in defining the scope of the assessment is to specify the nature of the risk that is being assessed, i.e. to specify the 'level of harm' that is to be estimated and evaluated.

The following definitions of harm are commonly used in health and safety risk assessments: risk of fatality; risk of major injury; risk of minor injury; and risk of receiving a defined level of exposure (e.g. a particular dose of toxic substance, radiation, heat or explosion overpressure). There are difficulties associated with whatever definition of harm is used. The use of injury definitions is problematic since they are difficult to define consistently, and there are problems associated with comparability between different types of injuries (e.g. thermal vs explosion vs toxic effects) (AIChemE, 1989). The use of fatality, although easy to define, has the following difficulties: society is concerned about risks of serious injury and damage to health as well as death; and there are technical difficulties in calculating the risk of death from a hazard to which individual members of the population may have widely differing vulnerabilities.

While it is generally not possible to add the risk from different harm criteria as this does not produce meaningful results, the concepts of 'dangerous dose' and 'equivalent fatalities' have been developed to allow the risk of different levels of harm to be combined in a meaningful way.

The HSE introduced the concept of 'dangerous dose' (HSE, 1989) as an alternative harm criteria to overcome some of the problems associated with using fatality as the harm criteria. The dangerous dose is defined as a dose that has the potential to cause fatality, but will not necessarily do so. The HSE dangerous dose is defined as a dose of toxic gas, or heat, or explosion overpressure that gives all the following effects:

- Severe distress to almost everyone;
- A substantial fraction of the population requires medical attention;
- Some people are seriously injured requiring prolonged treatment;
- Any highly susceptible people might be killed.

The HSE dangerous dose is often equated to a dose that would kill roughly 1% of a 'typical' exposed population.

The concept of 'equivalent fatalities' is widely used within the rail industry (Railway Safety, 2001). In this approach, a fatality is assumed to be equivalent to a specified number of major and minor injuries, so that the risk of fatalities, major and minor injuries can be combined to give the total number of 'equivalent fatalities'. Two different assumptions as to the equivalence of fatalities and injuries are commonly used: 1 fatality = 10 major injuries = 100 minor injuries; and 1 fatality = 10 major injuries = 200 minor injuries.

It is important that whatever definition of harm is to be used in a particular assessment, it is clearly specified at the outset. The choice of an appropriate definition of harm for any particular assessment should be informed by initial consideration of the criteria that are going to be used to judge the acceptability, or otherwise, of the estimated risk. It only makes sense to estimate the risk of a particular level of harm being realised if this will allow a judgement to be made as to whether the risk is acceptable, or whether more needs to be done to control or reduce the risk.

3.1.2 Risk 'to what?'

When considering health and safety risks, risk 'to what' really means risk 'to whom'. In defining the scope of the assessment, it is necessary to define the individuals or groups of people that form the subject of the risk assessment, i.e. that you are interested in estimating the risk to. For example, this may be employees, or particular groups of employees, 'atypical workers' such as contractors or cleaners, members of the public, or particular vulnerable populations.

Many of the regulations that require risk assessments to be carried out effectively define 'target populations' to whom the risks must be assessed, and this may be of assistance in defining an appropriate scope for the assessment. MHSWR (1999) require employers to carry out a suitable and sufficient assessment of the risks to both employees and to persons not in their employment who may be affected as a result of their undertaking. Hence there is a general requirement for workplace risk assessments to consider all those who may be affected as a result of the work activity.

If more than one 'target population' is to be considered, it is possible to carry out a number of separate assessments, rather than considering all populations in a single assessment. For example, it may be appropriate to carry out separate assessments of the risks to particular groups of employees, or to employees and members of the public or to consider on and off-site risks separately. However, if this approach is taken it is important to ensure that no gaps in the assessment process are introduced as a result of the subdivision.

3.1.3 Risk 'from what?'

Risks arise from the realisation of hazards. Work activity hazards can be present as a result of any one or a combination of the following: substances; machines/processes; work organisation, including people and circumstances in which the activities take place, and non-routine (such as emergency or breakdown activities) as well as routine, planned activities. It is important to clearly define which elements of the work activity are included within the scope of any assessment, and to take account of risks arising out of all reasonably foreseeable events and behaviour that are under the control of the duty-holder (HSE, 2001a), or that can have an effect on the duty holder's activities, i.e. that are a reasonably foreseeable source of harm (HSE, 2002a). Whether a reasonably foreseeable, but unlikely, event (such as an earthquake) should be considered depends on the consequences for health and safety of such an event (HSE, 2002a). It is not necessary to take account of risks that are trivial, or arising from routine activities associated with life in general, unless the work activity compounds the risks or there is evidence of significant relevance to the particular work activity.

MHSWR (1999) require employers to carry out an assessment of the risks arising from all work activities, hence there is a general requirement for all hazards to be considered within a risk assessment. However, it is possible to achieve this by carrying out a number of separate assessments that each focus on particular activities, or areas of the workplace, or types of hazards, rather than considering everything within a single assessment. If this approach is taken, it is important to ensure that no gaps in the assessment process are introduced as a result of the subdivision.

An approach that is commonly taken to assessing workplace risks where there are similar activities and hazards across different work activities, areas of the workplace, or at different sites owned by the same company is to carry out a generic risk assessment, considering all

common hazards in a single assessment. While this approach can appear attractive as it is less resource intensive than carrying out a series of site-specific assessments which would require some duplication of effort as common hazards would be assessed in each case, care is needed. For any particular work activity, area of the workplace, or site it is necessary to consider whether all hazards are included in the generic assessment, and also whether the circumstances are such that the generic assessment is valid, even if it has considered the appropriate hazards. It is probably most useful to use generic risk assessments as the starting point for site-specific assessments, and consider in a systematic way whether there are any additional hazards, or significant differences in the specific situation that require any additions or modifications to the generic assessment.

Case Study 2: 'Using a generic assessment when a site-specific assessment is needed'

Because of a lack of risk assessment competence in-house and in order to avoid duplicated effort, a company used a risk assessment prepared for a similar site, to show that, compared with risk tolerability criteria, the risks were as low as reasonably practicable. However there were significant differences between the natural features of the two sites. For example, the second one was immediately adjacent to a river, which was not a feature of the site with the risk assessment. As a result, accidental spillage of very toxic substances and subsequent contamination of the river had not been considered. The risk assessment was therefore incomplete and conclusions from the risk assessment were inappropriate.

Pitfall - When using a generic risk assessment it is essential that it is suitably tailored to take account of site-specific operations and locations.

3.2 DECIDING ON AN APPROPRIATE APPROACH AND DETERMINING THE LEVEL OF DETAIL NEEDED

The methodological approach to be adopted for the risk assessment should be determined at the outset. Risk assessment should be a structured, logical process. In order to decide what is an appropriate risk assessment approach in any particular situation, it is useful to consider the following issues:

- What level of detail is needed in the assessment;
- Is it appropriate to consider individual or societal risk or both?
- What is the most appropriate way to estimate the risk to individuals of interest.

Different approaches are also likely to be appropriate when estimating risks to safety, where the effects are acute or immediate, compared with estimating risks to health, where the effects are likely to be chronic or long-term. This report concentrates primarily on the assessment of risks to safety, however brief discussion of the differences in approach that would be appropriate when considering risks to health has been included where possible throughout the text.

3.2.1 Determining the level of detail needed in the assessment

HSW (1974) effectively defines the legally acceptable level of risks. The general duty on all employers is to reduce risks SFAIRP. HSE considers that this will be achieved if risks are reduced 'As Low As is Reasonably Practicable' (ALARP).

The level of detail in a risk assessment should be proportionate to the level of the intrinsic hazards. In general, the greater the magnitude of the hazards under consideration, and the greater the complexity of the systems being considered, the greater the degree of rigour and robustness (and hence the greater the level of detail) HSE requires in arguments to show that risks have been reduced ALARP. The level of risk arising from the undertaking should therefore determine the degree of sophistication needed in the risk assessment.

HSE starts with the expectation that controls in place must, at a minimum, achieve the standards of relevant good practice precautions, irrespective of specific risk estimates (HSE, 2001a). Relevant good practice provides duty holders with generic advice for controlling the risk from a hazard. In so far as they can adopt relevant good practice, duty holders are relieved of the need (but not the legal duty) to take explicit account of individual risk, costs, technical feasibility and the acceptability of the residual risk, since these will all have been considered when the good practice was established (HSE 2002a). In practice, if relevant good practice exists and is adopted for all workplace hazards, explicit evaluations of risk rarely need to be made in relation to day-to-day hazards; in these situations the risk assessment duty can be said to be discharged by the appropriate adoption of relevant good practice (see Case Studies 3 and 4). However, the good practice must be appropriate to the activity considered, relevant to the risks from the undertaking and cover all the risks from that undertaking (HSE, 2002a). The good practice must also be up to date, i.e. the guidance should reflect the current situation. Further guidance on recognised sources of good practice is given in HSE (2001a) (see Case Study 5).

The level of detail needed in the assessment should also be informed by initial consideration of the criteria that are going to be used to judge the acceptability, or otherwise, of the estimated risk. It only makes sense to estimate a particular risk in a particular way if it will allow a judgement to be made as to whether or not the risk is acceptable.

Case Study 3: 'Carrying out a detailed QRA without first considering whether any relevant good practice was applicable'

HSE's Railway Inspectorate (RI) issued a Prohibition Notice (PN) against a railway operator requiring them to include a specific safety procedure that involved staff carrying out physical checks on trains before carrying out a specific manoeuvre, as they were not satisfied that the procedures currently in place were adequate to ensure the safety of rail passengers. The railway operator appealed against the PN on the grounds that it was not reasonably practicable for them to physically check trains, as they argued that the cost of doing so was grossly disproportionate to the benefit. As part of their appeal, they prepared a QRA and CBA that they said demonstrated the gross disproportion.

Pitfall - It was not appropriate to carry out a detailed QRA without first considering whether there was any relevant good practice. Industry good practice in relation to the manoeuvre involved physically checking trains. In addition, the railway operator themselves carried out full physical checks on another part of their operation, effectively demonstrating the reasonable practicability of carrying out this action.

Other problems with RA - Prior to an operational change, the railway operator used to physically check trains prior to the manoeuvre considered, across all parts of their operation. By removing these checks, and not replacing them with an equally effective safety measure, the standard of safety was reduced - effectively a 'reverse ALARP' argument. There were also a number of problems with the QRA/CBA methodology. In the CBA, the railway operator had included in the cost side, costs to passengers in terms of the 'value' of passengers' time arising from delays they said would be introduced as a result of the additional physical checks. This is not a legitimate cost to include in the CBA; only costs to the duty holder should be included, not costs to others, as explained in HSE (2001a). In addition, the cost benefit analysis was not carried out such that there was a transparent bias on the side of safety. The analysis of benefits was based on historical data only; no account was taken of the known high levels of under-reporting or of near-miss information (see Section 3.8).

Case Study 4: 'Carrying out a detailed QRA without reference to relevant good practice'

A system to detect radioactive emissions and automatically shut down the reactor was introduced to all power stations. However, one power station had such a system that proved unreliable, as well as being difficult and costly to maintain. As a result it was decided to veto the system altogether and a review of power stations' safety cases claimed that it was not reasonably practicable to reinstate the system. This decision was supported by a quantified ALARP study and CBA which showed that the costs exceeded the benefits by a factor between 27 and 250 and were therefore grossly disproportionate. HSE challenged this conclusion on the basis that it was too reliant on CBA and did not give sufficient weight to good engineering practice. The safety case was rejected and a programme of work for the reinstatement of the system requested.

Pitfall - Good practice had not been considered. It is not appropriate to carry out a detailed QRA without first considering whether relevant good practice exists. The provision and continuing operation of the system at other similar power stations appeared to be industry good practice and therefore could be considered to be reasonably practicable.

Other problems with RA - HSE also questioned the validity of the cost estimates presented believing they were excessive and also some of the assumptions made by the operator in the CBA (see Section 3.8).

Case Study 5: 'Carrying out a risk assessment using inappropriate good practice'

A risk assessment was carried out to justify helipad refuelling at a hospital rather than at the nearby airport. Three arguments were used in this justification namely that it is safer, more environmentally friendly and more efficient. However, it is simpler and less contentious to refuel elsewhere. The risk assessment attempted to follow HSE guidance (HSG 146) which sets out a '5 steps to risk assessment' approach to petrol refuelling (no specific guidance exists for aviation fuel). However, the guidance makes it clear that 'where it is necessary to vary from normal industry practice then a more detailed risk assessment must be carried out'. It also stresses the importance of assessing surroundings, occupants etc. and explains that hospitals present particular problems in terms of evacuation.

Pitfall - The risk assessment does not recognise that it is not an industry standard to have refuelling stations on top of hospitals. Rather than follow inappropriate HSE guidance for a vaguely related type of operation, a detailed risk assessment specific to helicopter refuelling on a hospital helipad should have been carried out.

Other problems with RA - The second step of the risk assessment should consider what could go wrong and who could be affected - this does not appear to have been done. In particular no consideration of the potential for a conflagration in close proximity to a large number of vulnerable people (in the intensive care unit) is given. The risk assessment considers the benefits of being able to refuel on the helipad in terms of the additional number of people that can be transferred to the hospital by helicopter but neglects to take account of the additional risk to those already in hospital (see Section 3.4).

Risk assessment can be a qualitative, semi-quantitative or quantitative process. Any assessment should begin with a simple qualitative assessment, including consideration of whether any relevant good practice is applicable. In some cases it will be appropriate to supplement the qualitative assessment by a more rigorous semi-quantitative or quantitative assessment.

3.2.1.1 Qualitative risk assessment

Where the hazards presented by the undertaking are few or simple, for example in many small businesses, it is appropriate to just carry out a simple qualitative risk assessment. This can be a very straightforward process based on informed judgement and reference to appropriate guidance. Where the hazards and risks are obvious they can be addressed directly, and no complicated process or skills will be required (MHSWR, 1999).

A qualitative risk assessment should be a systematic examination of what in the workplace could cause harm to people, so that decisions can be made as to whether existing precautions or control measures are adequate or whether more needs to be done to prevent harm. HSE has published guidance on carrying out simple, qualitative assessments in the booklet 'Five steps to risk assessment' (HSE, 1999a).

In carrying out a qualitative risk assessment it is not necessary to follow explicitly all the steps outlined in the introduction to section 3. In a qualitative assessment it is appropriate to complete just the following steps:

- Identify the hazards;
- Identify the possible consequences (Decide who might be harmed and how);
- Evaluate risk (and decide whether the existing precautions are adequate or whether more should be done);
- Record the findings.

The text in brackets reproduces the way these steps are described in the '5 steps' booklet (HSE 1999a) where the wording is either different or expanded. The 'fifth step' in the '5 steps' booklet is described in this report as one of the post-assessment activities: 'review your assessment and revise it if necessary'. In carrying out the assessment it is also important to: ensure transparency throughout; ensure appropriate and adequate consideration of human factors; and ensure adequate handling of uncertainty.

As outlined in section 3.2.1, in practice, if relevant good practice exists and is adopted for all workplace hazards, explicit evaluations of risk rarely need to be made in relation to day-to-day hazards; in these situations the risk assessment duty can be said to be discharged by the appropriate adoption of relevant good practice (see Case Studies 3 and 4).

3.2.1.2 Semi-quantitative risk assessment

In many intermediate cases where the hazards are neither few and simple, nor numerous and complex, for example if there are some hazards that require specialist knowledge, such as a particular complex process or technique, it may be appropriate to supplement the simple qualitative approach with a semi-quantitative assessment.

In carrying out semi-quantitative risk assessments, simple qualitative techniques, supplemented by for example measurements to identify the presence of hazards from chemicals or machinery, or the use of simple modelling techniques may be appropriate. Simple modelling techniques may be used to derive order of magnitude estimates of the severity of the consequences and likelihood of realisation of hazards. These estimates can be combined to obtain estimates of the order of magnitude of the risk.

A number of different techniques for carrying out semi-quantitative risk assessments exist, including risk matrix approaches (Worsell and Wilday, 1997; Middleton and Franks, 2001) and lines of defence/layers of protection analysis (Franks, 2001).

3.2.1.3 Quantitative risk assessment

Where the hazards presented by the undertaking are numerous and complex, and may involve novel processes, for example in the case of large chemical process plants or nuclear installations, detailed and sophisticated risk assessments will be needed, and it is appropriate to carry out a detailed quantitative risk assessment in addition to the simple qualitative assessment.

Quantitative risk assessment (QRA) involves obtaining a numerical estimate of the risk from a quantitative consideration of event probabilities and consequences (in the nuclear industry the term 'probabilistic safety analysis' is used in place of QRA). Detailed guidance on the

application of QRA in the chemical process industries is given in IChemE (1996) and AIChemE (1989).

In carrying out quantitative risk assessments, special quantitative tools and techniques will be used for hazard identification, and to estimate the severity of the consequences and the likelihood of realisation of the hazards. Where such methods and techniques are used it is important that they are carried out by suitably qualified and experienced assessors. The results of the QRA will be numerical estimates of the risk, which can be compared to numerical risk criteria at the risk evaluation stage.

3.2.2 Deciding whether it is appropriate to consider individual or societal risk, or both

There are two commonly used types of risk measure: individual risk measures and societal risk measures. Individual risk is defined as the frequency at which an individual may be expected to sustain a given level of harm from the realisation of specified hazards (IChemE, 1985). Societal risk is defined as the relationship between frequency and the number of people suffering from a specified level of harm in a given population from the realisation of specified hazards (IChemE, 1985).

Individual and societal risks are different presentations of the same underlying combinations of incident frequency and consequences (AIChemE, 1989). Individual risk is a measure of the risk to a particular individual, while societal risk on the other hand is a measure of the risk to society as a whole.

Both individual and societal risk measures are of importance in assessing the benefits of risk reduction measures and in judging the acceptability of the particular process or operation being assessed. However in most cases decisions as to whether the level of risk is acceptable or whether more needs to be done to control or reduce the risk will be informed primarily by either one or the other. In deciding whether it is appropriate to consider individual or societal risk, or both, it is important to consider the number of people that could be affected at the same time, from realisation of the hazards considered.

It is always necessary to consider individual risk to determine whether the risk to the individuals of interest is acceptable, or not. If the hazards being considered have the potential to affect only individuals, or a few people at the same time it would be appropriate to only consider individual risk. However, if the hazards being considered have the potential to affect a large number of people at the same time, it would also be necessary to consider societal risk, even if the individual risk was estimated to be low.

Case Study 6: 'Making decisions on the basis of individual risk estimates when societal risk is the appropriate measure'

A safety case for a new rail terminus was presented to HSE. The safety case encompassed all aspects of the operation of the terminal and considered risks to employees, passengers and members of the public. The risk assessment considered individual risk to workers, which is acceptable, but then used similar methodologies to consider risks to the public and passengers. For example, risks from fire/ evacuation/ overcrowding were calculated as individual risk and not societal risk. It was then

argued that because a regular traveller will not be spending long (when calculated on a yearly basis) in the terminal, the risk could be considered negligible.

Pitfall - Societal risks must be determined when there is the potential for harm to large numbers of people. These risks may not be negligible due to the large numbers of people who may be exposed, even when each individual has minimal exposure.

3.2.3 Approaches to estimating individual risk

A number of different approaches to estimating individual risk may be adopted, and it is important to ensure that the approach taken is appropriate in any particular situation.

As stated in section 3.2.2, individual risk is defined as the frequency at which an individual may be expected to sustain a given level of harm from the realisation of specified hazards. The 'individual' that is the subject of the assessment needs to be carefully chosen and defined; initial consideration of this should have been made when defining the scope of the assessment, as discussed in section 3.1.2.

While risk assessments can be (and sometimes are) carried out to assess the risk to an actual person, i.e. the risk to an individual taking full account of the nature, extent and circumstances in which the exposure arises, such an approach is of limited use for managing risks generally, as explained in HSE (2001a).

A more commonly adopted and useful approach is to carry out an assessment of the risk to a hypothetical person. A hypothetical person describes an individual who is in some fixed relation to the hazard, e.g. the person most exposed to it or a person living at some fixed point or with some assumed pattern of life. To ensure that all significant risks for a particular hazard are adequately covered, it will usually be necessary to construct a number of hypothetical persons. For example, for each population exposed to the hazard, it will be necessary to construct a hypothetical person in order to determine the control measures necessary to protect that population.

The aim in considering risks to hypothetical persons is to provide a 'full picture' of the risks generated by a hazard. This will be achieved by creating enough hypothetical persons to enable control measures to be put in place to protect all those exposed from the undesirable consequences of the hazard, taking account of the different populations exposed and the circumstances of their exposure (HSE, 2001a).

Risk is often expressed in units of frequency, that is harm per unit time, e.g. the chance of receiving the specified level of harm per year. Sometimes however, other measures of activity, such as per journey, per shift or per passenger mile, may be more appropriate depending on the risk problem to be solved. It is important to ensure, whatever measure of activity is used, that common pitfalls are avoided, such as only considering the risk from one activity or dividing the time spent on the hazardous activity between several individuals.

In deciding what hypothetical persons will allow the determination of a 'full picture' of the risks, it may be helpful to consider, in general terms, the criteria that will be used to determine whether the existing control measures for a hazard are adequate or whether more needs to be

done to control or reduce the hazard. In general, when comparing the estimated individual risk with risk criteria, it will be important to ensure that both:

(i) The risk to all individuals exposed is below the unacceptable threshold (and hence tolerable if ALARP), taking account of the actual exposure of particular individuals to all the hazards to which they are exposed;
(ii) The risks from each individual hazards is below the unacceptable threshold (and hence tolerable if ALARP) irrespective of individual exposure to the hazard.

Consideration of these two aspects in the assessment should help to ensure that common pitfalls are avoided.

3.2.3.1 Estimating the risk to particular groups of persons

If a hypothetical person is to be constructed to represent a particular group of persons, such as those interacting with the hazard in a particular way (e.g. the most exposed group of workers), or those who may be particularly vulnerable to it (e.g. new or inexperienced workers), it is important that the hypothetical person is truly representative of the group of people they are designed to represent. For example, all the hazards to which the hypothetical person is exposed must be taken into account when estimating the overall individual risk. Recognition of the fact that exposure to many hazards is not uniform but comes in peaks and troughs should be factored in, by for example time-weighting the period of exposure of the hypothetical person to each of the different hazards, or to different attributes of the same hazard, to get an accurate reflection of the overall risk to the individual. It is also appropriate to take account of the overall exposure time of the hypothetical person, for example to estimate the risk over a typical working week. If the risk from any of the individual hazards is high, additional consideration of the risk from that hazard will also be needed, as explained in sections 3.2.3.2 and 3.2.3.3.

Case Study 7: 'Only considering the risk from one activity'

A large item of gas fired plant was fitted with an enclosure. Occasionally access by personnel was necessary for maintenance purposes. Should a leak of gas occur during the maintenance procedure and ignite, there would be a high likelihood of death or serious injury to the personnel. Data was available on the frequency of leaks from pipework and the probability of ignition. This showed the likelihood of this happening coincidentally with the presence of the personnel to be acceptably low compared to published tolerability criteria.

Pitfall - This was a worthless comparison because the personnel may have been at similar risk through other tasks carried out during their working day/week/year and it is the total risk from all activities that should be compared with the criteria. In addition, this example illustrates another common pitfall, that of using independent probabilities for events that are not independent.

Other problems with RA - Maintenance activities are a significant cause of releases from pipework and the use of averaged data is inappropriate in this case (see Section 3.6).

3.2.3.2 Estimating the risk in cases of 'transient exposure' - avoiding 'salami slicing'

When considering the risk from a particular activity, situation or process to which individuals are exposed for only a short time, care is needed to ensure that an accurate picture of the risk is obtained. It is not appropriate to divide the time spent on the hazardous activity between several individuals and estimate the risk on this basis (the 'salami slicing' technique). For example, if any one person is only exposed to the hazard for a short time, but someone is always exposed, it would give a misleading picture of the risk to estimate the risk from this hazard by taking into account the exposure time of each individual. Instead, a truer picture of the risk would be obtained by constructing a hypothetical person who is exposed to the hazard 100% of the time (or for the proportion of time that any individual is exposed). In this case it would be necessary to define precisely the hypothetical person's location/interaction with respect to the hazard.

Case Study 8: 'Incorrectly defining the appropriate hypothetical person - the 'salami slicing' approach to risk estimation'

An operator, when carrying out a risk assessment for a safety case, was considering the risks associated with a particularly 'high risk' activity. A hypothetical worker was defined as one who works 40 hours per week but only spends a proportion of their time working on the high risk task thus reducing their overall individual risk. However, when the first worker stops performing the high risk task, they are replaced by a second worker. In the extreme, there may always be someone performing the task, just not the same person.

Pitfall - The employer is confusing the individual risk of a specific person with the risk to anyone performing the high risk activity. The hypothetical person defined (when considering the high risk activity) should represent the worst case for occupancy of the task. That is, if the task is performed continuously by a team of workers in rotation, then the individual risk should be calculated on the basis of 100% exposure rather than the fraction of time any one of the team is performing the task.

3.2.3.3 Estimating the risk from 'risk hotspots'

The approach to be adopted when considering 'risk hotspots' is an extension of the approach that should be adopted to avoid salami slicing. Avoidance of salami slicing is particularly important when considering risk hotspots, as estimating the risk for each individual on the basis of a short exposure time in such situations could disguise an intolerable level of risk.

When considering the risk from high hazard activities, situations or processes (risk hotspots), the risk from the hazard should be estimated without taking into account the exposure time of individuals. The risk from the hazard needs to be below the unacceptable threshold (and tolerable if ALARP) at all times. For example, it would not be possible to say that the overall

risk to workers is in the tolerable (if ALARP) region if for even a short time they are engaged on a very high hazard activity with an unacceptable risk.

Case study 17 also illustrates this particular risk assessment pitfall.

3.3 DECIDING WHO SHOULD BE INVOLVED

A team approach to risk assessment should be adopted whenever possible. Pooling the knowledge, skills, expertise and experience of a range of people with different perspectives should ensure comprehensive coverage of all hazards. It is particularly important to involve employees who have practical experience of the particular process or activity being considered in the risk assessment as they will often have the best knowledge and understanding of any hazards. The risk assessment process should also always involve management or those with responsibility for health and safety (who should be aware of relevant legislation), whether or not advisors or consultants assist with the detail, as they are responsible for ensuring that the process is adequate (HSE, 1999a).

The number and range of people who need to be involved in the assessment will vary with the level of detail needed for the assessment. For example, in very small firms with few employees and for which simple qualitative assessments are appropriate, it may be appropriate for the owner/manager to carry out the assessment, provided they are confident they understand what is involved. In larger firms for which simple qualitative assessments are appropriate it may be appropriate for the owner/manager to lead a small assessment team including an employee representative or operator, and a safety representative or safety officer. In situations where more detailed quantitative assessments are appropriate, it will usually be appropriate to involve a team of people at each stage of the assessment, including operators, specialists, safety representatives and management.

Risk assessment often involves a multi-disciplinary approach, since it may cover a variety of areas of expertise or the systems being assessed may be too complex to be fully understood by one person. A group of people with different skills and expertise will therefore be needed to carry out a full assessment. It is important to involve a representative of the employees who carry out the task being assessed, as they know how the job is actually done, may have experience of abnormal as well as normal conditions and understand the scope for dangerous shortcuts. They can also provide useful information about aspects of the work that are difficult, and any ideas for making the work more efficient or safer. A teamwork approach enables a true picture of the activity to be built up and increases the potential pool of ideas for improvement.

The individuals or working group involved in the risk assessment process should be familiar with the assessment methods used, have a thorough knowledge of the subject under consideration and other necessary specialised knowledge should be provided and integrated into the assessment as required (BS 8444, 1996). That is, all those involved in the risk assessment process should be competent to undertake the task (see Case Study 9).

Competence means the ability to undertake responsibilities and to carry out activities to the necessary standard. Competence is a combination of knowledge, experience and expertise, usually involving both practical and theoretical skills. While training has a part to play in developing competence, competence means more than simply training. Experience of applying skills and knowledge is another important ingredient and needs to be gained under adequate supervision (HSE, 1997).

Although the number and range of people involved will vary with the level of detail needed for the assessment, the general principles outlined above apply in all situations. In summary, these are that:

- A team based approach should be adopted wherever possible;
- Input from employees with practical experience of the process or activity being assessed should be included;
- Management (or those with responsibility for health and safety) should be involved;
- All those involved in the risk assessment process should be competent to undertake the task.

Case Study 9: 'Not involving a team of people in the assessment/not including employees with practical knowledge of the process/activity being assessed'

An accident occurred at a factory in which an employee had an arm amputated by a machine he was cleaning. A risk assessment had been carried out by the factory manager, but he had not tested the interlocks (because he didn't have a detailed knowledge of how the machine was operated) or noticed that some of the interlocks and guard switches were missing from the machine. The manager had received no formal training on how to complete the risk assessment form.

Pitfall - Risk assessments should be carried out by a team of competent people which includes operators who are familiar with the machine or task being assessed. As a result of this failure significant hazards were missed and therefore not controlled.

The use of consultants to carry out risk assessments should also be treated with caution. Over-reliance on consultants (whether external or internal to the company) can mean that:

- Those responsible for the activity lose some of the value of the risk assessment process, as their own understanding of the issues is not developed to the same depth;
- A risk assessment that is unrealistic or with inappropriate conclusions is produced if the process is not managed to ensure that the consultants have adequate knowledge of the process/operation, and work closely with those responsible for the activity.

Where consultants are utilised, the process should be well managed by the company to ensure that the consultants work closely with the company to avoid such problems.

Case Study 10: 'Ineffective use of consultants'

When a rail operator first carried out a QRA on their operations, they employed a consultancy to develop a QRA model and do the risk assessment for them. Initially the operator failed to buy into the model and there was little consultation between the contractors and the operator during the development of the model. The model included a detailed fault tree and event tree approach which needed a certain level of QRA

expertise to understand. As a result of the lack of discussion between the two parties, the company ended up owning a detailed QRA model for which they had very little in house knowledge of its development and a lack of supporting documentation including the justifications for the data used and assumptions made.

Pitfall - Where consultants are used to carry out risk assessments, it is essential that the company works closely with the consultants to ensure that a number of staff within the company have a deep understanding of the risk assessment, including its scope and limitations.

3.4 IDENTIFYING THE HAZARDS

Hazards associated with work activities can be present as a result of any one or a combination of the following: substances, machinery/processes, work organisation, tasks, procedures and the people and circumstances in which the activities take place including the physical aspects of the plant and/or premises. It is important that all these different elements are addressed to ensure that the hazard identification process is as thorough as possible. As well as identifying the 'intrinsic hazards' in the workplace, the hazard identification process also involves the identification of all possible routes to failure, i.e. identification of the various mechanisms by which the hazards could be realised. Hazard identification is a critical step in the risk assessment process, as a hazard omitted is a hazard not analysed or controlled (see Case Studies 11-13).

A number of tools, techniques and approaches may be used to assist in the hazard identification process. Different approaches are likely to be appropriate depending on the type of assessment that is being carried out (e.g. whether a qualitative, semi-quantitative or quantitative assessment is being undertaken), but there are a number of sources of information that are likely to be of relevance whatever approach is adopted. Relevant sources of information include:

- Legislation and supporting Approved Codes of Practice which give practical guidance and include basic minimum requirements;
- HSE guidance;
- Process information;
- Product information from manufacturers and suppliers;
- Relevant British and International standards;
- Industry or trade association guidance;
- The personal knowledge and experience of managers and employees;
- Accident, ill health and incident data (including near miss data) from within the organisation, from other organisations or from central sources;
- Expert advice and opinion and relevant research.

It is important when carrying out the hazard identification process to adopt a structured, systematic approach and to ensure that the hazards identified reflect the current process or system or operation. For example, it is not appropriate to base the hazard identification solely on lessons learned from historical data. Some hazards present may never have been realised, and therefore may not show up in the historical data, or there may have been changes to the system/process/operation such that the historical data is not a good representation of the hazard potential of the current situation. Adopting a team based approach to hazard identification,

wherever possible, involving people with a range of knowledge, skills, expertise and experience should ensure comprehensive coverage of all hazards.

It is also important to ensure that an analysis of non-routine (such as emergency or breakdown maintenance/activities) as well as routine, planned activities is included in the hazard identification process, and that consideration of all those who may be affected by all activities considered is included. It is also important to ensure that there is adequate consideration of human factors; ways to ensure adequate consideration of human factors throughout the risk assessment process are discussed further in section 3.11.

Case Study 11: 'Failure to identify all relevant modes of operation'

During the hazard identification stage of a risk assessment, to identify faults that can occur during operation, the duty-holder did not include all operational states. In particular, a transient state, which involved reducing the cooling to the vessel such that some of the temperature alarm settings were exceeded, was not considered. This was important as it required a different approach to control for this stage of plant operations. It is particularly important to include all operational states, that require different control requirements, when identifying faults. In this case, the choice between operation with a standing alarm or changing the alarm setting, both of which raise issues of safety management, was not subject to proper assessment being left to local plant management to decide; nor was the possibility of alternative protection considered.

Pitfall - There are often transient states, such as start-up and shut down of plant, which in varying ways can mean that safety measures used in normal operations are not available or have to be changed. By omitting these states, the risks during them are not assessed and may become dominant.

It is also important to ensure when carrying out the hazard identification process that dependence between failures or common-cause failures are considered, e.g. organisational effects on individual operator errors. Failure to consider common-cause failures is a common pitfall, and represents a major way in which the development of actual incidents can diverge from any assumed independence of events in the analysis.

Case Study 12: 'Failure to consider common-cause failures'

A duty-holder had a fourfold redundant pumping system and assumed a low probability for the loss of all pumps. However all four pumps were sited close together near a roadway and consequently were vulnerable to a vehicle hitting them all. This scenario could potentially result in the loss of all four pumps. Protection in the form of a crash barrier improved the situation when a more detailed analysis (and a check on the actual plant situation) was carried out.

Pitfall - Failure to consider dependent failures is a common pitfall. Modelling of redundant systems must consider the system in sufficient detail such that common-cause failures are included.

As outlined in section 3.1.3, it is not necessary to take account of hazards arising from events and behaviour that are not reasonably foreseeable (HSE, 2001a). Nor is it necessary to take account of trivial hazards, such as those arising from routine activities associated with everyday life, unless they are compounded by the work activity, or there is evidence of significant relevance to the particular work activity.

In the simplest cases (for example where a simple, qualitative risk assessment is appropriate), hazards can be identified by observation and by comparing the circumstances in the workplace with the relevant information. The guidance in HSE (1999a) on how to look for the hazards is for the person carrying out the assessment to walk around the workplace and look afresh at what could reasonably be expected to cause harm. Trivial hazards can be ignored, and the hazard identification should concentrate on significant hazards that could result in serious harm or affect several people. It is important to ask employees or their representatives for their ideas, as they may have noticed things that are not immediately obvious. It is also important to ensure that work activities are analysed as they are actually carried out in practice, rather than as they should be carried out (e.g. as specified in procedures) if there are differences. Manufacturers' instructions or data sheets can also assist with the identification of hazards, as can accident and ill-health records.

In more complex cases, in addition to the simple analysis techniques outlined above, some measurements such as air sampling, or examining the method or methods of machine operation may be necessary to identify the presence of hazards from chemicals or machinery.

In the most complex or high risk cases (for example in the chemical process or nuclear industry, where detailed quantitative risk assessments are being carried out) special tools and techniques for hazard identification may be needed. There are a variety of methods and techniques for hazard identification, and no single method can be recommended for all circumstances. Different techniques have different strengths and weaknesses, and may be most appropriately applied at different phases of a project or process lifecycle. The following methods are commonly used for hazard identification, particularly in the chemical process industries:

- Hazard indices;
- Hazard and Operability (HAZOP) studies;
- What if? Analysis;
- Check-lists;
- Failure modes and effects analysis (FMEA);
- Preliminary Hazard Analysis;
- Fault tree analysis;
- Event tree analysis;
- Task analysis.

Specialist advice may be needed to choose and apply the most appropriate method; many of the methods listed above can only be used effectively by suitably qualified and experienced assessors. Guidance on the use and applicability of the majority of these hazard identification techniques is given in IChemE (1996), and a comprehensive review of these and many other

hazard identification techniques, including discussion of the advantages and disadvantages of each technique, and further references has been carried out by Glossop *et al.* (2000).

Case Study 13: 'Failure to identify all hazards associated with a particular activity'

An employee was undertaking maintenance work which involved the use of high pressure water jetting equipment whilst suspended by ropes. He slipped causing the water jet to come into contact with his body resulting in a deep cut. Prior to the maintenance work beginning, a risk assessment had been undertaken which identified hazards associated with both high pressure water and working at height. Consequently the team of workers were made aware of procedures for both rope access operations and water jetting before commencing work. The subsequent investigation found that the operator had not carried out a suitable and sufficient risk assessment. They had not identified the additional hazards associated with the use of rope access techniques in conjunction with high pressure water jetting and consequently had not identified any appropriate risk reduction measures, e.g. consideration of alternative access techniques or provision of adequate personal protective equipment.

Pitfall - The pitfall in this example was failure to properly analyse the operations being carried out, as a result of which all associated hazards were not identified.

3.5 IDENTIFYING THE POSSIBLE CONSEQUENCES

Following on from the identification of hazards, the next step in the risk assessment process is the identification of the possible consequences of the realisation of the hazards, that is, an assessment of who might be harmed, and of the effects and severity of the harm caused. This stage therefore involves the estimation of both the magnitude of the physical effects arising from realisation of each hazardous event and the estimation of the severity of the harm caused to all those who may be affected (within the scope of the risk assessment).

If individual risk is being calculated, an estimate of the magnitude of the consequences of the event can be obtained by determining the extent over which a specified 'level of harm' would be experienced. If societal risk is being calculated, an estimate of the magnitude of the consequences of the event can be obtained by determining the number of people who would experience a specified 'level of harm'. In estimating the severity of the harm caused, it is important to consider all those who may be affected (within the scope of the assessment), including employees, other workers in the workplace, and members of the public. It is important to remember to include those who may not be in the workplace all the time (e.g. cleaners, visitors, contractors, maintenance workers, etc.) and to identify groups of workers (or others) who may be particularly vulnerable to the hazards (e.g. young or inexperienced workers, new and expectant mothers, those who work alone, etc.) (MHSWR, 1999).

Different approaches to estimating the consequences arising from realisation of the hazardous events will be appropriate depending on the approach that is being adopted for the risk assessment, i.e. whether a qualitative or quantitative assessment is being carried out. However,

whether a qualitative or quantitative approach is adopted, the identification of the possible consequences should (Franks *et al.*, 2000):

- Be based on the hazardous events identified;
- Describe any consequences arising from the hazardous events;
- Take into account existing measures to mitigate the consequences (including consideration of their effectiveness) together with consideration of all relevant conditions that have an effect on the consequences;
- Consider both immediate consequences and those that may arise after a certain time has elapsed, if consistent with the scope of the study (consideration of delayed consequences is likely to be more appropriate when assessing ill-health rather than safety risks).

This stage of the assessment should also take into account any subsequent actions to mitigate the effects of the consequences.

Qualitative estimates of the consequences of the realisation of hazards can be derived based on analysis of relevant sources of information about the hazards, from accident and incident data, and from collation of the knowledge and experience of a wide range of employees. When estimates of the consequences are derived qualitatively, it is common to categorise the severity of the consequences of exposure using either numerical indicators or descriptions (such as catastrophic, major, minor). Estimating the likely severity of the consequences often involves an element of subjectivity (Simpson, 1996), but this can be reduced considerably by careful and explicit definition of the category indicators or descriptors.

Quantitative estimates of the consequences of the realisation of hazards can be derived using appropriate quantitative consequence models. Different models are usually used to estimate the physical effect of the hazard, the impact of these effects on the exposed population, and to take into account any mitigating factors. A range of models exists, from simple analytical models, to very complex computer models. Care should be taken to ensure that the methods and models used are appropriate to the problem being considered (BS 8444, 1996). It is also important to ensure that uncertainties in the consequence model results are properly accounted for; consideration of uncertainty, and approaches for dealing with it appropriately throughout the risk assessment process are discussed in section 3.12.

When considering the action of persons and systems following the occurrence of an initiating event, it is necessary to ensure that the whole range of possible behaviour with respect to time is considered.

Case Study 14: 'Reliance on archetypal descriptions of events and behaviours without also considering the range of other possible outcomes'

When considering the time that a vessel containing a liquefied flammable gas would take to rupture under fire attack it was noted that typically this would be much longer than the time it would take to evacuate a nearby office building. It was therefore concluded that there was no significant risk to the building occupants from this event. However if the full range of times to rupture and evacuation times had been considered, there was a significant probability that the rupture would occur when the persons evacuated from the building were still within the hazard range and in the open.

Pitfall - Reliance on archetypal descriptions of events and behaviours without also considering the range of other possible outcomes resulted in incorrect conclusions.

When considering events with more than one outcome, it is necessary to adopt cautious assumptions for each outcome.

Case Study 15: 'Failure to consider all possible outcomes in a cautious manner'

The release of a large quantity of flammable gas from containment may result in a fireball, vapour-cloud explosion, flash fire or no significant consequence depending on whether ignition is immediate, somewhat delayed, delayed, or if no ignition occurs, respectively. In one study it was noted that a significant source of ignition was located in the vicinity of the postulated release. When considering the likelihood that the consequence would be a fireball, cautiously the probability of immediate ignition was assessed as unity. Correspondingly the probability of other outcomes was set to zero.

Pitfall - The analysis did not consider all possible outcomes in a cautious manner.

3.6 ESTIMATING THE LIKELIHOOD OF THE POSSIBLE CONSEQUENCES

Once the possible consequences of the realisation of identified hazardous events have been identified, the next step in the risk assessment process is to estimate the likelihood that each of the identified hazardous events will occur. Estimates of the likelihood of hazardous events occurring are usually expressed as a frequency of occurrence over a specified time interval (for example, per year or per shift), but sometimes likelihood is expressed as a probability of occurrence, and is therefore dimensionless. The distinction between dimensionless probability and frequency values is important if these numbers are to be used in combination.

The choice of approach for likelihood estimation, and the detail of how the approach is adopted will depend on the approach that is being adopted for the risk assessment, i.e. whether a qualitative, semi-quantitative or quantitative assessment is being carried out.

If a qualitative risk assessment is being carried out, the step of assessing the likelihood of hazardous events may not be carried out explicitly. This step may instead be carried out implicitly as part of the step of risk evaluation. HSE (1999a) suggests that the process of risk evaluation begins with consideration of how likely it is that each hazard could cause harm, as this will determine whether or not you need to do more to reduce the risk.

Whether carried out explicitly or implicitly, qualitative estimates of hazardous event likelihoods are usually derived based on a combination of analysis of historical data, and informal expert judgement (i.e. experience). When estimates of the event likelihoods are derived qualitatively, it is common to categorise the event likelihood either numerically or using descriptions (such as likely, probable, possible, remote, improbable). Such categorisation introduces an element of subjectivity, which is harder to eliminate by explicit definition of the categories than in the case of qualitative consequence estimation, as it is hard to explicitly

define likelihood categories qualitatively. The more precise the definition, the more likely it will be necessary to possess accurate, quantitative predictive data.

Semi-quantitative estimates of hazardous event likelihoods (for example, order of magnitude estimates) can be determined from a combination of analysis of historical data, simple modelling and informal expert judgement. Franks (2001) describes the use of semi-quantitative likelihood assessment within the LOPA (Layer of Protection Analysis) technique. In this technique, frequencies are assigned to initiating events, given failure of all the protective layers. Initiating event frequencies may be obtained from public domain sources or through the use of simple fault or event trees. The data used should be appropriate to the industry or operation under consideration. The failure rate is then reduced by one or two orders of magnitude for each protective layer, or safety system that is effective, reliable and auditable. LOPA is intended to be a simplified approach giving order of magnitude risk estimates. A high degree of accuracy in the failure data is therefore not warranted.

Quantitative estimates of hazardous event likelihoods are usually derived using a combination of the following three approaches with full quantification of fault and event trees, and formal expert judgement:

1) Using relevant historical data;
2) Derivation via analytical modelling techniques;
3) Using expert judgement.

It is possible to derive incident frequencies directly from the historical record if there are sufficient and accurate data available, and the data are relevant and applicable to the particular process/hazard under review. The 'historical approach' to derivation of incident frequencies based on appropriate data has the advantage that it is not limited by the imagination of the analyst in deriving failure mechanisms (AIChemE, 1989) and therefore the assessment will not omit any significant routes to failure. The data should already encompass all common relevant contributory aspects including: reliability of equipment, human factors, operational methods, quality of construction, inspection, maintenance and operation, etc. (IChemE, 1996). However, the data may not include rare incidents, as they may not have occurred, unless the population of items is very large, and outdated failure modes that may not be relevant to the specific case under study may also be included, resulting in inaccurate estimates of the chance of the event. Statistical techniques exist for addressing the problem of sparse historical data that can be particularly useful for estimating the likelihood of rare events. The problems posed by having sparse historical data can be tackled by adopting a Bayesian approach (Williams and Thorne, 1997): from as wide a knowledge base as possible (including making use of relevant experience from similar situations and elicitation by expert judgement), a bounded or cautious failure frequency is proposed; as additional relevant information becomes available this initial estimate is modified in accordance with Bayesian Analysis, to refine the estimate.

One particular type of historical data, generic failure frequency data is commonly used in QRA studies. An explanation of this data and the implications of its use are discussed by Franks *et al.* (2000). An additional potential problem with using historical accident data is accounting for the degree of under-reporting in the data; the degree of under-reporting tends to increase as the severity of the consequences decreases (Adams, 1998).

Case Study 16: 'Inappropriate use of generic failure data'

Using generic failure data for large items of plant rather than carrying out a detailed analysis can lead to failure to recognise the importance of specific parts of the system. It is common practice to do this, for example, for cranes where an overall frequency is applied for dropping loads. There are several protective systems which can be used to reduce the likelihood of this happening and if a generic approach is used without first checking that the particular crane has the same systems as those on which the generic data is based, a crane which is not adequate may appear to be so.

Pitfall - This is a particular example of the dangers of using generic data rather than specific data. Reliability and failure data are often quoted with insufficient information regarding the environment in which the system is used, the maintenance required and the type of operation carried out. This can lead to significant errors if the data is used without due consideration.

Modelling techniques to estimate incident frequencies from more basic data are used when suitable historical data are not available or are inadequate. Various techniques are available, but the most common and widely used techniques are Fault Tree Analysis (FTA) and Event Tree Analysis (ETA). Both of these techniques are essentially qualitative in nature, providing models that form the basis for subsequent quantification.

FTA permits the hazardous event frequency to be estimated from a logical model of the failure mechanisms of the system (AIChemE, 1989). A 'fault tree' is developed that includes all events that can contribute to the hazardous event under consideration (the 'top event') and graphically illustrates all the logical sequences of sub-events that could result in realisation of the top event, indicating where only one or more than one sub-event needs to occur for the sequence of events to propagate to the top event. The fault tree model is based on the combinations of failures of more basic system components, safety systems and human reliability. The fault tree can provide powerful qualitative insight into the potential failure modes of a complex system. ETA allows the identification and quantification of possible event outcomes in a systematic, logical way following the initiating event. An 'event tree' is developed that graphically illustrates all possible outcomes following realisation of the initiating event; it depicts the chronological sequence of events that could occur following the initiating event, including escalation and mitigation.

Quantification of the fault trees requires numerical data on component failure rates, protective systems unavailability and human error rates. Several types of data can be used to quantify event frequencies or probabilities, including accident data, incident or near miss data, and event data (more detailed than near miss data; a record of all non-trivial events that could be a step towards an incident) (IChemE, 1996). The most appropriate data of any type for use in a given application are suitable data relating to the particular situation. If available data is not relevant/suitable, it is necessary to use data from other sources, wherever possible comparing this with any information available from the particular situation. There is often considerable uncertainty associated with the frequency estimates, which increases as hazardous event frequencies become rarer. It is important to ensure that this uncertainty is properly accounted

for; consideration of uncertainty, and approaches for dealing with it appropriately throughout the risk assessment process are discussed in section 3.12.

The main advantage of using modelling techniques such as FTA and ETA are that they allow a thorough understanding of an activity to be built up, even if quantification is not performed. In addition to identifying all potential paths that could lead to failure, the techniques can also serve to single out the critical events that contribute significantly to the likelihood of failure of a system, and reveal potential weak links in the system (Wang and Roush, 2000). However, errors can be made if the analyst is unaware of the theoretical basis underlying the construction, manipulation and evaluation of the fault and event trees, and sufficient expertise is needed before the techniques can be used (IChemE, 1996). In addition, considerable effort is usually required to develop and quantify the trees, and there is the potential for error if failure paths are omitted or manual calculation methods are incorrectly employed (AIChemE, 1989). As outlined in section 3.4, it is particularly important when identifying all potential paths to failure that common-cause failures are considered.

Expert judgement is also sometimes used as an approach for deriving incident frequencies (BS 8444, 2000). There are a number of formal methods for eliciting expert judgement that make the use of judgements visible and explicit and provide an aid to the asking of appropriate questions. Expert judgements should draw upon all relevant available information including historical, system-specific, experimental, design, etc.. Available methods include the Delphi approach, paired comparisons, category rating and absolute probability judgements.

Case Study 17: 'Inappropriate use of historical data'

A company employed workers who carried out tasks on top of stacked containers on board ships at heights of 15 metres or more. The company argued against the routine use of safety harnesses or guardrails on the basis that the risk of falling was negligible, the time spent at the edge of the containers was short, and that use of the safety equipment reduced productivity. They also tried to show, *via* a QRA and CBA, that implementation of additional control measures would be grossly disproportionate to the reduction in risk. However, they assumed that because no accident had occurred in recent history (the last 10 years) the risk was effectively zero. In addition they also failed to recognise that it was UK industry practice to protect workers at heights.

Pitfall - Accident statistics based on limited sample exposures must be used with care and established statistical techniques for deriving the likelihood of an event need to be used. Such techniques, in this case, showed that the risk was not negligible and therefore the company needed to put additional measures in place to control the risk if reasonably practicable.

Other problems with RA - The calculated risk was averaged over all tasks carried out on top of the containers and did not recognise the risk hotspot of working close to the edge of the container (see Section 3.2.3.3). In addition the company should have recognised the importance of good practice rather than carrying out a detailed CBA to show that implementation of risk control measures was grossly disproportionate (see Section 3.8).

Case Study 18: 'Inappropriate use of statistical averaging'

The range of heat loading and toxicities that would occur in normal operation of a particular plant was well defined. However, when considering the relevant controls needed within the plant, an average input for heat loading and toxicity was used. As a result, the control measures were not adequate for the full range of expected hazards, only the average.

Pitfall - This example shows that it is not appropriate to apply a probabilistic approach to the scale of the hazard that will be encountered. Had this deficiency not been identified the plant would have been at risk as it would not have been adequately protected for the full range of possible hazards.

3.7 ESTIMATING THE RISK

The risk from an activity is a function of the likelihood of occurrence of possible undesired events and the magnitude of the associated consequences. The process of risk estimation therefore involves examining the initiating events or circumstances, the sequence of events that are of concern, and the severity and likelihood of the consequences of the identified hazards (taking account of any mitigating features) to produce a measure of the level of the risk being estimated. The three steps of hazard identification, likelihood estimation and consequence estimation can all therefore be considered as part of the risk estimation process and have been described in sections 3.4 to 3.6. The remaining part of the risk estimation process is to combine the estimates of likelihood and consequence for the identified hazardous events to produce a measure of the level of risk.

A useful approach to estimating the risk that is often employed in more extensive analyses (e.g. typically in semi-quantitative or detailed quantitative assessments) where the number of possible sequences of events can become very large and unmanageable, is to carry out the assessment using a representative sample of events. Rather than analysing the whole system explicitly, a discrete set of scenarios is selected to represent the whole. Care must be taken if such an approach is adopted to ensure that the chosen set of scenarios is truly representative of the whole system. The process is not mathematically rigorous, and can only be validated by means of a sensitivity study addressing whether a different and/or more comprehensive sample would result in a significantly different answer. Simplifications should err on the side of caution.

Case Study 19: 'Inappropriate definition of a representative sample of events'

When carrying out an analysis involving toxic gas dispersion it is generally accepted that a few discrete combinations can be used to represent the whole of the possible combinations of wind speed and atmospheric stability. When this regime was applied to the dispersion of combustion products from fires containing toxic substances it was

found that the buoyancy of the fire plumes were sufficient to mitigate against a significant risk at ground level. The conclusion was that the risk was acceptable. However had high wind speed conditions (10 to 15 metres per second) been included, the plume would have remained close to the ground and the risks would have been found to be significant.

Pitfall - The hazard analysis was not complete. The representative set of events did not include the most hazardous conditions and scenarios. A sensitivity analysis had not been carried out.

There are a number of different methods for expressing risk; the choice of an appropriate method in any particular assessment will be influenced by:

- The scope of the assessment;
- The assessment approach adopted;
- The level of detail needed for the assessment, i.e. whether a qualitative, semi-quantitative or quantitative approach has been chosen.

The method chosen for risk estimation should also provide results in a form that enhances the understanding of the nature of the risk and how it can be controlled, and therefore some thought should also be given to the presentation of the risk results.

A number of different risk measures are commonly used to combine information on likelihood and magnitude of loss or injury. Three of the most commonly used types of measure are (AIChemE, 1989): risk indices, which are single numbers or tabulations that provide simple presentations of the risk; individual risk measures, which provide estimates of the risk to particular individuals; and societal risk measures which provide estimates of the risk to groups of people or society as a whole.

There are many uncertainties associated with the estimation of risk. A good understanding of the uncertainties and their causes is important to ensure effective interpretation of risk estimates. Consideration of uncertainty throughout the risk assessment process and approaches for dealing with it appropriately are discussed in section 3.12.

3.7.1 Risk indices

Risk indices are quantitative expressions of risk in the form of single numbers or tabulations, and they may be used in either an absolute or a relative sense. A number of different risk indices exist including: the Fatal Accident Rate; the Individual Hazard Index (IHI); the Average Rate of Death; the Equivalent Social Cost Index; the Mortality Index and the Economic Index. In addition, the NII Safety Assessment Principles (HSE, 1992b) use indices of radioactive release levels as a surrogate for individual (and societal) risk. Perhaps the most commonly used risk index is the Fatal Accident Rate (FAR) which is the estimated number of fatalities per 10^8 exposure hours (roughly 1000 employee working lifetimes) (Lees, 1996). The FAR is a single number index that is directly proportional to the average individual risk. Limitations on the use of risk index measures are that there may not be absolute criteria for accepting or rejecting the risk, and they lack resolution and do not communicate the same information as individual or societal risk measures (AIChemE, 1989).

3.7.2 Individual risk measures

Individual risk is defined as the frequency at which an individual may be expected to sustain a given level of harm from the realisation of specified hazards (IChemE, 1985). It is possible to estimate individual risk either qualitatively, semi-quantitatively or quantitatively depending on the approach that is adopted for the risk assessment. If a qualitative risk assessment approach is adopted, the risk estimation step may be carried out implicitly as part of the risk evaluation process, rather than explicitly as a separate step.

Qualitative estimates of individual risk usually involve categorising the risk from each identified significant hazard as e.g. either high, medium or low, after consideration of both the likelihood and consequences of realisation of the hazard. Risk matrix approaches are sometimes used to assist in this qualitative estimation and risk ranking. Risk matrices typically comprise a square divided into a number of boxes, with each box representing a different underlying estimation of risk (Middleton and Franks, 2001). Typically three to five (possibly more) categories are used for the possible consequences and a similar number for the possible likelihoods. The consequence and likelihood categories can be defined either numerically or by a description, and it is important that the category indicators or descriptors are defined as explicitly as possible as outlined in sections 3.5 and 3.6. Each risk box in the matrix represents the combination of a particular level of likelihood and consequence, and can be assigned either a numerical or descriptive risk value (the risk estimate). If numerical consequence and likelihood category indicators are used, it is common to estimate the risk values as the product of the likelihood and consequence values, as a convenient way of ranking the risks. Care should be taken if such an approach is adopted as, for example, hazards of low severity and high likelihood will receive the same risk value as hazards with high severity and low likelihood. Although the risk values may be the same, the response to these different hazards in terms of priority for correction may be very different (St John Holt, 1999), and care therefore needs to be taken to ensure the method for estimating risk results in values or categories that can be interpreted appropriately.

Risk matrix approaches can also be used to derive semi-quantitative risk estimates if either the likelihood or consequence categories can be quantified, or given order of magnitude estimates. Semi-quantitative order of magnitude risk-estimates can also be determined using the LOPA technique (Franks, 2001). In this technique, the individual risk is estimated by taking account of the frequency of the initiating event (for particular consequences), and the probability of failure for each layer of protection. In order to calculate the individual risk to a specific exposed person at a given location, the risk contributions from each of the scenarios with the potential to affect the individual are summed. Alternatively, the various parameters can be combined within a matrix or decision table containing the number of independent layers of protection as one of the parameters.

There are a number of different risk measures that can be used to estimate individual risk quantitatively, depending what is most appropriate for the study. As outlined in section 3.2.3, individual risk can be estimated to various different individuals or hypothetical persons, for example, the most exposed individual, for groups of individuals at particular places, or for a typical or ‘average’ individual within the effect zone. Whichever hypothetical person is the subject of the risk assessment, the individual risk is equal to the sum of the frequencies from all events that produce the specified level of harm for the individual of interest. Commonly used quantitative individual risk measures include: individual risk contours, which show the geographical distribution of individual risk; maximum individual risk, i.e. the individual risk to the person exposed to the highest risk in an exposed population; average individual risk

(exposed population), i.e. the individual risk averaged over the entire exposed population; average individual risk (exposed hours/worked hours), i.e. the individual risk for an activity calculated for the duration of the activity, or averaged over the working day. Further details of these and other individual risk measures commonly used in the chemical process industries, including details of their calculation are given in AIChemE (1989).

3.7.3 Societal risk measures

Societal risk is defined as the relationship between frequency and the number of people suffering from a specified level of harm in a given population from the realisation of specified hazards (IChemE, 1985). Societal risk is most usually expressed as a cumulative frequency distribution, either as 'F-N curves' or tables. F-N curves are plots of the cumulative frequency of N or more people receiving the specified level of harm per year, against the number of people (N) receiving the specified level of harm. It is usual for F-N curves to be plotted on a logarithmic scale, and for the specified level of harm to be fatality, i.e. societal risk of death is calculated. The calculation of societal risk requires the same frequency and consequence information as the calculation of individual risk, as well as information about the number of people exposed, often in the form of population density and distribution information. The societal risk expectation value (the average number of people receiving the specified level of harm per year) is also often calculated by summing the products of all events and their associated frequencies. However, this single value loses the distinction between low frequency/high consequence and high frequency/low consequence events and conveys much less information than the two dimensional F-N curve which shows the whole spectrum of the risk.

3.8 EVALUATING THE RISK

Consideration of the questions 'risk of what, to what, from what?' should have enabled the scope of the assessment to be clearly defined. An appropriate approach for the risk assessment should have been adopted and followed to enable the identification of all significant hazards within the scope of the assessment, and the estimation of the likelihood of their realisation and of the severity of the consequences to allow an estimate of the risk to be determined. Once an estimate of the risk has been obtained, the question that then needs to be asked (and answered) is 'so what?'. That is, consideration needs to be given to whether the level of risk arising from workplace activities (taking into account existing control and mitigation measures and their effectiveness) is acceptable or whether more needs to be done to control or reduce the risk.

As stated in section 3.2.1, HSW (1974) and other relevant legislation effectively defines the legally acceptable level of risk in any situation. The general duty on all employers is to reduce risks SFAIRP (HSW, 1974). HSE considers that this will be achieved if risks are reduced ALARP, as it considers that duties to ensure health and safety SFAIRP and duties to reduce risks ALARP call for the same set of tests to be applied (HSE, 2002a). A number of other regulations also include more specific requirements relating to the acceptability of risk levels in particular sectors of industry or in specific situations, sometimes expressed in different ways but the underlying requirement to reduce risks ALARP is basically the same in each case. For example, the COMAH (1999) regulations require operators to take 'all measures necessary' to control the off-site risks; this is interpreted by HSE to mean that off-site risks must be reduced ALARP. The COMAH (1999) regulations also include the requirement for operators to 'demonstrate' that all measures necessary have been taken. This has implications for the level of detailed argument that is needed in making the ALARP decision; where a demonstration that

risks have been reduced ALARP is required, detailed justification for decisions relating to the reasonably practicability (or otherwise) of implementing risk reduction measures is needed.

Reducing risks ALARP means that if it is reasonably practicable to implement a risk reduction measure, it must be implemented. When evaluating risks, and considering the level of risk reduction achieved by existing control or mitigation measures, duty-holders should consider what more could be done, and why it is not being done, i.e. whether or not it is reasonably practicable to implement any possible additional measures that are identified.

When considering what more could be done, i.e. how the risk could be reduced further, options should be considered according to the order outlined in the following 'hierarchy of risk control principles' (based on that given in MHSWR, 1999):

- Can the hazard/risk be eliminated by e.g. doing the work in a different way (that does not introduce new hazards);
- Can the hazard be substituted by something less hazardous, e.g. can a less hazardous substance, machine or process by used;
- Can the risk be controlled at source by the introduction of physical engineering controls;
- Can the risk be minimised by introducing procedural controls, such as safe systems of work, or using personal protective clothing and equipment; this option should only be used as a last resort.

The hierarchy reflects the fact that eliminating and controlling the risk by using physical engineering controls and safeguards is more reliable than relying on people to follow procedures etc. (MHSWR, 1999).

In determining whether the implementation of additional risk reduction measures is reasonably practicable or not, relevant case law has determined that an assessment should be made of the risk to be avoided, and of the sacrifice (in money, time and trouble) involved in taking measures to avoid that risk, and a comparison be made of the two (HSE, 2002a). For a measure to be not reasonably practicable, the degree of disproportion between costs (the sacrifice) and benefits (in terms of risk avoided) must be gross, i.e. the test of gross disproportion must be applied. That gross disproportion is required before a measure can be ruled out on the grounds of sacrifice can be interpreted as applying a bias on the side of safety (HSE, 2002a).

The ALARP principle is embodied in the HSE's 'Tolerability of Risk' (TOR) framework (HSE, 2001a), illustrated in figure 3.1. In addition to considering estimates of individual and societal risk when assessing and evaluating risks, where hazards also give rise to societal concerns, HSE may also require duty-holders to take these into account (HSE, 2001a). This is reflected in the scale of the TOR triangle which represents increasing levels of individual risk and societal concerns for a particular hazardous activity from the bottom of the triangle towards the top. Societal concerns arise in response to risks that can impact on society as a whole and have certain characteristics to which people are particularly averse. Societal risk is a subset of societal concerns, and represents that part of societal concerns that arises due to the occurrence of multiple fatalities in a single event. Societal concerns also arise in situations where (HSE, 2001a): it is difficult for people to estimate intuitively the actual threat; where exposure involves vulnerable groups, e.g. children; where the risks and benefits tend to be unevenly distributed, for example between groups of people with the result that some people bear more of the risks and others less, or through time so that less risk may be borne now and more by some future generation.

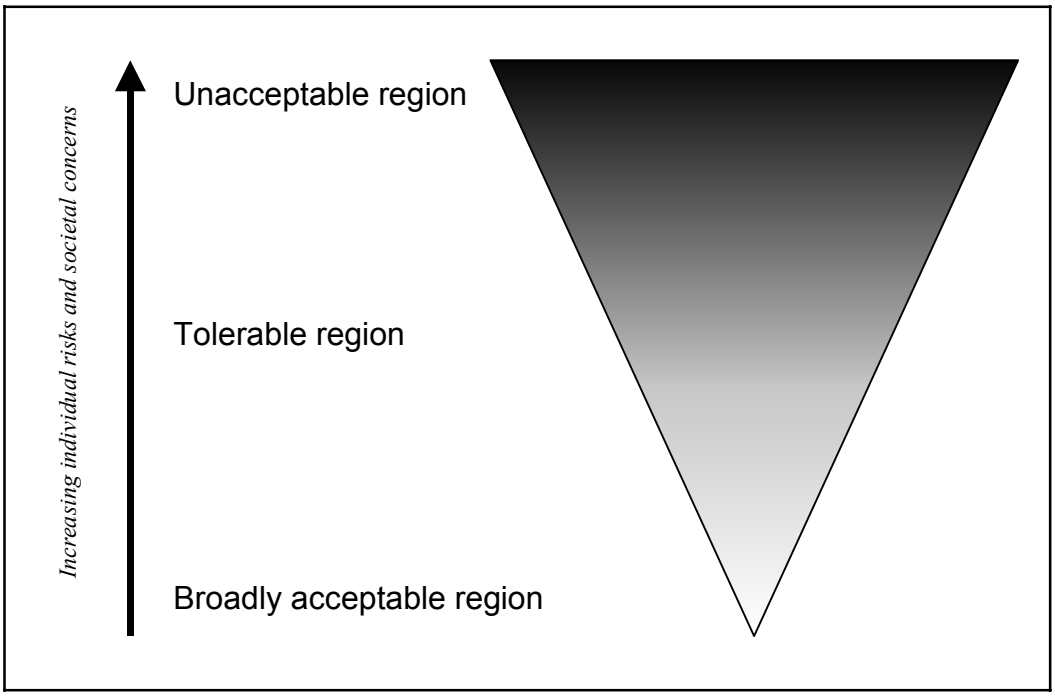

Figure 3.1 HSE Framework for the tolerability of risk (HSE, 2001a)

The TOR triangle includes three zones or regions, in which different approaches to evaluating risk and reaching decisions as to whether or not it is reasonably practicable to implement additional risk reduction measures are appropriate.

The dark zone at the top represents the 'unacceptable' region. For practical purposes, a particular risk falling into this region is regarded as unacceptable, whatever the level of benefits associated with the activity, and would be ruled out unless the activity or practice can be modified to reduce the degree of risk so that it falls into one of the two regions below (HSE, 2001a). The light zone at the bottom represents the 'broadly acceptable' region. The levels of risk characterising this region are comparable to those that people regard as insignificant or trivial in their daily lives, and are typical of the risk from activities that are inherently not very hazardous, or from hazardous activities that can be, and are readily controlled (HSE, 2001a). Within this region, additional risk control measures must be implemented if it is reasonably practicable to do so, but detailed arguments would not usually be needed to support the decision.

The zone between the unacceptable and broadly acceptable regions is the 'tolerable' region. Risks in this region are typical of the risks from activities that people are prepared to tolerate in order to secure certain benefits, such as employment, lower cost of production, personal convenience or the maintenance of general social infrastructure (HSE, 2001a). Within this region, additional risk control measures must be implemented if it is reasonably practicable to do so, and the greater the level of risk (i.e. the further up the TOR triangle the risk is situated), the greater the level of detail is needed in arguments to support the decision.

The TOR framework can in principle be applied to all hazards. When determining reasonably practicable measures for any particular hazard, whether the option being considered is

reasonably practicable or not depends in part on where the boundaries are set between the unacceptable, tolerable and broadly acceptable regions (HSE, 2001a). HSE has established indicative numerical individual risk criteria (applicable to estimates of the total risk from an activity) for the boundaries between the regions. These criteria, and the assumptions underlying their derivation are given in HSE (2001a). A basic societal risk criterion for the limit of tolerability from a single major industrial activity is also proposed. Developing criteria on the tolerability of risks giving rise to societal concerns has proved more difficult, but the criteria that have currently been adopted by HSE are also outlined in HSE (2001a).

Case Study 20: 'Inappropriate use of risk criteria'

A company's COMAH safety report tried to demonstrate that the site's risk was ALARP by comparing the risk from the identified hazardous scenarios to the tolerability criteria for individual and societal risk found in HSE (2001a). The risk from each scenario was found to be below the tolerable level and therefore judged to be broadly acceptable.

Pitfall - The criteria in HSE (2001a) are intended to be used for a site's overall individual risk or societal risk. They are therefore not suitable when trying to determine whether or not the risk from individual hazardous scenarios is ALARP. The risk from all hazardous scenarios should be summed before comparing to the criteria.

Other problems with RA - The societal risk criteria point in HSE (2001a) is for the unacceptable level. Risk below this point is in the ALARP region and therefore cannot necessarily be considered to be broadly acceptable.

In general therefore (for risks that are below the unacceptable region), in decisions relating to the reasonable practicability of implementing additional risk control measures, the greater the level of risk under consideration, the greater the degree of rigour and robustness (and hence the greater level of detail) HSE requires in arguments to show that risks have been reduced ALARP (HSE, 2002a). In addition, in comparing the risks and benefits of various options, and applying the test of gross disproportion, the greater the level of risk, the higher the proportion may be before being considered 'gross', but the disproportion must always be gross. This means that the gross disproportion factor that should be applied will increase as the level of risk increases (as you move up the TOR triangle).

In practice, for risks in the broadly acceptable region or towards the bottom of the tolerable region, for which simple qualitative risk assessments are likely to have been appropriate, HSE may accept the application of relevant good practice in an appropriate manner as a sufficient demonstration of part or all of a risk/sacrifice computation (HSE, 2002b). The working assumption is that the appropriate balance between risks and sacrifice will have been struck when the good practice was formally adopted, i.e. the reasonable practicability of the measures it entails has effectively already been demonstrated. However, the good practice must be appropriate to the activity considered, relevant to the risks from the undertaking and cover all the risks from the undertaking (HSE, 2002a). In judging and recognising good practice, HSE must be satisfied that it is correctly formulated, taking account of appropriate factors, and in an appropriate form as outlined in HSE (2002b). HSE considers authoritative sources of relevant

good practice to be those enshrined in prescriptive legislation, Approved Codes of Practice and guidance produced by Government (HSE, 2001a). Other sources that may be considered include standards produced by Standards-making organisations (e.g. BS, CEN, ISO, IEC) and guidance agreed by a body representing an industrial or occupational sector (e.g. trade federation, professional institution, sports governing body). It is important to remember that good practice may change over time, for example because of advances in technology or because of increased knowledge about the hazard, and what is appropriate good practice for a situation must therefore be kept under review.

For risks further up the tolerable region, where appropriate risk assessments are likely to have involved semi-quantitative or quantitative approaches, decisions relating to the reasonable practicability of implementing additional control measures are likely to require more detailed justification. It will be necessary to explicitly identify possible risk reduction options, and assess whether they are reasonably practicable. Where a number of options exist, all options, or combination of options that are reasonably practicable must be implemented. In measuring the risk to be reduced, and the sacrifice involved in measures to achieve that reduction, the starting point should be the present situation; if there are several options, they should each be considered as against the present situation (HSE, 2002a). The risks considered should be only those over which the duty-holder can exercise control or mitigate the consequences through the conduct of their undertaking. HSE starts with the expectation that controls in place must, at a minimum, achieve the standards of relevant good practice precautions, irrespective of specific risk estimates (HSE, 2001a). It is therefore important to consider whether there is any relevant good practice applicable to the situation under consideration before performing detailed cost benefit analyses to assess the reasonable practicability of implementing possible risk reduction measures. In carrying out comparisons of costs and benefits, only costs (in money, time or trouble) incurred by the duty-holder that are necessary and sufficient to implement the measures identified to reduce the risk can be taken into account. Details of the sorts of costs that typically need to be considered are given in HSE (2002a). Any additional benefits gained by the duty-holder as a result of the implementation of a safety measure should be offset against the costs incurred. If a measure results in a transfer of risk to other people (from the same hazard), the added risk to those people should be offset against the benefits the measure provides. It is also customary, when preparing formal cost benefit analyses to discount future costs and benefits to reflect the fact that people, on balance, prefer to have benefits now and pay for them later. Further details on approaches to discounting are given in HSE (2001a).

It is important to remember that what is reasonably practicable changes with time, for example advances in technology may make a higher standard of protection reasonably practicable to achieve. This means that, particularly for new plant or processes, comparison with standards in existing plant may not be appropriate. Reducing risks from an existing plant ALARP may still result in a level of residual risk which is higher than that which would be achieved by reducing the risks ALARP in a similar, new plant (HSE, 2002a). This means that it may not be reasonably practicable to apply retrospectively to existing plant what would be required to reduce risks ALARP for a new plant.

Case Study 21: 'No consideration of ALARP/further measures that could be taken'

This case involves the loading and off loading of tankers at a major dock terminal. The risk assessment was performed by the site safety officer and included a detailed description of the processes and procedures to follow when loading or off loading a tanker. A hazard checklist associated with these activities, a risk factor calculation table and a list of additional actions required to continually monitor the risk assessment were included. However the risk assessment did not show that all steps to stop accidents to people loading and unloading tankers had been taken and it was used to justify making no changes to the present system.

Pitfall - The risk assessment does not demonstrate that all the reasonably practicable steps to stop people falling off the tankers have been taken. Good practice has not been taken into account and examples from other industries have not been examined.

Other problems with RA - This case study was also an example of working backwards from the desired outcome and using a risk assessment to justify no changes (see Section 3). It also involved the use of a generic assessment without appropriate modification to take account of site or process specific considerations (see Section 3.1.3).

Case Study 22: 'Inappropriate use of CBA'

A cost benefit analysis was undertaken to determine whether expenditure on a gas detection system was reasonably practicable. Having determined the value of the potential losses, a detection system was described which was considerably more expensive. The study therefore concluded that the expenditure was not justified as the costs were grossly disproportionate to the benefits. However alternative cheaper gas detection systems, that were not grossly disproportionate to the potential losses, were not considered.

Pitfall - 'Gold plating' of the costs involved rendered the decision worthless. When carrying out a CBA, care should be taken to provide valid cost estimates and not only consider the most expensive option.

Case Study 23: 'Inappropriate use of the results of a CBA'

This case arose as a result of RI pressing for a level crossing near a busy holiday resort in the Southwest to be upgraded. The crossing was an automatic open crossing locally monitored, with flashing lights and warning signs but no barriers. There had been a history of car - train collisions and the potential for the crossing to become blocked with traffic queuing in the busier summer months. RI had requested that the crossing

be upgraded to one with barriers, following an accident the previous summer, but this had been resisted, with the reason that the risk did not warrant it. The risk assessment had illustrated that the risks to motorists were just on the limit of what is currently acceptable, therefore the protection measures should be increased unless it could be shown that it was not reasonably practicable to do so. However, the risk assessment used cost benefit analysis and showed that all three of the risk reduction measures considered were practicable.

Pitfall - The risk assessment shows that the risks were not reduced to ALARP as all three risk reduction measures considered in the CBA were shown to be reasonably practicable. Therefore the results of the CBA had not been used appropriately.

Attempts are sometimes made to use cost benefit analysis to argue for the removal of an existing risk control measure that achieves only a small reduction in the risk at significant cost. It is claimed that the cost benefit analysis shows that the cost of the existing measure is grossly disproportionate to the benefit achieved in terms of risk reduction, and hence that the continued implementation of the control measure is not reasonably practicable. However, such arguments are not acceptable as the removal of an existing control measure would result in an increase in the level of risk, and would mean that risks are no longer being controlled ALARP; the current implementation of the control measure effectively demonstrates its reasonable practicability. Arguments along such lines are sometimes termed 'Reverse ALARP' or 'Inverse ALARP' arguments.

Case Study 24: 'Reverse ALARP'

A public utility produced a new design for a piece of apparatus that formed part of a system. The new item posed a higher risk than the existing item. The utility tried to show that, compared to published tolerability criteria, the risks associated with the new design were still acceptable. However, when doing this the item was considered in isolation and not as part of the system. Also third party interference with the item was excluded from consideration, even though the new item was more vulnerable to this than the existing design.

Pitfall - This is an example of 'Reverse ALARP', it is not correct to use a risk assessment to justify a reduction in standards. As the old piece of apparatus had been installed and operational, it follows that this design was reasonably practicable to use and therefore the new design was not reducing risks ALARP.

Other problems with RA - Any risk assessment must clearly define the scope in terms of risk from what, risk to what, and risk of what. The criteria used to judge the result must be framed accordingly. The basis of the assessment did not match up with the basis of the criteria, and therefore any conclusion was worthless (see Section 3.1).

3.9 RECORDING THE FINDINGS

The findings of the risk assessment should be clearly documented. There is a legal requirement for employers who employ more than 5 people to document the significant findings of risk assessments (MHSWR, 1999). The record should represent an effective statement of the hazards and risks that leads management to take the relevant actions to protect health and safety. In relation to high hazard industries, where either Safety Reports or Safety Cases are required, these will contain a risk assessment that forms an integral part of any submission. In such regimes there is therefore a formal system for recording the findings of risk assessment; in addition, formal review dates are set.

The risk assessment findings should be documented in a way that is proportionate to the hazards and risks. The record of significant findings should include enough detail to show that a suitable and sufficient assessment has been made. It should therefore include:

- Details of the identified hazards associated with the work activity;
- Details of the severity of the possible consequences from realisation of the identified hazards, including identification of groups of people who may be affected, and any groups identified as especially at risk;
- Details of the precautions that are (or should be) in place to control the risks (with comments on their effectiveness);
- Improvements/further actions identified as necessary to control risks ALARP.

Details of those involved in carrying out the risk assessment, and others consulted should also be recorded, with the roles and responsibilities of all those involved clearly documented. In addition, the date of the assessment, and the date set for review of the assessment should also be recorded.

The record of the risk assessment may be kept in writing or recorded by other means (e.g. electronically) as long as it is retrievable and remains so even when, for example, the technology of electronic recording changes (MHSWR, 1999). Where appropriate, the risk assessment should be linked to other health and safety records or documents, such as the record of health and safety arrangements, and written health and safety policy (MHSWR, 1999) or relevant company rules, manuals and manufacturers' instructions and other documents and records describing procedures and safeguards.

In addition to recording the findings of the assessment, it is also important to record details of the assessment process itself, as outlined in section 3.4. The record should be kept for future reference and use. The risk assessment can help if an Inspector asks what precautions have been taken and can act as a reminder to keep an eye on particular hazards and precautions (HSE, 1999a). Documentation of the findings from and the details of the risk assessment can also help facilitate future reviews, and can be used as the basis for training and communication to employees. There is a requirement under MHSWR (1999) for employers to provide employees with comprehensible and relevant information on the risks to their health and safety identified by the assessment, and of the preventive and protective measures. The risk assessment will therefore help to identify any relevant information that should be provided to employees; the information should be pitched appropriately given the level of training, knowledge and experience of the employee, and provided in a form that takes account of any language difficulties or disabilities.

The risk assessment record should always be readily available, as should any supporting documentation which should also be up to date and comprehensive.

3.10 ENSURING TRANSPARENCY THROUGHOUT

The details of the assessment process should be recorded in a clear, auditable way, to ensure transparency of the approach. Any assumptions made throughout the risk assessment process should be clearly stated, and details of the date of the assessment and the people involved should be recorded.

The problem to be addressed should be clearly defined, particularly regarding the scope of the assessment and the aims and objectives should be clearly stated. The boundaries of the assessment should be stated and justified, and the approach to be taken specified.

The hazard identification stage should be considered systematically and comprehensively, using recognised techniques if a QRA approach is adopted, and carried out by suitably qualified and experienced assessors. Sources of data and any assumptions made in the consequence assessment and risk estimation stages should be documented and traceable, and any data used should be fit for purpose. The use of historical or predictive data should be clearly justified and acceptable. Any models used should be justified and clearly fit for purpose, and a thorough description should be given of all techniques used. Risk estimates should be expressed in understandable terms and the strengths and limitations of different risk measures used should be explained. The risk measure used should be the most suitable for the application.

Key uncertainties and assumptions should be identified and their effects considered and taken into account. The effects of uncertainty should be considered in detail, and any sensitivity or uncertainty analyses carried out should be detailed. Peer review of the assessment can also be helpful, particularly if a QRA approach is adopted to help avoid 'operator only' focus.

Details of the way in which the risks have been evaluated and a decision reached as to what additional actions (if any) are needed to ensure risks are controlled ALARP should be clearly outlined. Any criteria used to assist in the decision making process should be clearly stated and justified.

3.11 ENSURING APPROPRIATE CONSIDERATION OF HUMAN FACTORS

Human factors refers to environmental, organisational and job factors, and human and individual characteristics that influence behaviour at work in a way that can affect health and safety (HSE, 1999b). Accidents can, and often do occur as a result of human factors. As technical systems have become more reliable, the focus has turned to human causes of accidents, and it is estimated that 80% of accidents may be attributed, at least in part, to the actions or omissions of people (HSE, 1999b). It is therefore extremely important to ensure the potential for human factors contributions to the realisation of hazards is appropriately considered throughout the risk assessment process, so that suitable preventive and protective measures and controls can be put in place to ensure the risks are controlled to ALARP.

Human factors can be thought of as the influence of three inter-related factors (the job, the individual and the organisation) on human performance and health and safety related behaviour. The following guidance on how each of these factors can impact on health and safety related behaviour is extracted from HSE (1999b):

- **The job** - Tasks should be designed in accordance with ergonomic principles to take into account limitations and strengths in human performance. Matching the job to the person will

ensure that they are not overloaded and that the most effective contribution to the business results. *Physical match* includes the design of the whole workplace and working environment. *Mental match* involves the individual's information and decision-making requirements, as well as their perception of the tasks and risks. Mismatches between job requirements and people's capabilities provide the potential for human error.

- **The individual** - People bring to their job attitudes, skills, habits and personalities which can be strengths or weaknesses depending on the task demands. Individual characteristics influence behaviour in complex and significant ways. Their effects on task performance may be negative and may not always be mitigated by job design. Some characteristics such as personality are fixed and cannot be changed. Others, such as skills and attitudes, may be changed and enhanced.
- **The organisation** - Organisational factors have the greatest influence on individual and group behaviour, yet are often overlooked during the design of work and during investigation of accidents and incidents. Organisations need to establish their own positive health and safety culture. The culture needs to promote employee involvement and commitment at all levels, emphasising that deviation from established health and safety standards is not acceptable.

It is important to ensure that there is adequate consideration of human factors at each stage of the risk assessment process. In carrying out the assessment, it may be helpful at each stage to consider explicitly the task, the individual and the organisational factors.

At the hazard identification stage, when considering risks to health and safety, it is important to consider the full range of physical, chemical, biological and psychological hazards, and consider both immediate and longer-term effects. It is important to remember that people not only suffer as a result of hazards at work but that they can also contribute to the hazards themselves, for example through human error or failure to follow procedures. Consideration of attitudes to risk, safety culture, ergonomic design, and human error are all relevant to the hazard identification stage (HSE, 1999b) as is consideration of the knowledge, skills and experience of staff engaged on the work activity.

Assessing the consequences of the realisation of the hazards involves an assessment of who might be harmed, how the harm might arise, and its physical effects and severity. This is combined with an assessment of the likelihood that the harm will be realised. In considering how harm might arise, and how likely it is to occur you should not assume that people will always follow set procedures, and you should allow for the occurrence of human errors and violations (HSE, 1999b). It is also important to consider people's behaviour during abnormal and emergency situations as well as during planned, routine tasks.

The estimates of likelihood and consequence for the identified hazardous events must be combined to produce an estimate of the risk itself. In some situations it may be desirable to quantify the risks arising from human failures. A number of methods exist for quantifying the contribution made by human action or inaction to the overall risk from a system. These approaches, known as Human Reliability Assessment (HRA) include the process of task analysis which helps with the identification of all points in a sequence of operations at which incorrect human action, or the failure to act, may lead to adverse consequences (Hurst, 1998). HRA techniques assign a degree of probability on a numerical scale to each event, and by aggregating these, arrive at an overall figure for the probability of human error; HRA approaches therefore enable human error to be 'factored into' the risk assessment such that both engineering reliability and human error are considered together (Hurst, 1998). Further details of the range of HRA techniques that exist, including details of other relevant references are given

in Hurst (1998) and HSE (1999b). If such methods are used, it is important not to assume that they are a substitute for upgrading control measures against human failures (HSE, 1999b).

In evaluating the risk and deciding whether existing precautions are adequate or if more should be done to control or reduce the risk, you should not rely on individuals to control a hazard. As with all controls, the hierarchy of risk control principles outlined in section 3.8 should be followed. That is, it is important to first consider whether it is possible to eliminate the hazard at source and only rely on individual actions as a last resort. If possible, ways of making the situation more 'error tolerant' should be considered, for example by improving the ways in which people can detect and correct errors and mistakes before they lead to adverse consequences (HSE, 1999b).

It is likely that the risk assessment will have resulted in some recommendations for improvements and further actions to control and reduce risk. It is important that consideration of human factors continues at this stage and is incorporated into the design of appropriate new or additional risk control precautions. Further guidance on how to incorporate human factors into job design is included in HSE (1999b).

3.12 HANDLING UNCERTAINTY

Each stage in the risk estimation process involves an element of uncertainty. An understanding of uncertainties and their causes is therefore required to interpret risk estimates effectively. It is important that uncertainty is accounted for in the reporting of risk estimates so that the credibility of the risk assessment process is not undermined.

A substantial amount of literature on uncertainty analysis has been developed, a large part of which has arisen out of the work on nuclear 'probabilistic safety analyses'. Useful references are provided in AIChemE (1989). It is generally accepted that there are three main sources of uncertainty:

- Model uncertainty;
- Data or knowledge uncertainty;
- General quality uncertainties.

Model uncertainty reflects the weaknesses, deficiencies and inadequacies intrinsic to any model, and is a measure of the degree to which a model fails to represent reality. Data uncertainty arises from the fact that data is almost always incomplete in some respect, and gaps need to be filled through estimation, inference or expert judgement. General quality uncertainty relates to the completeness and comprehensiveness of the assessment; uncertainty arises from not knowing the combined risk contributions from those incidents that have been omitted.

Examining each stage of the risk estimation process in turn, uncertainty can arise as a result of:

1) Descriptions of processes, procedures and other site information being incorrect, out of date, or not representing actual operation;
2) Failure to identify all hazards and routes to failure;
3) Inappropriate selection of models for estimating the magnitude and severity of the consequences, e.g. selection of models with incorrect or inadequate physical basis, or with inadequate validation;
4) Uncertainties in consequence model input data;

5) Inappropriate or inadequate use of techniques for modelling event likelihoods;
6) Uncertainties in the failure rate data, e.g. inaccurate, incomplete or inappropriate data.

The first two sources of uncertainty in the risk estimation process are forms of general quality uncertainties. Such uncertainties cannot readily be quantified, and can be handled most effectively by ensuring they are minimised; this can be achieved by ensuring appropriate techniques for hazard identification are adopted, utilising all available relevant information and involving a team of competent people with appropriate knowledge, skills, expertise and experience.

The third and fifth sources of uncertainty in the risk estimation process are forms of model uncertainty. Model uncertainty is also difficult to quantify, but can be handled effectively by ensuring it is minimised and through the use of sensitivity analysis. Model uncertainty can be minimised by ensuring that the most appropriate models are chosen in any situation, by ensuring that the people making the choice have the appropriate knowledge and experience. Sensitivity analysis can be used to identify the contribution model uncertainty makes to the overall uncertainty in the risk estimate (AIChemE, 1989).

The fourth and sixth sources of uncertainty in the risk estimation process are forms of data uncertainty. Such uncertainties can be analysed quantitatively through the use of mathematical methods that combine the theories of probability and statistics (AIChemE, 1989). For example, statistical methods can be used to determine the following quantitative estimates of the data uncertainty: an expected value with upper and lower bounds; an expected value with a standard deviation; a probability distribution function; or an expected value with a confidence interval. Further details of how to derive such estimates of uncertainty is given in AIChemE (1989) which includes further references.

The uncertainty in the overall risk estimate arises from the combination of the uncertainties at each stage of the process. Various methods exist for combining uncertainties that have been quantified; these methods are briefly described in AIChemE (1989) and further relevant references are provided. Sensitivity analysis can also be a useful tool to identify which models, assumptions and data are important to the final risk estimate; further details of how to carry out sensitivity analyses are provided in AIChemE (1989).

Many of the approaches for handling uncertainty outlined above will only be applicable if a quantified approach to risk assessment has been adopted. Alternative approaches to handling uncertainty in a more general way that are applicable in situations where a qualitative approach to risk assessment is appropriate, as well as in situations where a quantitative approach is adopted, also exist. Examples of such approaches, are where the overall risk estimate is defined as being either a ‘best estimate’, a ‘cautious best estimate’ or a ‘worst case estimate’. At each stage of the risk assessment process, depending on which approach is adopted, where assumptions have to be made either the best estimate (based on judgement) or a modification to the best estimate that is judged to err slightly on the side of caution, or the worst case estimate of the outcome or parameter value should be used. If such an approach is adopted, the basis for the overall risk estimate should be clearly stated.

Whatever type of risk assessment approach is adopted (i.e. whether qualitative or quantitative), it is important to ensure that appropriate account is taken of the uncertainty in the risk estimate at the risk evaluation stage. As the degree of uncertainty increases, so should the extent to which you ‘err on the side of safety’ in making decisions as to the measures that are needed to ensure the risk is adequately controlled.

An extension of this idea is the 'precautionary principle' which describes the philosophy that should be adopted for addressing hazards subject to high scientific uncertainty, and rules out lack of scientific certainty as a reason for not taking preventive action (HSE, 2001a): 'where there are threats of serious or irreversible damage, lack of full scientific certainty shall not be used as a reason for postponing cost effective measures to prevent degradation'. HSE policy (HSE, 2001a) is that the precautionary principle should be invoked where: there is good reason, based on empirical evidence or plausible causal hypothesis to believe that serious harm might occur, even if the likelihood of harm is remote; and the scientific information gathered at this stage of consequences and likelihood reveals such uncertainty that it is impossible to evaluate the possible outcomes with sufficient confidence to move to the next stage of the risk assessment process.

3.13 ACTING UPON THE FINDINGS

It is essential that action is taken as a result of the findings of the risk assessment. Risk assessment should never be just a paper exercise; the entire process will have been a waste of time if the findings are merely noted, but no action taken as a result. Risk assessment should be an integral part of the company's safety management system, and should lead to the development of plans for improvement, and contribute to the robustness of the overall system.

The assessment will almost inevitably result in recommendations for improvements and further actions to control and reduce risk; any identified new or additional risk reduction measures or risk control systems must be implemented. Suitable risk control systems should be implemented that are proportionate to the hazards and risks. For example (HSE, 1997), the control of minor hazards can be dealt with by a number of simply stated general rules. The control of more hazardous activities may need more detailed workplace precautions and risk control systems, and the control of high hazard activities may demand detailed workplace precautions and a risk control system that needs to be strictly followed such as a permit to work system. Any procedures developed should be clearly documented and adequate; the balance of resources devoted to the various risk control systems should reflect the hazard profile of the business.

Case Study 25: 'Not doing anything with the results of the assessment'

A company had carried out risk assessments and had kept records of all previous risk assessments stacked up on shelves. However, the risk assessments were found to be inadequate as they did not identify suitable control measures (when such measures were required) and, in addition, the risks were difficult to prioritise.

Pitfall - A risk assessment is not a paper exercise. Risk assessments should be living documents and their results should be used to control risk by identifying additional control measures that could be implemented.

Other problems with RA - Although the people carrying out the risk assessment had been trained in the company's risk assessment process, they were not aware of the industry standards which they should be using to evaluate the risk (see Section 3.3).

Case Study 26: 'Not linking hazards with risk controls'

This case concerns the risk assessment section of the Railway Safety Case (RSC) for the operation of track maintenance vehicles. In this particular case it was a very small operation, and the operator had inappropriately adapted a generic system-wide QRA to do their risk assessment. The risk assessment did not reflect the local situation; it was not specific to the particular operation and was inappropriate in its scale and extent. As a result the findings of the risk assessment were not clearly presented, and identified hazards were not linked with risk controls.

Pitfall - The findings of the risk assessment were presented in a format which rendered them useless - a long list of risks and control measures which were not prioritised. No clear link was made between the identified hazards and risks and the control measures.

Other problems with RA - Many hazards had not been accurately identified, which led to confusion when identifying risks and the double counting of certain events. This undermines any attempt at quantification and prioritisation of risks (see Section 3.4). In addition, a generic risk assessment had been inappropriately adapted so that the risk assessment did not adequately reflect the local situation (see Section 3.1.3).

In practice, the implementation of risk assessment findings and recommendations may involve several stages, similar to those outlined by Bateman (2001). It is likely that there will first be a review of the assessment findings and recommendations by an independent team, probably involving senior management. Following this review, the identified actions will need to be costed and prioritised, leading to the development of an action plan. The action plan should include the identification of individuals with clearly allocated responsibility for each element of the plan and define timescales for completion of actions. It is good practice to build in to the plan the task of following up on the recommendations to ensure they have been implemented, and also to examine their effectiveness, and consider whether any unexpected risks have been inadvertently created. The assessment record should be annotated or revised to take account of any changes made. If recommendations have not been implemented, it would be necessary to refer to senior management for them to take action to identify and overcome the obstacles to progress.

The findings of the risk assessment, and any new working procedures or risk control systems that are implemented as a result must also be communicated to employees, as outlined in section 3.9. In devising new procedures or risk control systems, it is important to ensure that there is adequate consideration of human factors, as outlined in section 3.11, and where possible, 'error tolerant' systems and procedures should be devised. It is also important to recognise that the introduction of any changes will need to be properly managed. For example, training of staff may be required when new procedures or systems are introduced, and there may initially be a need for increased supervision.

3.14 REVIEWING THE ASSESSMENT

Risk assessment should not be a one-off activity, but should be part of the process of continuous improvement. It is therefore important that assessments are reviewed at appropriate intervals, and updated as necessary. Reviews of risk assessments should form part of standard management practice (MHSWR, 1999). In relation to high hazard industries, where either Safety Reports or Safety Cases are required, these will contain a risk assessment that forms an integral part of any submission. In such regimes there are formal systems for review in response to significant changes in operation, or after a certain length of time.

Risk assessments should be reviewed when there is any reason to suspect they may no longer be valid, for example: following an accident, incident or near-miss; as significant new information becomes available; or when there have been significant changes to working procedures. Such reviews are unplanned reviews that should be triggered by significant changes. Relevant new information may become available from various sources, including for example new staff with different expertise and experience, new manufacturers and suppliers of raw materials and equipment, and as a result of technological developments. Other information of relevance to the validity of risk assessments and assumptions within them will come from monitoring by way of inspection and from routine measurements, e.g. air quality measurements and medical surveillance (St John Holt, 1999).

Changes may be made in response to various factors including company policies, economic and market pressures and technological developments. Significant changes to working procedures may include such things as the use of new machines, substances or processes. It is not necessary to review and amend the assessment for each trivial change, but fresh assessments will be required when there are significant changes, and the risk assessment should be updated and modified accordingly. If a new job introduces significant new hazards of its own, these should be considered in their own right, and appropriate assessments carried out.

There should also be a well defined, formal system in place for regular review and update of risk assessments, i.e. planned reviews, in addition to unplanned reviews triggered by significant changes. It is good practice to review all assessments periodically to ensure they are still valid (MHSWR, 1999), to check that associated precautions are still working effectively and to identify any ‘creeping changes’. The time between such reviews should relate to the extent and nature of the risks involved, and the degree of change likely in the work activity. The process of conducting regular reviews may be aided by the application of document control systems, for example of the type used to achieve compliance with ISO 9000 and similar standards (Bateman, 2001).

4 BRIEF REVIEW OF GUIDANCE FOR HSE INSPECTORS ON RISK ASSESSMENT

An important aim for HSE is the transparency of the decision making process. As such, documents on how decisions are made, are either published or made publically available on the Internet. As described in section 3, HSE's "Reducing Risks, Protecting People" publication (HSE, 2001a) describes HSE's overall decision making process. This has recently been supplemented with guidance on ALARP decisions (HSE, 2002a and b). In addition, there is specific guidance for Inspectors on risk assessment relating to the assessment of safety cases in the nuclear, offshore, railway and chemical sectors. The guidance is generally in the form of principles or criteria to be used by Inspectors to assist them when assessing risk assessment aspects of safety cases. A high level review of the guidance is presented in this section.

Nuclear - Licencees are responsible for the safety of their plants and must demonstrate adequacy of safety operations by producing and submitting safety cases to cover all phases of a plant's life cycle. The Nuclear Installations Inspectorate (NII) agree, approve or consent to the operation or procedure based on their assessment of the supporting safety case. To guide inspectors in making judgements on the adequacy of the safety case, NII has set out Safety Assessment Principles (SAPs) that are published (HSE, 1992b) so that nuclear licencees are aware of the principles against which they will be judged. SAPs were first published in 1979 for nuclear power plant (HSE, 1979) and in 1983 for nuclear chemical plant (HSE, 1983). The current SAPs were published in 1992 (HSE, 1992b) and cover all nuclear installations and translate the concepts of tolerability of risk into criteria appropriate for such installations. The Tolerability of Risk from Nuclear Power Stations (TOR) document (HSE, 1992a) was produced to explain how NII regulate and what the SAPs meant in terms of risk to people, as a response to a recommendation of the Chairman of the Sizewell B Inquiry.

The SAPs contain some 330 principles, the vast majority of which relate to engineering or operational good practices as well as dose and risk criteria. The overall requirement is that risks should be reduced ALARP: the criteria are for the guidance of inspectors and are not design targets for licensees. To support SAPs, more detailed internal guidance in the form of Technical Assessment Guides has been produced, some of which are available on the HSE website.

Offshore - The publication 'Assessment principles for offshore safety cases' (HSE, 1998) has been produced by HSE's Offshore Safety Division (OSD) to ensure a greater understanding of the principles against which HSE assessors evaluate safety cases. A key part of an offshore safety case is demonstrating that all hazards with the potential to cause a major accident have been identified, their risk evaluated, and that measures have been or will be taken to reduce the risks to people affected by those hazards to the lowest level that is reasonably practicable (ALARP). Acceptable safety cases will demonstrate that a structured approach has been taken which:

a) identifies all major accident hazards. The identification methods should be appropriate to the magnitude of the hazards involved and a systematic process should be used to identify where a combination or sequence of events could lead to a major accident;

b) evaluates the risks from the identified major accident hazards. Any criteria for eliminating less significant risks should be explained and in deciding what is reasonably practicable, relevant good practice and sound engineering principles should be taken into account. In

addition, human factors need to be accounted for and safety critical tasks should be analysed to determine the demands on personnel;

c) describes how any quantified risk assessment has been used, and how uncertainties have been taken into account;

d) identifies and describes the implementation of the risk reduction measures. The reasoning for or against the choice of risk reduction measures to be implemented should be clear.

Railways - The Railways (Safety Case) Regulations 2000 (HSE, 2001b) require railway operators to prepare and submit safety cases to HSE. The safety case needs to provide sufficient specific information to describe the nature and extent of the operation and must demonstrate that the operator has undertaken adequate risk assessment for all operations, identified risk control measures, and has systems in place to ensure the measures are implemented and maintained. The Assessment Criteria for Railway Safety Cases (HSE, 2002c) provide guidance to help HSE inspectors form judgements about the completeness of a safety case and the adequacy of the arguments presented to show that risks have been properly controlled. The criteria represent what is currently accepted as good practice. The criteria have been published to make them widely known throughout the railway industry and help develop a common understanding of the requirements for producing safety cases and to make the process by which HSE assess them transparent. The criteria used to assess risk assessment aspects of Railway Safety Cases are:

1. the Safety Case should give details of the duty holder's organisation and arrangements for identification of hazards and assessment of risk;
2. the Safety Case should justify the methodologies used for the identification of hazards and assessment of risk with particular reference to any assumptions and data used, together with the methods of calculation;
3. the Safety Case should describe the significant findings of the risk assessments and demonstrate that the control measures are adequate to control the risk to a level as low as reasonably practicable;
4. the Safety Case should describe the duty holder's arrangements to review risk assessments in the light of new information, new technology, incidents, or other changes that may affect risks, and to ensure that the risk assessments remain valid.

Chemical - The Control of Major Accident Hazards (COMAH) Regulations 1999 (COMAH, 1999) are enforced by a joint Competent Authority (CA) comprising HSE and the Environment Agency (or the Scottish Environment Protection Agency in Scotland). The Safety Report Assessment Manual (HSE, 2002d) documents the CA's arrangements for the handling and assessment of safety reports submitted in accordance with the COMAH regulations. It sets down the CA's policies, procedures and guidance for the handling and assessment of safety reports. The assessment involves the exercise of professional judgement by Inspectors and the manual provides a framework within which these judgements are made. The manual is intended as a practical tool for Inspectors, to help achieve consistency in the approach to safety report assessment and is publicly available via the HSE Internet Web site.

A safety report has to contain certain information, which relates to the major accident hazards and how major accidents are prevented or how the consequences of such an accident are limited. The information provided has to be sufficient to meet the purposes of the report and to help assessors gather this information, the CA has drawn up assessment criteria. There are

about 130 criteria set out in 6 groups. The predictive criteria deal with the identification of major accident hazards and risk analysis. They cover:

- principles of risk assessment and the use of appropriate data;
- identification of major hazards and accident scenarios;
- likelihood of a particular major accident scenario or the conditions under which they occur including initiating and event sequences; and
- consequence assessment.

The safety report as a whole should enable a view to be taken on the suitability and sufficiency of the risk assessment for drawing soundly based conclusions. It should be clear that the operator's approach to demonstrating compliance with the 'all necessary measures' requirement, is fit for purpose.

5 SUMMARY OF MAIN FINDINGS

A comprehensive literature search was carried out to attempt to identify any published critiques of both general and specific risk assessment methodologies, or any references that include descriptions of risk assessment pitfalls. While there exists a large body of published material on the general topic of risk assessment and its application, only three references were found that include material of relevance to the critical review of methodologies or information on risk assessment pitfalls. A brief review of the identified references has been included in section 1.3.

A large number of examples of inadequate industry risk assessments that illustrate common pitfalls in the application of risk assessment were identified by collating experience from HSE operational divisions, and from HSL's own experience of carrying out support and research work for HSE. Twenty six of the examples identified have been written up as case study examples to illustrate the most common pitfalls. The illustrative case study examples have been included throughout section 3 of the report which outlines good practice in the use of risk assessment. The identified pitfalls are as follows:

- Carrying out a risk assessment to attempt to justify a decision that has already been made;
- Using a generic assessment when a site-specific assessment is needed;
- Carrying out a detailed quantified risk assessment without first considering whether any relevant good practice was applicable, or when relevant good practice exists;
- Carrying out a risk assessment using inappropriate good practice;
- Making decisions on the basis of individual risk estimates when societal risk is the appropriate measure;
- Only considering the risk from one activity;
- Dividing the time spent on the hazardous activity between several individuals - the 'salami slicing' approach to risk estimation;
- Not involving a team of people in the assessment or not including employees with practical knowledge of the process/activity being assessed;
- Ineffective use of consultants;
- Failure to identify all hazards associated with a particular activity;
- Failure to fully consider all possible outcomes;
- Inappropriate use of data;
- Inappropriate definition of a representative sample of events;
- Inappropriate use of risk criteria;
- No consideration of ALARP or further measures that could be taken;
- Inappropriate use of cost benefit analysis;
- Using 'Reverse ALARP' arguments (i.e. using cost benefit analysis to attempt to argue that it is acceptable to reduce existing safety standards);
- Not doing anything with the results of the assessment;
- Not linking hazards with risk controls.

A review of internal HSE guidance for Inspectors was carried out and it was found that there is relatively little published internal guidance on how to assess risk assessments. HSE has recently produced guidance on ALARP decisions (HSE, 2002a and b) which complements the general guidance in HSE (2001a). In addition, the only other published guidance found relates to the assessment of safety cases in the nuclear, offshore, railway and chemical sectors. A high level review of the guidance on the assessment of risk assessment aspects of safety cases has been carried out in section 4 of the report.

6 REFERENCES

Adams, J (1998). Risk. University College London Press, Fourth impression 1998. ISBN 1-85728-068-7.

AIChemE (1989). Guidelines for Chemical Process Quantitative Risk Analysis. CCPS, New York.

Bateman, M (2001). Tolley's Practical Risk Assessment Handbook. Butterworths Tolley, Reed Elsevier (UK) Ltd, Second Edition 2001. ISBN 0 75450 749-1.

Beale, CJ (2001). Appropriate risk assessment: Selecting the best methodology for solving the problem. Paper presented at the IBC conference on 'The application and use of risk assessment', 8-9 October, 2001, The Café Royal, London.

BS 8444 (1996). British Standard BS 8444 Risk Management, Part 3 - Guide to risk analysis of technological systems. BS 8444: Part 3: 1996 IEC 300-3-9: 1995.

CDM (2001). Managing health and safety in construction: A guide to the Construction (Design and Management) Regulations 1994, Approved code of practice and guidance. HSG 224. ISBN 0 7176 21391.

COMAH (1999). A guide to the Control of Major Accident Hazard Regulations 1999. L111. ISBN 0 7176 16045.

COSHH (1999). The Control of Substances Hazardous to Health Regulations 1999. SI 1999/437. ISBN 0 11 082087 8.

Franks, A (2001). Interim Report prepared for HSL: Lines of Defence/Layers of Protection Literature Review. Amey Vectra report 300-2017 R01, revision A, December 2001.

Franks, A, Hughes, G and Hanif, S (2000). A Comparison of Accident Experience with Quantitative Risk Assessment (QRA) Methodology. HSE CRR 293/2000.

Glossop, M, Ioannides, A and Gould, J (2000). Review of Hazard Identification Techniques. HSL report RAS/00/02.

HSE (1979). Safety Assessment Principles for Nuclear Power Reactors. ISBN 0 11 883642 0.

HSE (1983). Safety Assessment Principles for Nuclear Chemical Plant. ISBN 017176 01536.

HSE (1989). Risk criteria for land-use planning in the vicinity of major industrial hazards. HSE C30 3/89. ISBN 0 11 885491 7.

HSE (1992a). The Tolerability of Risk from Nuclear Power Stations. ISBN 0 11 886268 1.

HSE (1992b). Safety assessment principles for nuclear plants. ISBN 0 11 882043 5.

HSE (1997). Successful health and safety management. HSG65, second edition HSE 12/98. ISBN 0 7176 1276 7.

HSE (1998). Assessment principles for offshore safety cases, HSG 181. ISBN 0 7176 12384.

HSE (1999a). Five steps to risk assessment. HSE INDG163(rev1) 5/99 C2620.

HSE (1999b). Reducing error and influencing behaviour. HSG48. ISBN 0 7176 2452 8.

HSE (2001a). Reducing Risks Protecting People - HSE's decision making process. HSE 12/01. ISBN 0 7176 2151 0.

HSE (2001b) Railways (Safety Case) Regulations 2000 including 2001 amendments, Guidance on Regulations L52. ISBN 0 7176 21278.

HSE (2002a). Principles and guidelines to assist HSE in its judgements that duty-holders have reduced risk as low as reasonably practicable.

HSE (2002b). Assessing compliance with the law in individual cases and the use of good practice.

HSE (2002c). Safety Case Assessment Criteria. www.hse.gov.uk/railway/criteria/section4.htm

HSE (2002d). COMAH Safety Report Assessment Manual
www.hse.gov.uk/hid/land/comah2/index.htm

HSW (1974). Health and Safety at Work etc. Act, 1974. In: Redgrave's Health and Safety, J Hendy and M Ford. Butterworths, third edition 1998. ISBN 0 406 89969 X.

Hurst, NW (1998). Risk Assessment: The Human Dimension. The Royal Society of Chemistry. ISBN 0 85404 554 6.

IChemE (1985). Nomenclature for hazard and risk assessment in the process industries. ISBN 0 85295 184 1.

IChemE (1996). Risk assessment in the process industries. Second edition, edited by R Turney and R Pitblado. ISBN 0 85295 323 2.

Lees, FP (1996). Loss Prevention in the Process Industries. Butterworth Heineman, second edition, 3 volumes. ISBN 0 7506 1547 8.

MHSWR (1999). Management of Health and Safety at Work Regulations 1999, Approved Code of Practice and Guidance. HSC 2000, HSE L21. ISBN 0 7176 2488 9.

Middleton, M and Franks A (2001). Using risk matrices. The Chemical Engineer, pp 34-37, September 2001.

OSC (1998). A guide to the Offshore Installations (Safety Case) Regulations 1992. L30. ISBN 0 7176 1165 5.

Railway Safety (2001). Profile of Safety Risk on Railtrack Plc. controlled infrastructure. Report No. SP-RSK-3.1.3.11.

Simpson, GC (1996). Toward a rational approach to risk assessment. Mining technology, January 1996, volume 78 no. 893 pp 19-23.

St John Holt, A (1999). Principles of health and safety at work. IOSH services, fifth edition (revised). ISBN 0 9013 5724 3.

Stowe, J (2001). Managing health and safety risk assessments effectively. The Stationery Office, 2001. ISBN 0 11 702825 8.

Wang, JX and Roush, ML (2000). What every engineer should know about Risk Engineering and Management. Marcel Dekker, Inc., 2000. ISBN 0-8247-9301-3.

Williams, MMR and Thorne, MC (1997). The estimation of failure rates for low probability events. Progress in Nuclear Energy, Vol. 31, no. 4, pp 373 - 476.

Worsell, N and Wilday, J (1997). The application of risk assessment to machinery safety: Review of risk ranking and risk estimation techniques. HSL Report, RAS/97/12.

Wright, S, O'Hara, R and Dickety, N (2002). Risk Assessment: The Extent of Adoption by Industry. HSL Report, RAS/02/12.

Printed and published by the Health and Safety Executive
C1.25 08/03

MAIL ORDER
HSE priced and free
publications are
available from:
HSE Books
PO Box 1999
Sudbury
Suffolk CO10 2WA
Tel: 01787 881165
Fax: 01787 313995
Website: www.hsebooks.co.uk

RETAIL
HSE priced publications
are available from booksellers

HEALTH AND SAFETY INFORMATION
HSE Infoline
Tel: 08701 545500
Fax: 02920 859260
e-mail: hseinformationservices@natbrit.com
or write to:
HSE Information Services
Caerphilly Business Park
Caerphilly CF83 3GG

HSE website: www.hse.gov.uk

RR 151

£15.00

ISBN 0-7176-2732-2

9 780717 627325

Good practice and pitfalls in risk assessment

HSE BOOKS