9 Engineering a Safety Culture

The Scope of the Chapter

Few phrases occur more frequently in discussions about hazardous technologies than *safety culture*. Few things are so sought after and yet so little understood. However, should it be thought that the current preoccupation with safety culture is just another passing fad, consider the following facts. Commercial aviation is an industry that possesses an unusual degree of uniformity worldwide. Airlines across the globe fly much the same types of aircraft in comparable conditions. Flight crews, air traffic controllers and maintenance engineers are trained and licensed to very similar standards. Yet, in 1995, the risks to passengers (the probability of becoming involved in an accident with at least one fatality) varied by a factor of 42 across the world's air carriers, from a 1 in 260 000 chance of death or injury in the worst cases to a 1 in 11 000 000 probability in the best cases.[1] While factors such as national and company resources will play their part, there can be little doubt that differences in safety culture are likely to contribute the lion's share to this enormous variation.

We first encountered the term 'safety culture' in Chapter 2 when making the distinction between pathological, bureaucratic and generative organizations. It cropped up again in Chapter 6 in regard to the motive forces that drive an organization towards a state of maximum resistance to its operational hazards. The present chapter focuses mainly on three questions: What is an organizational culture? What are the main ingredients of a safety culture? And, most importantly, how can it be engineered? The term 'engineered' is deliberate. But it is not meant in the traditional sense of developing more sophisticated gadgetry. Rather, we will be discussing the application of *social engineering*.

This book has sought to argue that most of the effective solutions to human performance problems are more the province of the technical manager (and the regulator) than the psychologist since they concern the conditions under which people work rather than the human condition itself. The main message of this chapter is that the same general principle also applies to the acquisition of an effective safety culture (hereafter, the phrase 'safety culture' will be taken to mean an effective safety culture). Whereas national cultures arise largely out of shared values, organizational cultures are shaped mainly by shared practices (a claim that is developed in the next section). And it is these practices that will be the focus of this chapter.

Many people talk as if a safety culture can only be achieved through some awesome transformation, akin to a religious experience. This chapter takes the opposite view, arguing that a safety culture can be socially engineered by identifying and fabricating its essential components and then assembling them into a working whole. It is undoubtedly true that a bad organizational accident can achieve some dramatic conversions to the 'safety faith', but these are all too often shortlived. A safety culture is not something that springs up ready-made from the organizational equivalent of a near-death experience, rather it emerges gradually from the persistent and successful application of practical and down-to-earth measures. There is nothing mystical about it. Acquiring a safety culture is a process of collective learning, like any other. Nor is it a single entity. It is made up of a number of interacting elements, or ways of doing, thinking and managing that have enhanced safety health as their natural byproduct.

What is an Organizational Culture?

To those with a 'hard' engineering background, many attempts to describe the nature of organizational culture must seem to have the definitional precision of a cloud. There is no standard definition, but here is one that captures most of the essentials with the minimum of fuss:

> Shared values (what is important) and beliefs (how things work) that interact with an organization's structures and control systems to produce behavioural norms (the way we do things around here).[2]

Until the 1980s, 'culture' was a term applied more to nationalities than to organizations. 'Organizational culture' became an essential part of 'management speak' largely as the result of two widely read books: *Corporate Culture* (by Terrence Deal and Allan Kennedy)[3] and *In Search of Excellence* (by Thomas Peters and Robert Waterman),[4] both published in 1982.

The latter book introduced the notion of *cultural strength*, and observed 'Without exception, the dominance and coherence of culture proved to be an essential quality of the excellent companies'.[5] Fifteen years later, there are some who might doubt this assertion (particularly those laid off from the 'excellent' companies), but few would argue with the idea that a strong culture is one in which all levels of the organization share the same goals and values. To quote Peters and Waterman again, 'In these [strong culture] companies, people way down the line know what they are supposed to do in most situations because the handful of guiding values is crystal clear'.[6]

Organizational theorists have described a number of negative or dysfunctional cultures. One such 'bad' culture is characterized by what psychologists have termed *learned helplessness*, describing a condition in which people learn that attempts to change their situation are fruitless so that they simply give up trying: 'The energy and will to resolve problems and attain goals drains away.'[7] Another counterproductive organizational strategy is *anxiety-avoidance*. When such an organization discovers a technique for reducing its collective anxiety, it is likely to be repeated over and over again regardless of its actual effectiveness.

> The reason is that the learner will not willingly test the situation to determine whether the cause of the anxiety is still operating. Thus all rituals, patterns of thinking or feeling, and behaviours that may originally have been motivated by a need to avoid a painful, anxiety-provoking situation are going to be repeated, even if the causes of the original pain are no longer acting, because the avoidance of anxiety is, itself, positively reinforcing.[8]

Both learned helplessness and repetitive anxiety-avoidance are likely to assist in driving the blame cycle, described in Chapter 7. People feel helpless in the face of ever-present dangers and, while the familiar reactions to incidents and events, such as 'write another procedure' and 'blame and train,' may not actually make the system more resistant to future organizational accidents, they at least serve the anxiety-reducing function of being seen to do something—and blaming those at the sharp end deflects blame from the organization as a whole.

There is a controversy among social scientists as to whether a culture is something an organization 'has' or whether it is something an organization 'is'. The former view emphasizes management's power to change culture through the introduction of new measures and practices, while the latter sees culture as a global property that emerges out of the values, beliefs and ideologies of the organization's entire membership. The former approach is favoured by managers and management consultants, while the latter is preferred by academics and social scientists. This chapter stands with the managers and agrees with the organizational anthropologist, Geert Hofstede, when he wrote:

> On the basis of [our] research project, we propose that practices are features an organization *has*. Because of the important role of practices in organizational cultures, the ['has' approach] can be considered as somewhat manageable. Changing collective values of adult people in an intended direction is extremely difficult, if not impossible. Values do change, but not according to someone's master plan. Collective practices, however, depend on organizational characteristics like structures and systems, and can be influenced in more less predictable ways by changing these.[9]

Although the idea of a safety culture has existed since 1980, it was given an authoritative boost by the International Atomic Energy Agency when they published a report in 1988, elaborating the concept in detail. They defined safety culture as: '... that assembly of characteristics and attitudes in organizations and individuals which establishes that, as an overriding priority, nuclear plant safety issues receive the attention warranted by their significance'.[10] Unfortunately, this is something of a 'motherhood' statement specifying an ideal but not the means to achieve it. A more useful definition, which is worth quoting in full, has been given by the UK's Health and Safety Commission in 1993:

> The safety culture of an organization is the product of individual and group values, attitudes, competencies, and patterns of behaviour that determine the commitment to, and the style and proficiency of, an organization's health and safety programmes. Organizations with a positive safety culture are characterized by communications founded on mutual trust, by shared perceptions of the importance of safety, and by confidence in the efficacy of preventive measure.[11]

While remaining in sympathy with this definition, this chapter emphasizes the critical importance of an effective safety information system—the principal basis of an *informed culture.* It must be stressed again that our primary concern in this book is not with traditional health and safety measures that are directed, for the most part, at the prevention of *individual* work accidents. Our focus is upon the limitation of *organizational* accidents, and it is this that has shaped the arguments set out below.

The Components of a Safety Culture

The main elements of a safety culture and their various interactions are previewed below. Each subcomponent will be discussed more fully in succeeding sections.

- As indicated in Chapter 6, an ideal safety culture is the engine that continues to propel the system towards the goal of maximum safety health, regardless of the leadership's personality or current commercial concerns. Such an ideal is hard to achieve in the real world, but it is nonetheless a goal worth striving for.
- The power of this engine relies heavily upon a continuing respect for the many entities that can penetrate and breach the defences. In short, its power is derived from not forgetting to be afraid.
- In the absence of bad outcomes, the best way—perhaps the only way—to sustain a state of intelligent and respectful wariness is to gather the right kinds of data. This means creating a safety information system that collects, analyses and disseminates information from incidents and near-misses as well as from regular proactive checks on the system's vital signs (see Chapter 7). All of these activities can be said to make up an *informed culture*—one in which those who manage and operate the system have current knowledge about the human, technical, organizational and environmental factors that determine the safety of the system as a whole. In most important respects, an informed culture *is* a safety culture.
- Any safety information system depends crucially on the willing participation of the workforce, the people in direct contact with the hazards. To achieve this, it is necessary to engineer a *reporting culture*—an organizational climate in which people are prepared to report their errors and near-misses.
- An effective reporting culture depends, in turn, on how the organization handles blame and punishment. A 'no-blame' culture is neither feasible nor desirable. A small proportion of human unsafe acts are egregious (for example, substance abuse, reckless non-compliance, sabotage and so on) and warrant sanctions, severe ones in some cases. A blanket amnesty on all unsafe acts would lack credibility in the eyes of the workforce. More importantly, it would be seen to oppose natural justice. What is needed is a *just culture*, an atmosphere of trust in which people are encouraged, even rewarded, for providing essential safety-related information—but in which they are also clear about where the line must be drawn between acceptable and unacceptable behaviour.

- The evidence shows that high-reliability organizations—domain leaders in health, safety and environmental issues—possess the ability to reconfigure themselves in the face of high-tempo operations or certain kinds of danger. A *flexible culture* takes a number of forms, but in many cases it involves shifting from the conventional hierarchical mode to a flatter professional structure, where control passes to task experts on the spot, and then reverts back to the traditional bureaucratic mode once the emergency has passed. Such adaptability is an essential feature of the crisis-prepared organization and, as before, depends crucially on respect—in this case, respect for the skills, experience and abilities of the workforce and, most particularly, the first-line supervisors. But respect must be earned, and this requires a major training investment on the part of the organization.
- Finally, an organization must possess a *learning culture*—the willingness and the competence to draw the right conclusions from its safety information system, and the will to implement major reforms when their need is indicated.

The preceding bullet points have identified four critical subcomponents of a safety culture: a *reporting culture*, a *just culture*, a *flexible culture* and a *learning culture*. Together they interact to create an *informed culture* which, for our purposes, equates with the term 'safety culture' as it applies to the limitation of organizational accidents.

Engineering a Reporting Culture

On the face of it, persuading people to file critical incident and near-miss reports is not an easy task, particularly when it may entail divulging their own errors. Human reactions to making mistakes take various forms, but frank confession does not usually come high on the list. Even when such personal issues do not arise, potential informants cannot always see the value in making reports, especially if they are sceptical about the likelihood of management acting upon the information. Is it worth the extra work when no good is likely to come of it? Moreover, even when people are persuaded that writing a sufficiently detailed account is justified and that some action will be taken, there remains the overriding problem of trust. Will I get my colleagues into trouble? Will I get into trouble?

There are some powerful disincentives to participating in a reporting scheme: extra work, scepticism, perhaps a natural desire to forget that the incident ever happened, and—above all—lack of trust and, with it, the fear of reprisals. Nonetheless, many highly effective reporting programmes do exist. What can we learn from them? How have they engineered their success?

In what follows, we will briefly look at the 'social engineering' details of two successful aviation reporting programmes, one operating at a national level and the other within a single airline. These are NASA's Aviation Safety Reporting System (ASRS) and the British Airways Safety Information System (BASIS). In this, we will rely heavily on the work of two people—Dr Sheryl Chappell of NASA, and Captain Mike O'Leary of British Airways[12]—each of whom has been closely involved in the design and management of these programmes. Our purpose here is not to consider the programmes themselves in any detail, but to abstract from them the 'best practices' for achieving a reporting culture. Throughout, we will concentrate on the issue of how valid reporting may be promoted. Although we are dealing exclusively with aviation reporting schemes, the basic 'engineering' principles can be applied in any domain. Indeed, reporting programmes in other domains—particularly in medicine—are partly modelled on pioneering aviation schemes such as critical incident reporting.

Examination of these successful programmes indicates that five factors are important in determining both the quantity and the quality of incident reports. Some are essential in creating a climate of trust, others are needed to motivate people to file reports. The factors are:

- Indemnity against disciplinary proceedings—as far as it is practicable.
- Confidentiality or de-identification.
- The separation of the agency or department collecting and analysing the reports from those bodies with the authority to institute disciplinary proceedings and impose sanctions.
- Rapid, useful, accessible and intelligible feedback to the reporting community.
- Ease of making the report.

The first three items are designed to foster a feeling of trust. O'Leary and Chappell explain the need:

> For any incident reporting programme to be effective in uncovering the failures which contribute to an incident, it is paramount to earn the trust of the reporters. This is even more important when there is a candid disclosure of the reporter's own errors. Without such trust, the report will be selective and will probably gloss over pivotal human factors information. In the worst case—that in which potential reporters have no trust in the safety organization—there may be no report at

> all. Trust may not come quickly. Individuals may be hesitant to report until the reporting system has proved that it is sensitive to reporters' concerns. Trust is the most important foundation of a successful reporting programme, and it must be actively protected, even after many years of successful operation. A single case of a reporter being disciplined as the result of a report could undermine trust and stop the flow of useful reports.[13]

The rationale for any reporting system—and a recurrent theme throughout this book—is that valid feedback on the local and organizational factors promoting errors and incidents is far more important than assigning blame to individuals. To this end, it is essential to protect informants and their colleagues as far as possible from disciplinary actions taken on the basis of their reports. But there will be limits upon this indemnity. These limits are defined most clearly by the Waiver of Disciplinary Action issued in relation to NASA's Aviation Safety Reporting System. Below is an excerpt from the FAA Advisory Circular (AC No. 00-46C) describing how the immunity concept applies to pilots making incident reports.

> The filing of a report with NASA concerning an incident or occurrence involving a violation of the Act of the Federal Aviation Regulations is considered by the FAA to be indicative of a constructive attitude. Such an attitude will tend to prevent future violations. Accordingly, although a finding of a violation may be made, neither a civil penalty nor a certificate suspension will be imposed if:
>
> - The violation was inadvertent and not deliberate;
> - The violation did not involve a criminal offence, or accident or ... a lack of qualification or competency;
> - The person has not been found in any prior FAA enforcement action to have committed a violation of the Federal Aviation Act, or of any regulation promulgated under the Act for a period of 5 years prior to the date of the occurrence; and
> - The person proves that, within 10 days after the violation, he or she completed and delivered or mailed a written report of the incident or occurrence to NASA under ASRS.[14]

This formula appears to work. The ASRS reporting rate was high, even at the outset. In the beginning, it averaged approximately 400 reports per month. It now runs at around 650 reports per week and more than 2000 reports per month. In 1995 ASRS received over 30 000 reports.

BASIS has been extended over the years to cover a wide variety of reporting schemes. All flight crew are required to report safety-related events using Air Safety Reports (ASRs). ASRs are not anonymous. To encourage the filing of ASRs, British Airways Flight Crew Order, No. 608 states:

> It is not normally the policy of British Airways to institute disciplinary proceedings in response to the reporting of any incident affecting air safety. Only in rare circumstances where an employee has taken action or risks which, in the Company's opinion, no reasonably prudent employee with his/her training and experience would have taken, will British Airways consider initiating such disciplinary action.[15]

Again, the formula seems to work. Its success is suggested by two statistics. First, the ASR filing rate more than trebled between its inception in 1990 and 1995. Second, the combined number of reports assigned to the severe and high risk categories has decreased by two-thirds between the first six months of 1993 and the first half year of 1995.[16]

Another important component of BASIS is the British Airways Confidential Human Factors Reporting Programme, instituted in 1992. While the ASRs provided good technical and procedural information, a need was felt for an information channel that was more sensitive to human factors issues. Each pilot filing an ASR is now invited to complete a confidential human factors questionnaire relating to the incident. The return of the questionnaire is voluntary. The following assurance was given by the senior manager in charge of BA's Safety Services on the front page of the initial version:

> I give an absolute assurance that any information you provide will be treated confidentially by Safety Services and that this questionnaire will be destroyed immediately after the data is processed. This programme is accessible only to my Unit.

In its first year of operation, the human factors reporting programme received 550 usable responses.[17] The issues raised in the reports are communicated to management on a regular basis, but great care is taken to separate the important safety issues from the incidents in order to preserve the anonymity of the reports.

Another important input to BASIS comes from the Special Event Search and Master Analysis (SESMA). This by-passes the need for human reporting by monitoring directly the flight data recorders (FDRs) of BA's various aircraft fleets, while at the same time guaranteeing the flight crews complete anonymity. The FDR for each flight is scanned for events that are considered to lie outside safe norms. All events are stored in a BASIS database and the more serious are discussed at a monthly meeting of technical managers and the pilots' union representatives. If the incident is considered to be sufficiently serious, the union representative is required to discuss the matter with the flight crew involved—while still withholding their identities from the management.

When a report is received by NASA's ASRS staff it is processed in the following manner, with great care being taken to preserve the anonymity of the reporter.[18]

- An initial analysis screens out reports involving accidents, criminal behaviour, or those classified as 'no safety content'.
- The report is coded and the reporter de-identified. At this stage, the reporter is also contacted by telephone to confirm receipt and de-identification.
- After a quality check, the information is entered into the ASRS database and the original report destroyed.

The most obvious way of ensuring confidentiality is to have the reports filed anonymously. But, as O'Leary and Chappell point out, this is not always possible or even desirable.[19] The main problems with total anonymity are as follows:

- Analysts cannot contact the informant to resolve questions.
- It is more likely that some managers will dismiss anonymous reports as the work of disaffected troublemakers.
- In small companies, it is almost impossible to guarantee anonymity.

O'Leary and Chappell conclude that removing identities from reports at a later stage—as described above for ASRS—is probably the most workable means of maintaining confidentiality. At a national level, complete de-identification means removing not only the people's names, but also the date, the time, the flight number and the airline name. The criteria for de-identification must be known and understood by all potential reporters.

Another important measure for engendering trust is to separate the organization receiving the reports from both the regulatory body and from the employing company. As in the case of ASRS, the system analysts should ideally have no legal or operational authority over the potential reporters. Reporting systems run by disinterested third parties—such as universities—can also help to earn the trust of reporters. If, like BASIS, the reporting system is internal to a company, the receiving department should be perceived as being completely independent of operational management, thus giving the necessary assurance of confidentiality.

Apart from a lack (or loss) of trust, few things will stifle incident reporting more than the perceived absence of any useful outcome. Both the ASRS and BASIS place great emphasis on the rapid feedback of meaningful information to their respective communities. If an ASRS report describes a continuing hazardous situation—for example, a defective navigation aid, a confusing procedure, or an incorrect chart—an alerting message is sent out immediately to the appropriate authorities so that they can investigate the problem and take the necessary remedial action. (As mentioned earlier, ASRS has no legal or operational authority of its own.) Some 1700 alert bulletins and notices have been issued by the ASRS team since the programme began in 1976. In 1994 there was a 65 per cent response rate to alert bulletins and FYI notices.

The information assembled in the ASRS database is made available to the aviation and research communities in a variety of ways. First, targeted searches can be carried out at the request of companies, agencies or researchers. The information is also disseminated via a newsletter, *Callback*, whose extended readership is estimated at over 100 000, and other ASRS reports. Such newsletters describe safety issues and highlight improvements that have been made as the result of incident reporting. This serves the double function of both informing the reporters and congratulating them on their collective contribution to aviation safety.

British Airways Safety Services also disseminate their BASIS information in a variety of ways. In addition to unit reports and journal articles, they issue *Flywise*, an 18–20 page monthly bulletin that includes trend analyses relating to selected events and brief accounts of incidents broken down by fleets. Each incident is assigned both a

SEVERITY **RISK MATRIX**

HIGH	C	B	A (severe)
MEDIUM	D	C	B (high)
LOW	E (minimal)	D (low)	C (medium)
	LOW	MEDIUM	HIGH

PROBABILITY OF RECURRENCE

Figure 9.1 The British Airways risk management matrix used to evaluate the future risk to the company of the recurrence of an event
The matrix generates risk categories on a scale from A (severe risk) to E (minimal risk).

risk category (based on a risk matrix—see Figure 9.1) and one of the following action categories:

- *Active investigation*—actions to prevent recurrence not fully understood.
- *Action required*—preventive measures have been identified but not yet implemented.
- *Action monitored*—preventive measures have been implemented and their effects are being monitored.
- *Report monitored*—action taken without need for further investigation by Safety Services. Rates of occurrence are being monitored.

The last factor to be considered here is ease of reporting. The format, length and content of the reporting form or questionnaire are extremely important, as is the context in which respondents are expected to make their report. Privacy and a labour-free returning mechanism are all important incentives—or, to put it the other way round, their absence could be a deterrent. O'Leary and Chappell make the following observations regarding the design of the reporting form:

> If a form is long and requires a great deal of time to complete, reporters are less likely to make the effort. If the form is too short, it is difficult to obtain all the necessary information about the incident. In general, the more specific the questions, the easier it is to complete the questionnaire; however, the information provided will be limited by the choice of the questions. More open questions about the reporter's perceptions, judgements, decisions and actions are not subject to this limitation and give the reporter a greater chance to tell the full story. This method is more effective in gathering all the information about an incident, but takes longer and usually requires more analytic resources within the reporting system.[20]

A certain amount of trial-and-error learning may be necessary before an organization hits upon a format that is best suited both to its purpose and to its potential respondents. In this regard, we can learn from the experience of British Airways Safety Services with their confidential human factors questionnaire. In its initial form, it asked a limited number of very specific questions. Some of the questions sought to establish whether either a slip or a mistake had occurred (see Chapter 4 for the technical descriptions of these unsafe acts) and, if so, what were the contributing factors. The latter were listed below each item and required 'yes/no' responses: in the case of action slips and lapses, they included tiredness, time pressure, lack of stimulation and flight deck ergonomics; in the case of mistakes, they included misleading manuals, misleading displays, insufficient training and crew cooperation.

Even though one of the questions asked what went right, several respondents complained about the negative flavour of the questions overall, and it was realized that this could well deter some potential respondents from completing the questionnaire at all. In addition, the BA analysts were unhappy with the validity of some of the data, since the technical distinctions between slips, lapses and mistakes were not well understood by the flight crew respondents. As a result, BA Safety Services launched a new jargon-free questionnaire in 1995. This asked open-ended questions covering a range of factors from local flight deck influences to the effectiveness of training. O'Leary gives two questions as examples of this new approach:

- How did you and the crew initially respond to the event, and how did you establish what technical and personal issues were involved?
- Was all the relevant flight, FMS and system information clearly available and were all the controls and selectors useful and unambiguous? If not, how could these be improved?[21]

As O'Leary observes, this style of questioning moves the analytic workload from the reporter to the human factors analysts. Although the change has increased demand on the limited resources available to process the data, it has made the questionnaire sensitive to a variety of issues not previously covered. In addition, the reliability of the analysis has improved dramatically: 'Previously, some 3500 pilots and engineers may report events idiosyncratically. Now we use a team of only a dozen volunteer flight crew analysts'.[22]

With the multiple choice format of the previous form of the human factors questionnaire, it was relatively simple to convert 'yes/no' responses directly into bar charts. With the second, more open-ended version, the BA analysts had to develop an agreed classification structure in order to identify the issues with human factors significance. They were interested in two major categories: crew performance and the influences upon that behaviour. Crew behaviour was subdivided into error descriptors (action slip or mistake), crew team skills and task specifics, such as automatic or manual handling. Influences too were divided into three groups: organizational factors (training, procedures, commercial pressure and the like), environmental factors (airport facilities, weather conditions and the like) and personal factors (automation, complacency, morale and the like). Future developments are directed at establishing causal links between the various factors—a shift from addressing the 'what?' question to tackling the 'why?' question. The aim here is to identify 'resident pathogens' that may contribute to a variety of different problems on the flight deck.

Figure 9.2 gives an idea of what this development might yield in the way of causal analysis. The figure summarizes a fictional incident involving a rejected takeoff. Here, operational stress was created by a busy airfield and communicating with the company to establish load-sheet figures while taxiing out. The climate on the flight deck was poor. The co-pilot felt overloaded but was not able to communicate this to the captain. He focused on one task at a time and did not cross-check what the captain was doing. As a result, he omitted the 'before takeoff' checks. Takeoff clearance was given as the aircraft approached the runway and, as takeoff power was set, the configuration warning horn indicated that no flap had been selected and the takeoff was aborted.

Finally in this context of engineering a reporting culture, it is worth asking whether there is any scientific evidence to support the efficacy of near-miss accident reporting. In one Swedish study,[23] relating

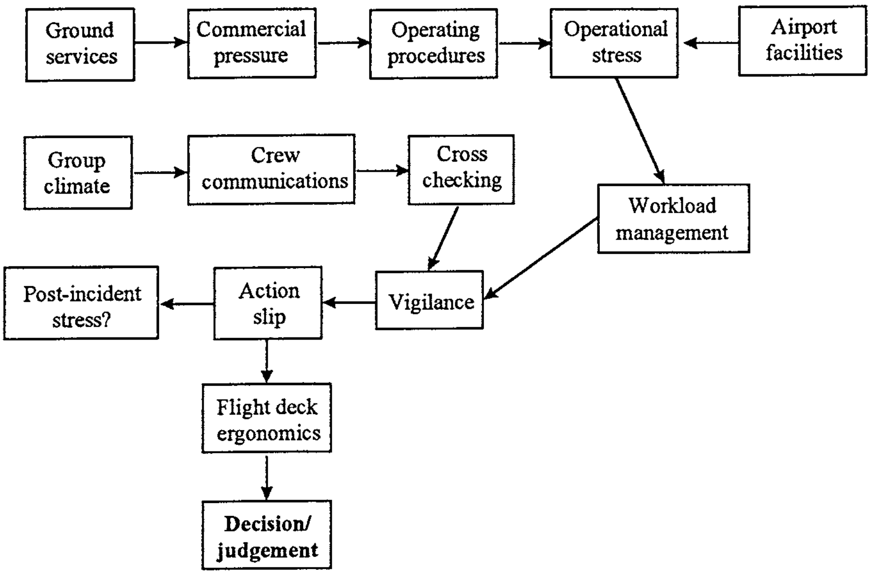

Figure 9.2 Flow diagram of the rejected takeoff incident showing judged causal linkages (after O'Leary)[24]

primarily to individual accidents, the implementation of an incident reporting scheme resulted in an increase in the number of remedial suggestions from the workforce but no significant reduction in the accident rate. In a follow-up study, participants received training in how to recognize and interpret critical incidents. This resulted in a 56 per cent reduction in the severity of injuries but no drop in the accident frequency rate. The main message of these findings is that potential respondents need to be very clear about what constitutes an incident. In some situations, this is not always intuitively obvious.

Engineering a Just Culture

A wholly just culture is almost certainly an unattainable ideal. However, an organization in which the majority of its members share the belief that justice will usually be dispensed is within the bounds of possibility. Two things are clear at the outset. First, it would be quite unacceptable to punish all errors and unsafe acts regardless of their origins and circumstances. Second, it would be equally unacceptable to give a blanket immunity from sanctions to all actions that could, or did, contribute to organizational accidents. While this book has strongly emphasized the situational and systemic factors leading to the catastrophic breakdown of hazardous technologies, it would be naïve not to recognize that, on some relatively rare occasions, accidents can happen as the result of the unreasonably reckless, negligent or even malevolent behaviour of particular individuals. The difficulty lies in discriminating between these few truly 'bad behaviours' and the vast majority of unsafe acts to which the attribution of blame is neither appropriate nor useful.

A prerequisite for engineering a just culture is an agreed set of principles for drawing the line between acceptable and unacceptable actions. To this end, we will start by outlining some of the psychological and legal issues that must be taken into account when making this judgement. Figure 9.3 sets the scene.

All human actions involve three core elements:

- An *intention* that specifies an immediate goal and—where these goal-related actions are not wholly automatic or habitual—the behaviour necessary to achieve it.
- The *actions* triggered by this intention—which may or may not conform to the action plan.
- The *consequences* of these actions—which may or may not achieve the desired objective. The actions can be either successful or unsuccessful in this respect.

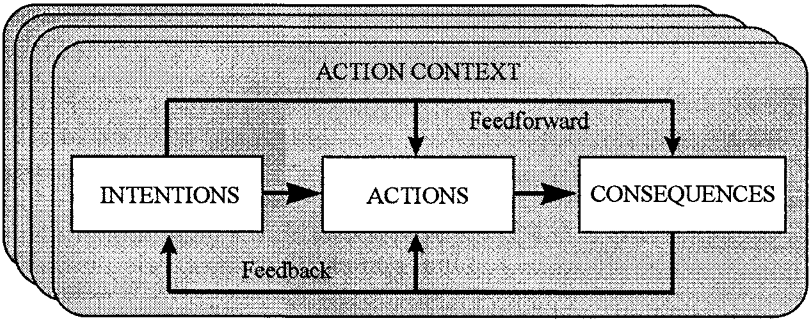

Figure 9.3 The basic elements of human action

The feedforward arrows shown in Figure 9.3 indicate that, in formulating an intention, actions are selected in the belief that they will achieve the goal (or at least provide useful feedback to ensure the success of future actions), but this belief is not always justified. The feedback arrows complete the loop by providing information about the success or otherwise of the preceding actions and their outcomes. Human actions are embedded in a context that contains both the immediate physical environment and the purpose of the behavioural sequence of which a particular action forms a part. The historical context is shown symbolically in Figure 9.3 by the three preceding action frames in the background.

Both of these issues have a close bearing on individual responsibility. In the case of hazardous technologies it is inevitable that all physical situations will contain an element of risk. But is also likely that individual actors will have been—or should have been—trained to foresee and to minimize these risks. This brings us back to the distinction made in Chapter 4 between successful and unsuccessful behaviour on the one hand, and correct and incorrect behaviour on the other. Although success is determined solely by whether the planned actions achieve their immediate objectives, success does not necessarily mean correctness. Successful actions may be incorrect. That is, they could achieve their local purpose and yet be either reckless or negligent.

In the law, a person who acts recklessly is one who takes a deliberate and unjustifiable risk (that is, one that is foreseeable, and where a bad outcome is likely though not certain). However, as Smith and Hogan point out:

> The operator of aircraft, the surgeon performing an operation and the promoter of a tightrope act in the circus must all foresee that their acts might cause death; but we should not describe them as reckless, un-

> less the risk taken was unjustifiable. Whether the risk is justifiable depends on the social value of the activity involved, as well as on the probability of the occurrence of the foreseen evil.[25]

Negligence, on the other hand, involves bringing about a consequence that a 'reasonable and prudent' person would have foreseen and avoided. One can also be negligent with regard to a circumstance: 'A person acts negligently with respect to a circumstance when a reasonable man would have been aware of the existence of the circumstance and, because of its existence would have avoided acting in that manner.'[26] In the latter case whether the person failed to foresee the bad outcome and was unaware of the circumstance is irrelevant. For example, X picks up a gun, believing it to be unloaded, points it at Y and pulls the trigger. If any reasonable person would have realized that the gun might possibly be loaded, and thus avoided acting in this way, then X was negligent with regard to circumstance. If the gun was loaded and kills Y, then X was negligent with regard to consequence. In a court of law, it is not necessary for the prosecution to prove anything at all about the person's state of mind at the time of the act. It is enough to establish that particular actions were carried out in certain circumstances. Negligence is historically a civil rather than a criminal law concept, and has a much lower level of culpability than recklessness.[27]

Those involved in the operation of hazardous technologies are often perceived as carrying an additional burden of responsibility by virtue of their training and of the great risks associated with human failure. For example, in the case of *Alidair* v. *Taylor* in 1978, Lord Denning ruled that:

> There are activities in which the degree of professional skill which must be required is so high, and the potential consequences of the smallest departure of that high standard are so serious, that one failure to perform in accordance with those standards is enough to justify dismissal.[28]

This 'hang them all' judgement is unsatisfactory in many respects. It ignores the ubiquity of error as well as the situational factors that promote it. Nor is it sensitive to the varieties of human failure and their differing psychological origins. Pushing this judgement to an absurd conclusion, it could be claimed that, since all pilots, control room operators and others with safety-critical jobs in hazardous technologies are fallible, they will all, at some time or another, inevitably fall short of Lord Denning's 'high standards' and so should all be sacked. Even wise and distinguished judges do not get it right all of the time.

A much sounder guideline is Neil Johnston's *substitution test.*[29] This is in keeping with the principle that the best people can make the worst errors. When faced with an accident or serious incident in which the unsafe acts of a particular person were implicated, we should perform the following mental test. Substitute the individual concerned for someone else coming from the same domain of activity and possessing comparable qualifications and experience. Then ask the following question: 'In the light of how events unfolded and were perceived by those involved in real time, is it likely that this new individual would have behaved any differently?' If the answer is 'probably not' then, as Johnston put it, '... apportioning blame has no material role to play, other than to obscure systemic deficiencies and to blame one of the victims'. A useful addition to the substitution test is to ask of the individual's peers: 'Given the circumstances that prevailed at that time, could you be sure that you would not have committed the same or similar type of unsafe act?' If the answer again is 'probably not', then blame is inappropriate.

So much for the background. We will now turn to the task of grading unsafe acts according to their blameworthiness. In this, as in jurisprudence, a crucial discriminator is the nature of the intention. A crime has two key elements: the *mens rea,* or 'guilty mind' and the *actus reus,* or 'guilty act'. Both are necessary for its commission. Except in very specific instances (as, for example, in the case of negligence), the act without the mental element is not a crime. While, for the most part, we are not concerned here with criminal behaviour, we will adopt the 'mental element' principle as a basic guideline. But hereafter we will approach the issue more from a psychological perspective than from a legal one.

Figure 9.4 sketches out the bare essentials of a decision tree for discriminating the culpability of an unsafe act. It is assumed that the actions under scrutiny have contributed either to an accident or to a serious incident in which a bad outcome was only just averted. In an organizational accident, there are likely to be a number of different unsafe acts, and the decision tree is intended to be applied separately to each of them. Our concern here is with individual unsafe acts committed by either a single person or by different people at various points in the accident sequence.

The key questions relate to intention. If both the actions and the consequences were intended, then we are likely to be in the realm of criminal behaviour and that is probably beyond the scope of the organization to deal with internally. Unintended actions define slips and lapses—in general, the least blameworthy of errors—while unintended consequences cover mistakes and violations. The decision tree usually treats the various error types in the same way, except with regard to the violations question. For mistakes, the question

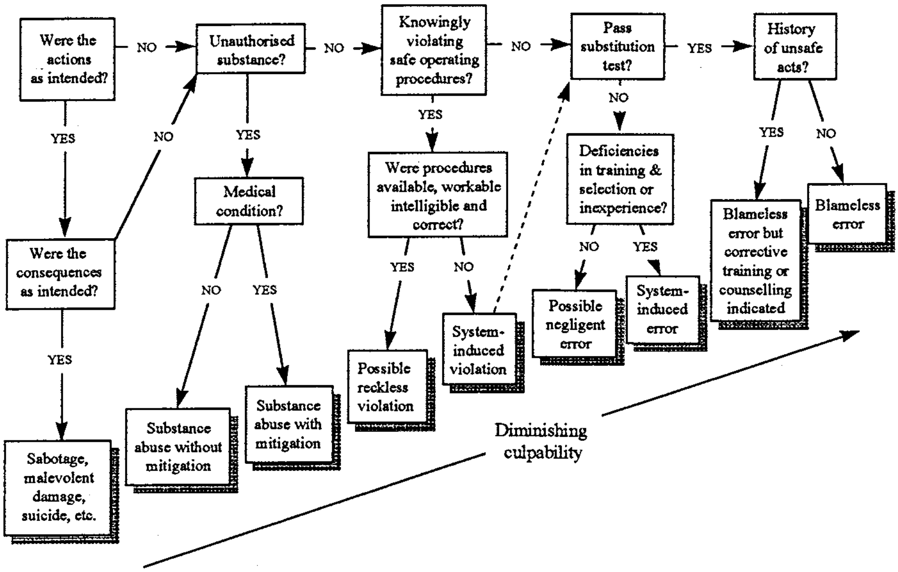

Figure 9.4 A decision tree for determining the culpability of unsafe acts

reads as shown in Figure 9.4, but for slips and lapses, the question relates to what the person was doing when the slip or lapse occurred. If the individual was knowingly engaged in violating safe operating procedures at that time, then the resulting error is more culpable since it should have been realized that violating increases both the likelihood of making an error and the chances of bad consequences resulting (see Chapter 4).

The 'unauthorized substance' question seeks to establish whether or not the individual was under the influence of alcohol or drugs known to impair performance at the time the unsafe act was committed. Since the ingestion of unauthorized substances is usually a voluntary act, their involvement would indicate a high level of culpability. But the matter is not entirely straightforward. In 1975, during a descent towards Nairobi, the co-pilot of a Boeing 747 misheard an air traffic control instruction. Instead of 'seven five zero zero', he heard 'five zero zero zero' and set the autopilot to level out at 5000 feet.[30] Unfortunately, that would have placed the aircraft in a tunnelling mode since it was around 300 feet below the unusually high airfield. When the aircraft broke cloud, the flight crew saw the ground a little more than 200 feet below them. Prompt action by the captain prevented this from being the first major disaster involving a Boeing 'jumbo' jet. It later transpired that the co-pilot had picked up a large tapeworm on a holiday in India and was dosing himself with unauthorized drugs that had, among their side-effects, drowsiness and nausea. Taking unauthorized medication as the result of a medical condition, while clearly reprehensible, is less blameworthy than taking drugs or alcohol for 'recreational purposes' and, as such, in Figure 9.4 it has been assigned to the category of 'substance abuse with mitigation'. The degree of mitigation will, of course, depend upon the local circumstances.

Except when non-compliance has become a largely automatic way of working (as sometimes happens in the case of routine short-cuts), violations involve a conscious decision on the part of the perpetrator to break or bend the rules. However, while the actions may be deliberate, the possible bad consequences are not—in contrast to sabotage in which both the act and the consequences are intended. Most violations will be non-malevolent in terms of intent, so the degree to which they are blameworthy will depend largely on the quality and availability of the relevant procedures. These, as discussed in Chapter 4, are not always appropriate for the particular situation. Where this is judged to be the case—perhaps by a 'jury' of the perpetrator's peers—the problem lies more with the system than with the individual. However, when good procedures were readily accessible but deliberately violated, the question must arise as to whether the behaviour was reckless in the legal sense of the term. Such actions are clearly more culpable than 'necessary' violations—that is, non-compliant actions necessary to get the job done when the relevant procedures are either wrong, inappropriate or unworkable.

It seems appropriate to apply Johnston's substitution test once the issues of possible substance abuse and deliberate non-compliance have been settled, although something like it clearly has a part to play in judging the culpability of system-induced violations (as indicated by the dotted arrow in Figure 9.4). The issue is reasonably straightforward. Could (or has) some well motivated, equally competent and comparably qualified individual make (or made) the same kind of error under those or very similar circumstances? If the answer given by a 'jury' of peers is 'yes', then the error is probably blameless. If the answer is 'no', then we have to consider whether there were any system-induced deficiencies in the person's training, selection or experience. If such latent conditions are not identified, then the possibility of a negligent error must be considered. If they are found, it is likely that the unsafe act was a largely blameless system-induced error.

Such a category would apply to the technician whose miswiring of a signal box significantly contributed to the Clapham Junction rail disaster (see Chapter 5). His actions would not pass the substitution test, since he was largely self-taught and had acquired his bad work practices in the absence of adequate training and supervision. But, as the Inquiry established, the underlying problems were those of the system rather than the individual, who was hardworking and highly motivated to do a good job.

In legal jargon, the last major question at the top right-hand corner of Figure 9.4 could be rephrased as 'Any previous?'. People vary widely and consistently in their liability to everyday slips and lapses. For example, some individuals are considerably more absentminded than others. If the person in question has a previous history of unsafe acts, it does not necessarily bear upon the culpability of the error committed on this particular occasion, but it does indicate the necessity for corrective training or even career counselling along the lines of 'Don't you think you would be doing everyone a favour if you considered taking on some other job within the company?'. This is the way that management acquires some of its most distinguished members. Absentmindedness has nothing whatsoever to do with ability or intelligence, but it is not a particularly helpful trait in a pilot or control room operator.

So where should the line be drawn on Figure 9.4 between acceptable and unacceptable behaviour? The most obvious point would be between the two substance abuse categories. Both malevolent damage and the dangerous use of alcohol or drugs are wholly unacceptable and should receive very severe sanctions, possibly administered by the courts rather than the organization. Between 'substance abuse with mitigation' and 'possible negligent error' lies a grey area in which careful judgement must be exercised. The remaining categories should be thought of as blameless—unless they involve aggravating factors not considered here. Experience suggests that the majority of unsafe acts—perhaps 90 per cent or more—fall into this blameless category.

What should happen to the small proportion of individuals whose unsafe acts are justly considered culpable? It is not within the competence of this chapter to advise on the nature of the sanctions. Although this is a matter for the organizations concerned, we can say something about the value—or otherwise—of punishments.

Unfortunately, a large amount of psychological research concerned with the issues of reward and punishment has involved the white rat, and is not especially relevant. Figure 9.5 summarizes in a very simplified way what psychologists know about the effects of reward and punishment in the workplace.[31] The principal issue here is the effectiveness of 'sticks and carrots' in enhancing the likelihood of

	Immediate	***Delayed***
Reward	Positive effects	Doubtful effects
Punishment	Doubtful effects	Negative effects

Figure 9.5 Summary of the effects of reward and punishment on behavioural change in the workplace

desired behaviour and reducing the chances of unwanted behaviour. Rewards are the most powerful means of changing behaviour, but they are only effective if delivered close in time and place to the behaviour that is desired. Delayed punishments have negative effects: they generally do not lead to improved behaviour and can induce resentment in both the punished and the could-be-punished. The cells labelled 'doubtful effects' mean that, in each case, there are opposing forces at work. Hence, the results are uncertain.

But there are other factors that argue strongly in favour of punishing the few who commit egregious unsafe acts. In most organizations the people in the front line know very well who the 'cowboys' and the habitual rule-benders are. Seeing them get away with it on a daily basis does little for morale or for the credibility of the disciplinary system. Watching them getting their 'come-uppance' is not only satisfying, it also serves to reinforce where the boundaries of acceptable behaviour lie. Moreover, outsiders are not the only potential victims. Justified dismissal protects the offender's colleagues. Perhaps more than other possible victims, they are likely to be endangered by the person's repeated recklessness or negligence. Their departure makes the work environment a safer place and also encourages the workforce to perceive the organizational culture as just. Justice works two ways. Severe sanctions for the few can protect the innocence of the many.

David Marx, an aircraft engineer who was one of the principal architects of the Boeing's Maintenance Error Decision Aid (see Chapter 7), made the following comments on the relationship between reporting and disciplinary systems. Though he is writing about the aviation industry, the points are widely applicable:

> Many of us have found today's disciplinary systems to be a significant obstacle to asking an employee to come forward and talk about his or

> her mistake. Consequently, as an industry, we have begun to re-evaluate the inter-relationship of employee discipline and event investigation. Many programs have been developed, both internal to an airline and in association with the FAA ... Whether it is called immunity, amnesty or 'performance-related incentive'—each program attempts to encourage the erring employee to come forward. Yet, as more incentive programs enter the marketplace of ideas, the disciplinary landscape becomes increasingly complex and confusing. With all the programs today, the individual employee needs to be a lawyer to assess whether it is safe to come forward.[32]

David Marx has recently taken a law degree and one of the most interesting products of this marriage between engineering and the law has been the computerized incident investigator, the Aurora Mishap Management System (AMMS). AMMS has a number of elements. For our present purposes, its most important aspect is a structured methodology for establishing the applicability of disciplinary action. This investigative tool is used by an organization's disciplinary review board to aid their decision-making. It applies a common and consistent approach to the issue of determining whether or not disciplinary action is warranted. To date, it has been used in the field of aircraft maintenance by a number of US airlines and has the backing of the Machinists Union.

Engineering a Flexible Culture

Organizational flexibility means possessing a culture capable of adapting effectively to changing demands. Flexibility is one of the defining properties of what a highly influential Berkeley research group—led by Todd La Porte, Karlene Roberts and Gene Rochlin—have termed *high-reliability organizations* (HROs). The group has conducted field research in a number of highly complex, technology-intensive organizations that must operate, as far as humanly possible, to a failure-free standard. The systems of interest here are air traffic control and naval air operations at sea.

The operational challenges facing these (and comparable) organizations are twofold:

- to manage complex, demanding technologies, making sure to avoid major failures that could cripple, perhaps destroy, the organization;
- at the same time, to maintain the capacity for meeting periods of very high, peak demand and production whenever these occur.[33]

The organizations studied by the Berkeley group had the following characteristics:

- They were large, internally dynamic and intermittently intensely interactive.
- Each performed complex and exacting tasks under considerable time pressure.
- They have carried out these demanding activities with a very low error rate and an almost complete absence of catastrophic failure over a number of years.

On the face of it, both of the organizations to be considered here—the US Navy nuclear aircraft carrier and the air traffic control centre—had highly bureaucratic and hierarchical organizational structures, each with a clear line of authority and command. Both organizations relied heavily on tested standard operating procedures (SOPs). Both organizations invested a great deal of effort in training people in the use of these procedures. It was almost the case that, under routine operating conditions, the only decision necessary was which SOP to apply.

Actions in these HROs were closely monitored so that immediate investigations—termed 'hot washups' in the US Navy—were conducted whenever errors occurred. Over the years, these organizations have learned that there are particular kinds of error, often quite minor, that can escalate rapidly into major, system-threatening failures. Trial-and-error learning in these critical areas was not encouraged, as it was elsewhere, in case it should become 'habit-forming.' Also, as La Porte and Consolini describe it: 'there is a palpable sense that there are likely to be similar events that cannot be foreseen clearly, and that may be beyond imagining. This is an ever-present cloud over operations, a constant concern'.[34] In short, these organizations suffer chronic unease. The following quotation from the same source captures this intelligent wariness and its cultural consequences very eloquently:

> The people in these organizations know almost everything technical about what they are doing—and fear being lulled into supposing that they have prepared for any contingency. Yet even a minute failure of intelligence, a bit of uncertainty, can trigger disaster. They are driven to use a proactive, preventative decision making strategy. Analysis and search come before as well as after errors. They try to be synoptic while knowing that they can never fully achieve it. In the attempt to avoid the pitfalls in this struggle, decision making patterns appear to support apparently contradictory production-enhancing and error-reduction strategies. The patterns encourage

- reporting errors without encouraging a lax attitude toward the commission of errors;
- initiative to identify flaws in SOPs and nominate and validate changes in those that prove to be inadequate;
- error avoidance without stifling initiative or (creating) operator rigidity; and
- mutual monitoring without counter-productive loss of operator confidence, autonomy and trust.[35]

So how do HROs respond to bursts of high-tempo operations? Lying in wait beneath the surface of the routine, bureaucratic, SOP-driven mode is quite another pattern of organizational behaviour. Here is what happened aboard the aircraft carrier when some 70 of its 90 aircraft were flying off on missions:

> Authority patterns shift to a basis of functional skill. Collegial authority (and decision patterns) overlay bureaucratic ones as the tempo of operations increases. Formal rank and status decline as a reason for obedience. Hierarchical rank defers to technical expertise often held by those of lower formal rank. Chiefs (senior non-commissioned officers) advise commanders, gently direct lieutenants and cow ensigns. Criticality, hazards, and sophistication of operations prompt a kind of functional discipline, a professionalization of the work teams. Feedback and (sometimes conflictual) negotiations increase in importance; feedback about 'how goes it' is sought and valued.[36]

A similar kind of flexibility was evident in the air traffic control centre. Sudden wind shifts can impose a high additional burden on already busy controllers. Reorienting the flight paths of a large number of aircraft in relation to what, in this instance, were three major airports, two large military airbases and five smaller general aviation airfields becomes a major programme for the controllers on duty. La Porte and Consolini described what happened:

> The tempo at the approach-control facility and the enroute center increases, and controllers gather in small groups around relevant radar screens, plotting the optimal ways to manage the traffic as the shift in [wind] direction becomes imminent. Advice is traded, suggestions put forward, and the actual traffic is compared with the simulations used in the long hours of training the controllers undergo... . While there are general rules and controllers and supervisors have formal authority, it is the team that rallies round the controllers in 'the hot seats'. It will be the experienced controller virtuosos [rather than the supervisors] who dominate the decision train. 'Losing separation'—the key indicator of controller failure—is too awful to trust to rules alone.[37]

When the high-tempo period slackens off, authority reverts seamlessly to its previous bureaucratic, rank-determined form. A very similar type of flexibility was evident in an anecdote which I came across concerning one of the most highly rated US Army units of the Korean War. The senior NCOs of the unit recognized that they lacked the qualities to lead men in action. When the unit went into combat, local command passed to a small group of enlisted men. Afterwards, these 'combat leaders' were quite happy to follow the orders of the NCOs, whose skills in everyday soldiering they fully recognized.

There is, then, convincing evidence that an organization's ability to switch from a bureaucratic, centralized mode to a more decentralized professional mode is an important determinant of reliability—or even survival. But how can it be engineered? Karl Weick—whose work has been cited at various points throughout this book—has made a number of important observations in this regard. In order to achieve effective decentralization—of the kind described earlier—Weick argues that:

> ... you first have to centralise so that people are socialised to use similar decision premises and assumptions so that when they operate their own units, these decentralised operations are equivalent and co-ordinated. This is precisely what culture does. It creates a homogeneous set of assumptions and decision premises which, when they are invoked on a local and decentralised basis, preserve co-ordination and centralisation. More important, when centralisation occurs via decision premises and assumptions, compliance occurs without surveillance. This is in sharp contrast to centralisation by rules and regulations or centralisation by standardisation and hierarchy, both of which require high surveillance. Furthermore, neither rules nor standardisation are well equipped to deal with emergencies for which there is no precedent.[38]

It is probably no coincidence that the HROs studied by the Berkeley group were either military or had many key personnel with a military background—this applies equally to the third HRO not discussed above, a Californian nuclear power plant in which many operators and supervisors had been in the nuclear Navy. The acceptance of a disciplined approach to working, well founded trust in SOPs, and a familiarity with the ways of rank-based structures would all help to forge the shared values about reliability that permit effective decentralized action when the occasion demands.

Weick makes another point of considerable relevance here. All hazardous technologies face the problem of requisite variety—the variety that exists in the system exceeds the variety of the people who must control it (see Chapter 4). As a result, 'they miss important information, their diagnoses are incomplete, and their remedies are short-sighted and can magnify rather than reduce a problem'. But this problem, can be reduced by a culture that encourages 'war stories'. Since the nature of these systems allows little scope for trial-and-error learning, maintaining reliability depends on developing alternatives for trial and error. These could include imagination, vicarious experience, simulation, stories and story-telling.

> A system that values stories and storytelling is potentially more reliable because people know more about their system, know more of the potential errors that might occur, and they are more confident that they can handle those errors that do occur because they know that other people have already handled similar errors.[39]

Other ways of reducing the gap between the variety of the system and the variety of its human controllers include:

- *A culture that favours face-to-face communication.* 'One way to describe (admittedly stereotype) engineers is as smart people who don't talk. Since we know that people tend to devalue what they don't do well, if high reliability systems need rich, dense talk to maintain complexity, then they may find it hard to generate this richness if talk is devalued or if people are unable to find substitutes for talk (e.g., electronic mail may be a substitute).'[40]
- *Work groups made up of divergent people.* 'A team of divergent individuals has more requisite variety than a team of homogeneous individuals.'[41] It matters less what makes up this diversity—different speciality, different experience, different gender, and the like—than the fact that it exists. 'If people look for different things, when their observations are pooled they collectively see more than any one of them alone would see.'[42] By the same token, groups made up of very similar people tend to see very similar things, and so lack requisite variety.

The decentralization of authority under certain conditions was a crucial feature of the German military concept of *Auftragssystem*—(mission system) discussed in Chapter 4. Its essence was that a subordinate commander, a subaltern or senior NCO, should be trained to a level where he (or, very rarely, she) could achieve the tactical goals of superior officers, with or without orders. Translating this into a civilian context, it means selecting and training first-line supervisors so that they are able to direct safe and productive working without the need for SOPs. Such a localized system of behavioural guidance makes heavy demands on the personal qualities of the supervisors. A prerequisite is an extensive experience of the jobs carried out in the workplace and the conditions under which they are likely to be performed. Supervisors need to be 'sitewise' both to the local productive demands and to the range of obvious and less obvious hazards. Equally important is a personal authority derived both from the respect of the workforce and the support of management—a key feature in the success of the German Army.

Not all activities in hazardous technologies are carried out in supervised groups. When people are relatively isolated, the onus shifts from group to self-controls. Crucial among these are the techniques designed to enhance hazard awareness and risk perception, These are the measures that seek to promote 'correct' rather than merely 'successful' performance. A number of hazard evaluation programmes are being developed or have already been implemented. However, as Willem Albert Wagenaar has observed,[43] risk appraisal training is of little value once the incorrect actions have become habitual. When this happens, people are not *taking risks* deliberately, they are *running risks* in a largely thoughtless and automatic fashion. To be effective, such training must occur in the initial phase of employment and then be consolidated and extended by on-the-spot supervisory guidance. By the same token, it is mainly through local supervisory interventions that long-established pattern of incorrect behaviour can be modified.

In summary, high-reliability organizations are able to shift from centralized control to a decentralized mode in which the guidance of local operations depends largely upon the professionalism of first-line supervisors. Paradoxically perhaps, the success of this transformation depends on the prior establishment of a strong and disciplined hierarchical culture. It is the shared values and assumptions created by this culture that permit the coordination of decentralized work groups. Effective teams, capable of operating autonomously when the circumstances demand it, need high-quality leaders. This, in turn, requires that the organization invest heavily in the quality, motivation and experience of its first-line supervisors.

Engineering a Learning Culture

Of all the 'subcultures' so far considered, a learning culture is probably the easiest to engineer but the most difficult to make work. Most of its constituent elements have already been described—observing (noticing, attending, heeding, tracking), reflecting (analysing, interpreting, diagnosing), creating (imagining, designing, planning) and acting (implementing, doing, testing). The first three are not so difficult. It is the last one—acting—that is likely to cause most of the problems. Echoing the rueful remark by the man from Barings Bank after the collapse—there always seemed to be something more pressing to do.

Beyond what has already been written,[44] there is little more that a book can do to give top managers the will to put in place the reforms indicated by their safety information systems, except to bring to their attention the chilling observation of the organizational theorist, Peter Senge:

> Learning disabilities are tragic in children, but they are fatal in organizations. Because of them, few corporations live even half as long as the person—Most die before they reach the age of forty.[45]

Senior managers should not need to be reminded that an organizational accident can brutally cut short even that brief span.

Safety Culture: Far More than the Sum of its Parts

At this point, I have in mind an imaginary technical manager from an organization with a good safety record (probably measured in LTIFs) who starts to count off the cultural elements that have so far been considered. Yes, he or she might decide, we have an incident reporting system of sorts. Yes, we have a reasonably fair and straightforward method of deciding whether or not disciplinary action is warranted. Yes, we have, on occasions, allowed our first-line supervisors a good deal of latitude and backed up their decisions afterwards—when things turn out all right, of course. And, yes, we have implemented a number of fairly expensive safety improvements on the basis of both reactive and proactive information, so it could be said that we have a learning culture. Does all of this mean that we have an informed culture—or, in more usual terms, a safety culture?

As any engineer knows, assembling the parts of a machine is not the same thing as making it work. And the same is even more true of social engineering than of its more mechanical counterparts. In order to answer our hypothetical manager, we would have to pose some questions in return:

- Which board members have responsibility for organizational safety—as opposed to conventional health and safety at work concerns?
- Is information relating to organizational safety discussed at all regular board meetings—or their high-level equivalent?
- What system, if any, do you have for costing the losses caused by unsafe acts, incidents and accidents?

- Who collates, analyses and disseminates information relating to organizational safety? By how many reporting levels is this individual separated from the CEO? What annual budget does this person's department receive? How many staff does he or she oversee?
- Is a safety-related appointment seen as rewarding talent (a good career move) or as an organizational oubliette for spent forces?
- How many specialists in human and organizational factors does the company employ?
- Who decides what disciplinary action should be meted out? Are the 'defendant's' peers and union representatives involved in the judgement process? Is there any internal appeals procedure?

The potential list is endless. The point is this—the mere possession of the 'engineered' externals is not enough. A safety culture is far more than the sum of its component parts. And here—perversely perhaps, considering what was said at the beginning of the chapter—we must acknowledge the force of the argument asserting that a culture is something that an organization 'is' rather than something it 'has'. But if it is to achieve anything approaching a satisfactory 'is' state, it first has to 'have' the essential components. And these, as we have tried to show, can be engineered. The rest is up to the organizational chemistry. But using and doing—particularly in a technical organization—lead to thinking and believing.

Finally, it is worth pointing out that if you are convinced that your organization has a good safety culture, you are almost certainly mistaken. Like a state of grace, a safety culture is something that is striven for but rarely attained. As in religion, the process is more important than the product. The virtue—and the reward—lies in the struggle rather than the outcome.

Postscript: National Culture

Every organizational culture is shaped by the national context in which it exists—and this is especially true for multinational organizations. It is not within the scope of this chapter to deal with the differences in national culture. For this, the reader is directed to the seminal books by Geert Hofstede.[46] The interested reader is also strongly advised to seek out the work of Robert Helmreich[47] and his colleagues at the University of Texas, and of Najmedin Meshkati[48] at the University of Southern California.

Notes

1 Data from the Flight Safety Foundation Icarus Committee, cited by Skandia International. I am grateful to Lars Högberg of the Swedish Nuclear Power Inspectorate (SKI) for sending me this information.
2 B. Uttal, 'The corporate culture vultures', *Fortune*, 17 October 1983.
3 T.E. Deal and A.A. Kennedy, *Corporate Cultures: The Rites and Rituals of Corporate Life*, (Reading, MA: Addison-Wesley, 1982).
4 T.J. Peters and R.H. Waterman, *In Search of Excellence: Lessons from America's Best-Run Companies*, (New York: Harper & Row, 1982).
5 Ibid.
6 Ibid., p. 76.
7 P. Bate, 'The impact of organizational culture on approaches to organizational problem-solving', in G. Salaman (ed.), *Human Resource Strategies*, (London: Sage, 1992), p. 229 cited by N. Thompson, S. Stradling, M. Murphy and P. O'Neill, 'Stress and organizational culture', *British Journal of Social Work*, **26**, 1996, pp. 647–65.
8 Cited by Thompson *et al.*, op. cit. p. 651.
9 G. Hofstede, *Cultures and Organizations: Intercultural Cooperation and its Importance for Survival*, (London: Harper Collins, 1994), p. 199.
10 International Nuclear Safety Advisory Group (IAEA), *Safety Culture*, (Vienna: IAEA, 1991).
11 Cited by R. Booth, 'Safety culture: concept, measurement and training implications', *Proceedings of British Health and Safety Society Spring Conference: Safety Culture and the Management of Risk*, 19–20 April, 1993, p. 5.
12 M. O'Leary and S.L. Chappell, 'Confidential incident reporting systems create vital awareness of safety problems', *ICAO Journal*, **51**, 1996, pp. 11–13. Dr Chappell is now Program Manager, Human Factors Services, at TransQuest Inc. in Atlanta, Georgia.
13 Ibid., p. 11.
14 S.L. Chappell, 'Aviation Safety Reporting System: program overview' in *Report of the Seventh ICAO Flight Safety and Human Factors Regional Seminar*, Addis Ababa, Ethiopia, 18–21 October, 1994, pp. 312–53.
15 J.A. Passmore, 'Air safety report form', *Flight Deck*, Spring 1995, pp. 3–4.
16 M. O'Leary and N. Pidgeon, 'Too bad we have to have confidential reporting programmes', *Flight Deck*, Summer 1995, pp. 11–16.
17 M. O'Leary and S. Fisher, *British Airways Confidential Human Factors Reporting Programme. First Year Report, April 1992–March 1993*, (Hounslow: British Airways Safety Services, 1993).
18 S. Chappell, op. cit.
19 M. O'Leary and S. Chappell, op. cit.
20 Ibid., p. 12.
21 M. O'Leary, 'New developments in the British Airways Confidential Human Factors Reporting Programme', *Flight Deck*, Summer 1996, pp. 19–22.
22 Ibid., p.21.
23 Cited by S.J. Guastello, 'Do we really know how well our occupational accident prevention programs work?', *Safety Science*, **16**, 1993, pp. 445–63.
24 O'Leary, op. cit. 1996, p. 22.
25 J.C. Smith and B. Hogan, *Criminal Law*, (3rd edn), (London: Butterworths, 1975), p. 45.
26 Ibid., pp. 45–6.
27 D. Marx, Personal communication, 9 January 1997. I am most grateful to David

Marx for his helpful advice on the distinction between recklessness and negligence.
28 Cited by M. O'Leary and N. Pidgeon, op. cit., p. 16.
29 N. Johnston, 'Do blame and punishment have a role in organizational risk management?', *Flight Deck*, Spring 1995, pp. 33–6.
30 Air Accident Investigation Branch (AAIB), *Boeing 747-136 G-AWNJ. Report on the Incident near Nairobi Airport, Kenya on 3 September 1974*, Aircraft Accident Report 14/75, (London: HMSO, 1975).
31 J.M. George, 'Asymmetrical effects of rewards and punishment: the case of social loafing', *Journal of Occupational and Organizational Psychology*, **68**, 1995, pp. 327–28.
32 D. Marx, Personal communication, 1 October 1996.
33 T.R. LaPorte and P.M. Consolini, 'Working in practice but not in theory: theoretical challenges of "high-reliability" organizations', *Journal of Public Administration Research and Theory*, **1**, 1991, p. 21.
34 Ibid., p. 27.
35 Ibid., p. 29.
36 Ibid., p. 32.
37 Ibid., p. 34.
38 K.E. Weick, 'Organizational culture as a source of high reliability', *California Management Review*, **24**, 1987, pp. 112–27, see p. 124.
39 Ibid., p. 113.
40 Ibid., p. 115. See also K.M. Eisenhardt, J.L. Kahwajy and L.J. Bourgeois, 'Conflict and strategic choice: How top management teams disagree', *California Management Review*, **39**, 1997, pp. 42–62.
41 K.E. Weick, op. cit., 1987, p. 116.
42 Ibid., p. 116.
43 W.A. Wagenaar, 'Risk-taking and accident causation' in J. Yates (ed.), *Risk-Taking Behaviour*, (Chichester: Wiley, 1992).
44 The interested reader is recommended to start with: B. Toft and S. Reynolds, *Learning from Disaster: A Management Approach*, (London: Butterworths Heinemann, 1994).
45 P.M. Senge, *The Fifth Discipline: The Art and Practice of the Learning Organization*, (London: Century Business, 1990).
46 G. Hofstede, op. cit., 1994. See also G. Hostede, *Culture's Consequences: International Differences in Work-Related Values*, (Beverly Hills, CA: Sage Publications, 1980).
47 R.L. Helmreich, A. Merritt and P. Sherman, 'Research project evaluates the effect of national culture on flight crew behaviour', *ICAO Journal*, **51**, 1996, pp. 14–16 (and many other sources).
48 N. Meshkati, 'Cultural factors influencing safety need to be addressed in design and operation of technology', *ICAO Journal*, **51**, 1996, pp. 17–18 (and many other sources).