

of participation are not worthless, and we have shown that they bring real benefits, but these benefits remain tightly constrained.

As for public policy, participatory approaches are clearly more strongly embedded in countries such as Germany or Sweden than they are in the UK or the USA. If policy is to promote participation, then the conditions where it thrives need attention. Such an approach is very different from that of much public policy, which entails the identification of 'best practice' and exhortation to emulate it. If conditions are not appropriate, emulation will be absent or will fail. One key feature of Germany is that institutions such as works councils are well established, which means that firms have to adjust how they behave: the institutional context shapes assumptions and the ways in which choices are made. By contrast, in the UK and the USA, assumptions about 'shareholder value', the need to retain commercial confidentiality, and the right and obligation of management to take key decisions are equally strongly embedded. Public policy would need to find ways to shift such assumptions as well as to promote the economic conditions which sustain organizations that stress status over contract.

Why Do Disasters Happen?

Organizations are rationally designed systems pursuing defined ends. So why do things go wrong?

A familiar list of spectacular failures would include the *Challenger* shuttle disaster in the USA in 1986 (7 crew were killed), the Piper Alpha oil rig in the North Sea in 1988 (167 dead), and the explosion of the Union Carbide chemicals works at Bhopal in India in 1984 that left between 2,500 and 10,000 people dead (estimates vary widely as to the exact number; the catastrophe is probably the largest single industrial disaster, leading to more immediate deaths and illnesses than the Chernobyl nuclear disaster of 1985, which killed immediately only 30 people). Less spectacular cases of organizational failures include computer systems that fail to work: the Taurus electronic trading system at the London Stock Exchange was initiated in 1986 and abandoned seven years later having run substantially over budget. In none of these cases was there a deliberate attempt to destroy the operation, and they all occurred either in prestigious public organizations or in large MNCs.

There are two distinct processes at work here. The first concerns operational mistakes and miscalculations that can have catastrophic consequences. The second covers decision-making processes that lead to failure. In the first case, the goal is clear and unquestioned (running an oil rig safely) but the execution fails. In the second case, there is a choice of goals and in retrospect it appears that the wrong one ended up being chosen. These two issues are addressed, respectively, in this chapter and in Chapter 8.

The criminologist Box (1983: 26-8) put the issue of workplace danger in sharp perspective. He calculated the number of deaths in the UK between 1973 and 1979 from fatal accidents and occupational diseases. The number was 11,436, which he set against the number of homicides (3,291). Allowing for differences in populations at risk (occupational hazards apply only to those in employment), the ratio was of the order of 7:1. Box also argues that there may be serious underrecording of the occupational data. Box appears to have intended the comparison to be taken literally, as comparing different sets of avoidable deaths. The degree of

intentionality and the social processes involved are, however, very different. Nonetheless, the data underline the extent of industrial injury and how little attention is paid to it.

Apart from its inherent importance, the issue of mistakes and errors is valuable for several reasons. First, it is rarely addressed in studies of organizations. Not only is it absent from many standard textbooks, but it is also not discussed in a volume aiming to lay out the key approaches to strategy, even though overcoming errors might be seen as central to strategy (Mintzberg et al. 1998). Nor is it mentioned in a book directed explicitly at what it terms organizational misbehaviour (Ackroyd and Thompson 1999). The reasons for this need not detain us here, but anyone using this book to study analytical approaches to organizations may care to think of some. Second, it throws light on many important themes of how people see the world and how they relate to each other in organizations. For example, it brings into very sharp focus what Weick (1993) calls false hypotheses: when we expect to see or hear certain things, we interpret evidence in the light of such expectations, and fill in gaps to tell a consistent story. This tendency can have severe consequences in certain conditions. Third, it helps us fill out and develop the themes of power and organizational negotiation introduced in Chapters 5 and 6. A key point about disasters is that they are attempts to keep organizations working, and thus illustrate power as a capacity. But they are also about struggles over different definitions of organizational processes, thus reflecting power in the sense of domination of one group by another.

One other remark is needed. Weick identifies a psychological process. Such a process is likely to operate anywhere. Indeed, as shown below, the need for interpretative schemes characterizes science as much as it does everyday life. They also need to be placed in the context of structural factors to do with technology and the distribution of power. A focus on psychology can lead to a tendency to 'blame the operator', whereas we also need to ask why the operator was placed in a certain position.

We thus begin by examining systems and then placing them in a wider political economy perspective. Our remark above, that there is a clear goal but a failure of execution, is a first approximation, for it is often the case that safety is much more contested and political than it implies. Analyses based on political economy often stress conflicts of interest almost exclusively, while more conventional approaches neglect them or treat them as unfortunate interferences with the pursuit of joint goals. Research that stops with the approximation is illuminating and often neglected, and it should be considered before wider issues are addressed.

We proceed as follows. First, some writers see science or organizational rationality as having run away from social control, with danger being inherent in modern society. We address two versions of this perspective in the first two sections of the chapter. We then turn to the opposite position, that high reliability is feasible. The rest of the chapter develops a more adequate analysis, first by placing disasters in the context of technology and then examining two specific perspectives, 'man-made disasters' and 'normal accidents'. The final section draws the argument together through a political economy of safety and risk.

Administrative Evil?

The broadest attempt to explain failures locates them in what has been termed administrative evil, which is in turn seen as an inherent component of the technical rationality of modern societies (Adams and Balfour 1999). Technical rationality is a key issue for large organizations, since they are the archetypical example of this approach. It embraces predictability, the following of rules, and formal procedures and standards of behaviour. Administrative evil is defined as a process through which outcomes causing suffering to others arise even though people may not be aware of these results because their actions are masked by technical rationality. Simply 'following orders' may entail evil.

For these authors, a critical case is the Holocaust. Along with others (e.g. Bauman 1989), they see the Holocaust, not as a dreadful aberration, but as inherent in modernity and technical rationality. For Bauman (1989: 17), it grew out of a bureaucracy true to its 'form and purpose' and was a central part of modernity and indeed could not have occurred under any other system. Adams and Balfour (1999: 167) stress that the administrators were following orders and procedures and that everything was legally sanctioned.

Explaining how it was that the Holocaust could happen and how 'normal' people were drawn to engage in it is an important activity. It is true that an administrative machine was put in place and that people who played their part did not ask why they were, for example, driving railway trains to extermination camps, focusing instead on the fact that they were 'just doing their job' or 'following orders'. Yet the causal chain here is too tight. Not all technical-rational organizations have inherent in them the risk of evil on a large scale. Much has been made of the 'following orders' argument, but under the exigencies of a repressive regime in wartime it is understandable why people would make such an argument, not least because refusal to follow orders could lead to severe punishment. In

other conditions, it is reasonable to expect that people will be able to question orders. This is not to say that they will, as we discuss below. But any bureaucracy with formal procedures is likely to contain some means by which orders may be questioned and reservations noted. The whole point of professional authority is that there are standards of behaviour, and that professionals do not follow orders blindly.

Indeed, an ironic result of the model of administrative evil is that everything is explained by it and genuine evil is forgotten. Means and ends become confused. Administrative evil involves the rational pursuit of means without any questioning of ends. But the ends were prescribed by some people. This is not to say that the shape of the 'final solution' existed fully formed in prior plans; like any other complex process, its nature evolved over time. But there was nonetheless a stated end. The Holocaust surely needs explanation in terms of specific policies and choices, rather than being treated as the direct result of 'modernity'.

At this point, we should be clear what 'going wrong' means. One review of the literature observes that 'routine non-conformity, mistakes, misconduct and disasters are not anomalous events' (Vaughan 1999: 298). This is true, but it is unhelpful to lump different phenomena together. 'Misconduct' is deliberate flouting of rules or expectations, while mistakes and disasters are unintentional. Just how to identify misconduct is also much less clear than might appear, as we discuss below. This point has been made in respect of Vaughan's own analysis of the *Challenger* case (Perrow 1999: 379-80; the reference is to Vaughan 1996). The case is now widely used in studies of organizations. Its essence is simple and is set out in Box 7.1.

For Adams and Balfour, and for Vaughan, the disaster can be explained in terms of the social construction of reality (see Chapter 5), wherein a bureaucracy allowed the 'normalization' of deviations from safe procedures. Yet, as Perrow argues, to rely on social construction is to minimize the distinctive events in this case, and to treat it as an emblem of any form of organizational practice anywhere. The deviation from procedure (launch under cold conditions) was in fact unprecedented and not normal. And the launch was opposed by a group of engineers, who knew that conditions were outside parameters that they had investigated, rather than being deeply embedded in an unquestioned culture. What stands out for Perrow is the exercise of power by other people over the objections of the engineers: 'we miss a great deal when we substitute culture for power' comments Perrow dryly.

Let us pursue social construction in relation to data on industrial injury. It is true that all kinds of definitional processes are involved in deciding just how severe an injury needs to be before it is reported. And there are

Box 7.1 The *Challenger* events

The shuttle exploded soon after take-off. The immediate cause was found to be a failure of an 'O' ring that was meant to seal the rocket fuel; leaking fuel was ignited. Further investigation found that problems with the rings had been identified, and addressed by including two sets of rings. On this occasion, however, the launch took place in very low temperatures, which meant that the rings were less flexible than under standard conditions.

Reasons offered as to why the launch took place instead of being postponed include the fact that there had been several delays already and that this was the first flight to feature a civilian, which was important for establishing the political legitimacy of the shuttle programme, not least because the President of the USA was scheduled to speak to the civilian when she was in orbit.

Source: Vaughan, (1996); Adams and Balfour (1999).

convincing arguments that the weakening of trade unions in many countries reduces the strength of workers and hence their willingness to risk reporting injuries; the result is that reported declines in injury rates may be exaggerations (Nichols 1997). Yet it is also accepted that the more severe injuries will tend to be reported, and that figures on them are likely to be relatively reliable. An interesting illustration given by Nichols is summarized in Box 7.2: the nature of data can be explained and their value grasped, rather than being relabelled as mere exercises in social construction. (Box 7.4 provides another example.)

In short, to rely too much on social construction is to play down distinctive structural influences and processes. Disasters are neither unique nor direct reflections of the modern condition. The politics of their occurrence deserve attention.

Box 7.2 Interpreting injury statistics

In 1969 iron ore miners in Sweden went on strike to end the piecework pay system. Following their success, recorded severe injuries fell while minor ones rose. This can be accounted for by two effects. Piecework is well known to induce risk because it encourages the cutting of corners to raise earnings, and its ending cut the rate of severe injury. Workers now also had the time to report less severe injuries, since they would no longer suffer a pay penalty from being away from production. Statistics need to be questioned and interrogated but not rejected out of hand.

Source: Nichols (1997: 84).

Cultures of Fear

There is a widespread and often misplaced 'culture of fear' leading people to be 'afraid of the wrong things' (Glassner 1999). Scares about issues as diverse as airline crashes, medical procedures, and crime are often exaggerated; for example, fear of crime increases when the available evidence of crime rates suggests that the risk is falling. Glassner explains this in terms of the following factors:

- *Media hype and exaggeration.* Although some parts of the media correct the errors of others, in some cases a media frenzy can be identified. For example, cases of child abuse have led to possible solutions of unproven effect and unintended consequences such as mob violence against people suspected (often wrongly) of being abusers.
- *Self-interest.* People selling security systems, for example, have an interest in preying on fears of crime.
- *Preying on uncertainty and doubt.* A classic example is the *War of the Worlds* radio programme of 1938, in which a portrayal of a Martian invasion was taken by many people as being a real event. One reason for this was that the radio was perceived as an authoritative and reliable source. But Glassner also makes the more subtle point that some people saw the programme as an analogy: with fear of war at the time, the Martians were taken to represent Germany or Japan.
- *Scapigoating.* The point about analogies links to the attribution of blame. Glassner quotes an example from fourteenth century Europe when Jews were accused of poisoning wells. Impure water had been a long-standing problem, but when for some reason it came into focus it was desirable to find an explanation, and blaming a group such as the Jews was a convenient solution.

In developed countries, at any rate, various indicators (such as average life expectancy and numbers of deaths in road traffic accidents despite rising traffic densities) have improved. This is not a Panglossian view that everything is for the best in the best of all possible worlds: new technologies bring risks which need discussion and debate, and in many cases a specific innovation may have broad and unexpected effects that need to be taken into account. The point at present is to flag the danger of 'technophobia' (see Wajcman 2004) and to stress that our interest in disasters is not driven by the pessimism of such a position.

Risk takes on particular forms, and, crucially, there are hidden aspects of organizational functioning that help us to understand the sources of

organizational failures. This point links to debates about the 'risk society', a concept introduced by Beck (1992). Beck does not claim that modern societies are necessarily more hazardous than those of the past. Instead, in pre-modern societies dangers were taken as pre-given. With the rise of instrumental rational control, unknown dangers become calculable risks. For Beck, 'risks always depend on decisions' (quoted in Elliott 2002: 295).

Given the popularity of Beck's ideas, we have here a clear rationale for the focus in this chapter and in Chapter 8 on decisions in organizations, for it is organizational decisions, rather than those of individuals, that shape the modern world. Yet Beck, like many writers identifying a shift from 'organized' to 'disorganized' capitalism or from 'modernism' to 'post-modernism', also argues that groups such as families and social classes are no longer stable sources of moral order. The individual is increasingly responsible for her actions and is able to reflect on the situation in which she finds herself and to act as a result. She is not the prisoner of structural forces but is able to act in relation to them. Hence the concept of 'reflexive modernity', meaning for Beck not just reflection on the state of the world but confrontation with oneself as to how one should act.

At this point, we need to be more precise about the meaning of 'risk'. There is a well-established distinction between risk and uncertainty.

- Risk refers to the probability of an event from a known distribution. A gambler on a roulette wheel faces a known set of odds. The chances of road traffic accidents can also be stated as one in so many thousand miles travelled, though here there is no defined probability distribution as there is in the roulette case. Similarly, it makes sense to estimate the probability that women, as compared to men, will attain top positions in organizations, or that a certain proportion of new businesses will fail within a year of start-up.
- Risk becomes uncertainty when the distribution is not known. The risks at roulette would be uncertainty if the table were fixed and the croupier could at will determine where the ball landed. Many aspects of organizational operation are stated in terms of risk, when they are in fact closer to uncertainty. It is, for example, common to estimate the probability of two events occurring as one in M million. But we would first need to be sure that there was reasonable evidence in relation to each event and, crucially, that there was an underlying statistical theory at work. If events are truly independent, the probability of them both occurring can be calculated by multiplying the two probabilities. But if the events are not independent

(e.g. if two safety devices are connected in an unexpected way, or even more broadly if there is an influence common to them both, they occur in a company with a poor safety culture) improbable events become much more likely. Even more fundamentally, the chances of some events are not known. The safety record of nuclear power is often stated in terms of numbers of years of accident-free operation. But power plants are not roulette wheels. There are many different designs, so that the life of any particular one may be short. They have parts which interact in complex and, as we will see, unpredictable ways. And they are subject to human intervention that can interfere with apparently objective odds.

As we will see, debate on safety often turns on treating uncertainty as risk. Whether this approach is defensible should be considered carefully. The ready translation makes safety assessment seem scientific when it may not be. The 'risk society' is more accurately termed the 'uncertainty society'.

While Beck's project may be a useful source of ideas for social theory, it is hard to make any judgement as to the accuracy or otherwise of what he says. As one commentator remarks in a matter-of-fact and routine way, writers such as Beck 'do not seek empirical support for their arguments' (Savage 2000: 105). Yet this is surely key. We would expect either evidence of long-term historical trends, to enable us to say how reflexive modernity differs from a presumed past, or a detailed anatomy of the current condition, indicating, for example, what countries or economic sectors are more or least prone to reflexive modernity. In one respect, what follows can be read as an attempt to apply Beck's ideas, in that we aim to examine in detail the nature of decisions and their political context. We addressed parallel issues in Chapter 4. But we will argue that there are very clear limits on individuals' power to comprehend and act on the world, precisely because it is structured by organizations and structures of power.

Reliable Systems

Complex systems are remarkably reliable: plane crashes and nuclear accidents are far from everyday occurrences. A line of theory, characterized usefully as high reliability organizational theory (or HROT to give it an unflattering acronym) by Sagan (1993) stresses four points:

- Safety is prioritized as a goal by leaders of organizations.
- There are high levels of 'redundancy', meaning that complex systems are designed with at least primary and back-up safety systems.

- A high reliability culture can be created, wherein decisions are decentralized to appropriate levels, people are trained, and an expectation of commitment to safety is established.
- Trial and error leads to the elimination of mistakes and to organizational learning.

Any critical organizational theorist could easily ridicule assertions stated this starkly. But Sagan does not follow this simple route. He carefully sifts evidence, and accepts that there are examples consistent with the theory. It is, for example, true that there has never been an accidental nuclear war and that safety systems in nuclear weapons programmes have functioned to prevent potentially catastrophic failures. A programme of research has demonstrated how high reliability systems can work (e.g. Roberts 1990). At a more mundane level, we have observed plants engaged in dangerous processes (aluminium smelters, where not only is the molten metal hazardous but also the raw material and some of the waste products are potentially carcinogenic, and chemicals factories) where there was a genuine safety culture. How one might identify such a culture is discussed when the reader has more evidence available.

The point of entering criticisms and qualifications is not to swing from complacency to panic. In many respects, work processes are safer than they were in the past. Yet they are not necessarily safe, and the purpose of a critical analysis is to help in thinking about what may be hidden, not to suggest that disasters are lurking wholly unchecked.

There is one key point in how we might think about failures. A thought experiment might entail a random sample of a set of a given 'type' of organization (however 'type' is defined). One might then assess cases where failures had occurred and investigate what they had in common. Yet the available evidence starts from the other end of the causal chain: cases that are big enough, or uncommon enough, to attract attention, or that simply happen to emerge because an investigator brought them to light. Committees of inquiry into disasters often tell a story of how one decision, perhaps inconsequential in itself, interacted with others to produce a chain of events that led to the disaster in question. This is perfectly sensible if we wish to explain a particular event, and indeed it is not just cool *explanation* that is sought, for those involved in a tragic event need to gain *understanding* (that is, embracing feeling and emotion as well as logic). But it can suggest, at the same time, two logically incompatible implications. The first is that this was a unique set of events. The second is that a similar set of conditions would spark off an identical chain reaction.

What is needed is an account of causal chains and appreciation that at various key points there were alternatives, together with consideration of why these were not taken. We may then be in a position to understand how people in organizations can recognize potential disasters and even do something about them.

Understanding Technology

Why things go wrong is often to do with the relationship between social and technical systems. Understanding disasters will be helped by considering the nature of technical systems, for these are often approached as though they are hard realities. The situation is more subtle and complex than is often thought.

There are two simple stories about the links between technology and society. The conventional one says that technology advances through its own logic: inventions occur as a result of the pursuit of scientific proof (and in the language of economics they can be treated as 'exogenous', with 'technological progress' being outside social science interest). A common corollary is that technology is generally beneficial, though a rather different position states that technology is socially neutral and that its uses in practice can vary between the good and bad. The opposite story sees technical progress as a battle between those using technology to advance their own interests. According to this version, for example, the introduction of computer-controlled machine tools and computer-aided design was a device for capitalists to wrest control of production from the skilled pattern-makers, toolmakers, and designers who used to perform the relevant tasks.

There is a substantial literature that moves beyond these two simple views. It shows that scientific discovery is itself a socially shaped process. For the purposes of this chapter we draw selectively from this literature, focusing on the politics of technical innovation. This enables us to extend the discussion of power in Chapter 6: the politics of technology is about the power to shape definitions, and not about directly opposed interests (as the second story above would have it).

Mackenzie (1996: ch. 1; see also Mackenzie and Wajcman 1999) summarizes the central conclusions to emerge from the field of science and technology studies. First, a technology that eventually emerges is not the 'best' in any clear sense, for we need to ask 'best for whom?' This question has also been refined, for it is not simply a matter of overt dispute: even 'homogeneous' groups have differing definitions of goals and of how to attain them. In the terms of the present discussion, there can be broad

agreement on a final objective, but the goals needed to meet it may be unclear and thus open to different definitions; and what they are is likely to emerge only over time.

Second, technologies are often productive only when they are widely used. Particular technologies are 'best' because they won out in contests with others and therefore became generalized, rather than being best and therefore winning. It is an open question whether the benefits of a technology are intrinsic to it.

Third, belief in the potential success of one technology can lead to a concentration of effort on it. MacKenzie acutely gives the celebrated example of Moore's Law, which states that microchip processing power tends to double annually. This is not a fixed law of nature: because it became accepted, people invested heavily in new innovations in the belief that increased processing power would emerge.

Fourth, knowledge is central, and social science helps us understand how knowledge is produced and defined. We should stress that this argument is not that knowledge is simply socially constructed and relative, or that choices are wholly open. To say that technology is socially shaped is not to say that all processes of social definition are equally effective. It is to point to politics and negotiation as key processes through which technical possibilities were or were not put into practice.

These points are illustrated by MacKenzie's (1990) study of guidance systems for nuclear missiles. The key relevance here is that the systems were presented as offering more and more precise targeting of weapons; yet such claims had little basis in evidence, so that the systems were far from infallible. If things were to go wrong (and fortunately we have yet to find out) it would not be a question of people responding inappropriately to well-designed machines, but of how the machines were conceived.

- The invention of inertial guidance was not a question of scientific inspiration, for the method was widely considered impossible and it was only when a 'need' emerged from the military that scientific attention was directed at the problem.
- The ability of ballistic missiles to reach a target was never established. Test conditions are not indicative of what might occur under wartime conditions. For example, it was not possible to show conclusively that missiles would survive intense cold and then intense heat without exploding before reaching the target.
- And even if a missile survived, the accuracy of hits cannot be known, once varying trajectories and conditions and random error, together with differences between tests and real conditions, are allowed for.

A final issue here concerns technology and capitalism. We have drawn on some of Marx's ideas. Further elaboration for those interested in this issue is given in Box 7.3.

Box 7.3 Marx and machinery

A standard reading of Marx is that he was a technological determinist. What this means is a view that technical change determines social development. It matters because such determinism is faulty (though still not uncommon), hence if Marx took such a view much of his work could be dismissed. It also matters because, once the false reading is rejected, we can think about technology and society in more constructive ways.

The basis of the reading is the famous 1859 Preface to *A Contribution to the Critique of Political Economy*. MacKenzie (1996: ch. 2) offers a cogent alternative view (see also Cohen 1978: chs 2 and 6).

First, Marx's 'determinism' states that the forces of production determine the social relations of production. Yet the forces for Marx include labour and skills, and therefore cover much more than non-social technology. 'Determine', moreover, is best taken to mean 'influence' or 'set the conditions for', rather than directly cause. (Similar arguments that the forces determine in the sense of setting limits on the development of social relations, or facilitating some developments rather than others, are given by Katznelson (1986), and Edwards (1986: 61–2, 151–3)).

Second, Marx saw the linkage between machinery and the organization of work (the labour process) as follows. MacKenzie here highlights *Capital, Volume 1*, Parts 3 and 4.

- The origins of capitalism lie not in machines as such but in the emergence of a class of propertyless wage labourers. Simple cooperation within the labour process not only led to benefits of economies of scale but also tended to increase the capitalist's authority because of his central role in the coordination of production. (In the terms of Chapter 6, power to achieve goals and power over other people are intertwined. The many arguments as to whether the rise of capitalism was about productive capacity or intensified control of workers thus miss the point. It was about both. The two terms should, however, be specified more exactly.)

- 'Productive capacity' does not mean simply 'progress' because there will be different preferences in terms of how the capacity is used, not only between workers and capitalists (e.g. preferences for labour-saving or labour-using technologies) but also between capitalists (development down one particular path rather than another). Capacities can be put to use in different ways, and choices between these ways reflect political struggles.

- Control of workers' should not be envisaged in terms of more or less. As capitalism evolves, new balances of control and cooperation are developed, and the question of whether any given form is more exploitative than another is strictly meaningless. The question is meaningful only when tight comparisons can be made, so that, for example, it does make sense to say that a given group of workers was working harder than previously and that, if other conditions were unchanged, it was more exploited than it had been.

- The further division of labour reinforced subordination, and this set of social relationships created the space for mechanization. The limits of organizational change possible under less mechanized systems created the necessity for this process of mechanization.

Marx's account is a theory, and not an attempt at a detailed history. The theory nonetheless broadly fits the historical evidence.

Man-Made Disasters

Two useful approaches to understanding why disasters happen, then, are Turner's (1976) model of man-made disasters (MMD); the gendered terminology is apparently unconscious, though in view of the masculine assumptions of correctness and unwillingness to listen to criticism revealed in the evidence, the label may be more appropriate than Turner intended) and Perrow's (1999) normal accident theory (NAT). They have much in common, though Perrow and other users of NAT refer neither to Turner nor to the examples used by him, and they thus need to be described in turn.

Drawing on three examples from the UK, Turner identifies a set of common features. The key ones are as follows:

- Rigid perceptions and beliefs: This includes inattention to warning signs and a lack of communication. For example, in the case of the disaster at Aberfan, South Wales in 1966 (in which 144 people were killed when a spoil heap from a coal mine slid down a mountain side onto the village) a memo of 1939 had anticipated the causal conditions but the document was restricted and people were unaware of its significance.
- The 'decoy problem': a focus on a well-defined but minor issue.
- Disregard of outsiders to the organization, who were written off as inexperienced.
- Informational and coordination difficulties.

- Involvement of untrained people who behave in unexpected ways. For example, in the case of a fire parents of children who were trapped tried to enter the burning building.
- Failures to comply with existing regulations.

A key observation by Turner concerns the later official inquiries into the disasters:

[E]ach dealt with the problem that caused the disaster as it was later revealed and not as it presented itself to those involved beforehand. The recommendations, therefore, treat the well-structured problem defined and revealed by the disaster, rather than with [sic] preexisting, ill-structured problems. (1976: 93)

(Recall the quotation from Mintzberg in Chapter 5: much decision-making occurs under conditions of *ambiguity*, not risk or even uncertainty). For the understanding of why things go wrong, it is important not to fall into the same position as these reports. They have the benefit of hindsight and tend to imply that if certain errors were avoided—if, for example, communication had been less ambiguous—a disaster would not have happened. But Turner's point is that ill-defined situations necessarily exist. Plugging one gap will not necessarily resolve an issue. What inquiries discover are particular chains of events which are unique. They are certainly useful in dealing with particular classes of events so that, for example, spoil heaps in other Welsh valleys could be investigated, and action taken. But disasters can occur in any situations where two conditions hold: the background conditions render behaviour potentially dangerous and there is then a chain of events containing the features identified by Turner.

Normal Accidents

As to the relevant background conditions, NAT offers important insights. It is a subtle approach, throwing light not just on disasters but on many aspects of organizational functioning under more 'usual' conditions. We therefore discuss it in some detail.

Perrow (1999) begins with a homely example. You have an important job interview one morning and your housemate starts making coffee for you before leaving. But the coffee pot cracks; you make coffee again; and, now late, you shut the door with your car keys locked inside the house. You normally have a spare set but you have lent them to a friend (a design *redundancy* has failed, says Perrow). You seek a lift from a neighbour but he has to stay in today to have his generator fixed (this is a *loosely coupled* situation since the lost key and the generator are not connected). He tells

you there is a bus strike, which means that when you try to call a cab none is available (the strike and the lack of cabs is *tightly connected*). You fail to make it to the interview on time.

It is the interaction of multiple failures that causes the problem, rather than a single event. Tightly coupled processes require special attention, whereas loosely coupled ones are more contingent and hard to predict. This framework establishes the core of the approach, except that, says Perrow, in this example the person involved can understand these interdependencies. In complex systems, the linkages are too many, and the possible ways in which parts affect each other too complicated, for an operator on the ground to understand them. The interactions cannot be seen and 'even if they are seen, they are not believed... seeing is not necessarily believing; sometimes, we must believe before we can see' (Perrow 1999: 9). Perrow here refers to many writers on organizations who stress that we understand the world through mental maps. Sense data are not simply received but have to be interpreted and rendered significant. In one neat experiment, subjects were asked to register at a desk. The 'receptionist' ducked behind the desk and appeared to emerge with a form to be completed. The majority of the subjects did not notice that one person ducked down and another, wearing a markedly differently coloured shirt, appeared. This was because the details of the receptionist, though 'seen' literally, did not register as important.

Note that Perrow is using 'normal' ironically, to stress that accidents are embedded in organizational practice rather than being exceptional. As discussed below, he says that 'system accident' would be a more exact term, since the analysis turns on the structuring of systems. Some *common* accidents are not system accidents because they were simply the result of an individual failure. As we will also see, Perrow considers that the spectacular Bhopal case was not a system accident but simply a gross failure of attention to safety.

There are five main differences between tightly and loosely coupled systems:

1. Tightly coupled systems have time-dependent processes that cannot stand by until they are attended to. Reactions in chemicals plants cannot be delayed or extended.
2. The sequences in tightly coupled systems are relatively invariant. To produce a chemical entails a given sequence, whereas in an assembly operation the order of fitting parts can be varied more readily.
3. Tight coupling entails one way to meet a goal: an aluminium smelter can produce only aluminium, whereas in looser systems different

- product mixes can be included and different elements can be included in the mix (e.g. using plastic rather than metal).
- 4. There is little slack in tightly coupled systems. A failure entails a shutdown of the process.
- 5. Safety calls for buffers and redundancies (alternative systems if a main system fails such as a reserve parachute). In tightly coupled systems these have to be deliberately planned whereas in looser systems there are more ways of replacing one operation with another.

Examples of tightly coupled systems include dams and nuclear power plants. Loosely coupled systems include post offices and similar agencies and universities.

Complexity is Perrow's short hand for interactions in an unexpected sequence. In complex systems, the state of an operation can be confusing, as for example when a monitoring device gives ambiguous data. A response to a signal can set off a series of events whose cumulative path is hard to follow. Operators, moreover, are likely to develop a mental picture of what is happening and to ignore disconfirming data. A ship's pilot, for example, may interpret a set of lights as indicating another ship moving in the same direction, and take other data to confirm this, even though in fact he has misread the lights and they indicate a ship moving in the opposite direction (Perrow 1999: 215-17 gives a real-world example of this). Other aspects of complex systems include the fact that one component can interact with one or more other components either by design (for example, when a heater serves two other components) or by accident (as when oil from a ruptured tank enters an engine compartment and ignites). Perrow gives several other examples of the nature of complexity and the problems that can arise as a result. A safety device can itself contribute to accidents. For example, in 1966 an engineer at the nuclear plant in Monroe, Michigan, introduced a new device aimed to add to redundancy, a zirconium plate; but the plate broke and blocked a coolant pipe. A further issue here is that the plate was added late in the design and was not included in the specifications, which made it hard to establish the cause of the problem (Sagan 1993: 160).

Complexity does not, in itself, render a system accident-prone. A university is a complex system because there are multiple functions which can interact in unexpected ways. But it is also loosely coupled because there is time for recovery when things go wrong and consequences of errors are not usually dangerous. Perrow's classification of his two dimensions is shown in Figure 7.1.




FIG. 7.1 Interaction and Coupling
Source: Abbreviated from Perrow (1999: 97).

One key implication emerges from this approach to disasters in terms of the structuring of systems. Perrow notes that, in inquiries into the causes of disasters, up to 80 per cent attribute a problem to operator error. Yet there is usually more to the story. Long working hours and tiredness are one feature. Another is the fact that operators are under pressure to save time or costs, and may thus be led to cut corners. Perrow cites one of the first oil tanker disasters, that of the *Torrey Canyon* in the UK in 1967; part of the context was pressure on the captain to reach port as soon as possible and he therefore chose a perilous route. But, perhaps most fundamentally, signals from the environment can be misleading, so that incorrect actions are encouraged. If systems are tightly coupled, this can set off a chain of events that is unpredictable and unmanageable, especially when operators have to respond in a short time to unfamiliar sets of events.

We can further elucidate NAT and compare it with HROT. Sagan (1993) considers the four elements of HROT and shows how NAT would see them.

- *Safety as a goal.* People in organizations have differing objectives, those at the top may not be directly exposed to some hazards (e.g. those in the inner workings of chemicals plants), there will be pressures to maintain production, and elites may be misinformed as to the prevalence of near accidents. Collinson (1999) provides a particularly sharp example of the underreporting of accidents (Box 7.4).
- *Redundancy.* Safety systems can interact in unpredictable ways, as in the Monroe case cited above.
- *Decentralization.* Can all events be anticipated and trained for, and can local decision-makers be given accurate information to allow them to make judgements, particularly when choices at one location interact with those at others?
- *Learning.* Causes of accidents are often unclear, politicized environments often lead to blame cultures, and some organizations are riddled with secrecy. In such contexts learning will be constrained.

Box 7.4 Hiding injuries on North Sea oil rigs

Several oil companies operate drilling platforms far out in the North Sea. Workers typically stay on them for two-week periods, working twelve-hour shifts. In 1990, Collinson (1999) studied two of these rigs owned by one firm. The company had a reputation for a safety culture, and had an accident rate of about one-tenth the industry average. There were safety committees, and each crew going through a year with no accident received a bonus. There were two key groups of workers, company employees, and contractors.

About half the company employees interviewed said that they had concealed accidents or near-accidents in order to protect their appraisal rating. (For analysis of performance appraisal, see Chapter 5). Attempts were also made to downgrade the recorded severity of incidents that could not be concealed.

The contractors worked for specialist firms and did the most dangerous work, notably drilling and erecting scaffolding. They suffered the most accidents but were under greatest pressure to conceal them: a feared system of reports meant that they could be refused further work, if 'Not Required Back' was recorded on a work report.

When these results were reported to senior management they were treated with disbelief, in line with Sagan's remarks. Several efforts were made to change the existing blame culture, but the contract bidding system remained in place.

Sagan goes on to examine the US nuclear weapons programme. The information is necessarily old, since much of it was originally classified (and Sagan was able to use the Freedom of Information Act to have documents released, using a degree of openness unknown in many other countries including the UK). But it is extremely detailed, since numerous reports were compiled and investigations held, such that similar information would not be available in much of the private sector. Though the examples are old, moreover, they reveal processes that are likely to continue to exist.

Sagan's initial expectation was that the tight discipline of a military organization, together with the lethal dangers of nuclear weapons, would lead to the confirmation of HIROT. As he says, this is a very tough test for NAT. But he finds several results that cannot be understood without reference to NAT. In 1962, the USA discovered the construction of Soviet nuclear missiles in Cuba and during the ensuing crisis the defence forces were put on high alert. At the time, it was standard procedure to keep aircraft continuously in the sky, so that they could respond to any surprise nuclear attack. A very large number of missions was flown without mishap, as HIROT would expect. Yet there were several false reports of missile launches that were potentially dangerous. Importantly the system was less tightly coupled than it was designed to be, with several cases where bombers should have been launched according to the information received but where other choices were made.

A particularly telling case occurred six years later. A B-52 bomber carrying nuclear weapons crashed near the Thule base in Greenland. (Among other things, the event influenced an election in Denmark, which administers Greenland, in that the presence of nuclear weapons on its territory had been kept secret.) The plane was flying a sortie around the base. The purpose was to add redundancy to the warning system, for if communication from the base disappeared there would be no way of distinguishing a Soviet attack from a systems failure; the B-52 sorties were designed to provide independent information. However, if the plane had crashed into the Thule base instead of near it, the crash would have set off false reports of an attack. The system added complexity and opacity as well as redundancy. A telling fact is that civilians in Washington did not know of the existence of the sorties: had they been faced with indications of an attack, they lacked the knowledge to consider the alternative possibility, of a B-52 crash.

The Politics and Economics of Safety and Risk Assessments

HROT was developed on US aircraft carriers and there is little or no reference to financial or budgetary constraints. In many other cases, however, costs and budgets loom large. The pressures in the *Torrey Canyon* case were noted above.

- Consider one example analysed from a basically HROT position. Weick (1993) studied an accident at Tenerife airport in 1977 in which two planes collided, killing 583 people. (This remains the world's biggest single airline disaster.) The planes had been diverted there and were much delayed. The crew of one of them was nearing the legal limits on its flying time, and was thus under pressure to leave as soon as possible. This economic pressure may have contributed to its decision to take off without proper clearance. Note also that very similar issues occurred in a collision of two planes on the runway at Milan's Linate airport in 2000 (BBC, *Horizon* programme, broadcast 10 August 2003). Here, standard procedures were changed without explanation, and the pilots became confused. Our point is that to explain such cases one certainly needs to understand how people make mistakes and misread information, but in addition what are the conditions allowing these errors to occur? Why were lessons from Tenerife not learned, why was ground radar bought but not installed at Linate (were financial constraints involved) and why was a new system designed in preference to using an existing Norwegian one (were organizational politics to blame)?
- In the case of the Union Carbide's Bhopal plant (see Box 7.5), Perrow (1999: 355–60) stresses that the disaster was not a case of a normal (i.e. systems) accident, for such an accident entails the unexpected interaction of small failures. Bhopal was much simpler. The plant had been starved of investment, maintenance workers were laid off, and safety systems were neglected. Alarms did not go off and evacuation procedures were not followed. Finally, whereas in other cases good luck reduced the number of deaths (e.g. the wind direction at Chernobyl took radioactive material away from the city of Kiev), here the effects were about as devastating as they could be. Perrow also notes the tendency of some observers to invoke Indian 'culture' as a convenient explanation, an explanation that appeared very hollow when a similar accident occurred only months later at Union Carbide's US plant. Rather than explain the accident in terms of culture,

Box 7.5 The Bhopal disaster

The best account of Bhopal, which is both analytical and moving, is by the journalists Lapierre and Moro (2003). The key facts are that on the night of 2–3 December 1984 there was an explosion in the Union Carbide chemicals plant. The explosion released toxic gases which killed immediately about 8,000 people, with the final death toll being in the region of between 16,000 and 30,000. The explosion was caused by a failure to manage the production of methyl isocyanate (MIC), a lethal and highly volatile substance that was intended for use in the pesticide industry. The reasons for the explosion lay in a catalogue of failures of safety procedures, which meant that on the fatal night forty tons of MIC were in storage tanks and all three safety systems were non-operational. Union Carbide no longer exists as an entity, many of its sites having been taken over by other firms.

[W]e might look at plain old free-market capitalism that allowed the Indian plant to be starved and run down... [A]n economic system that runs... risks for the sake of private profits, or a government system that runs them for the sake of national prestige, patronage or personal power, is the more important focus and culprit.... [T]he issue is not risk but power. (Perrow 1999: 360)

- Studying cases of injury, Nichols (1975) found that underlying them was pressure to keep up production. He reported that industry was estimated to spend 0.05 per cent of its research and development budget on safety. Is there any evidence that this proportion has increased? Much recent research has tried to measure 'human resource outcomes' of forms of work organization, but it is notable that safety performance is often excluded, for example in *The HR Scorecard* (Becker et al. 2001).

What, then, can be done to improve safety? Perrow begins by offering a critique of the apparently scientific approach of risk assessment (Perrow 1999: 306–24). This approach is essentially a cost–benefit analysis. It estimates the likelihood of an event, puts a monetary value on it, and then compares this with other events and selects the option which will maximize net benefits.

This thinking was exemplified by the widely discussed (and now dated) the company can reasonably claim to have learned from the case) Ford Pinto case. It was found that the fuel tank of the car in question was prone to explode; the company placed a dollar value on the number of liability suits it was likely to face and concluded that this cost was lower than the cost of recalling the cars and correcting the problem. In this

particular case the issues are obvious: the callous equation of purely private monetary costs (of fixing the cars) with the emotional and other costs of death and injury borne by other people.

But there are broader issues from less clear-cut cases. Given that decisions have to be made and resources are finite, why not try to make decisions on a rational basis? Three points stand out from the answer given by Perrow.

- Risk analysis treats any death or injury as equivalent, whereas the collective costs of major disasters are shared and are greater than those of a similar number of individual events. How does one compare the 144 deaths in the small village of Aberfan with the same number in unconnected road accidents?
- Some risks involve choice and a sense of being able to exert a degree of control, for example driving a car. Numbers of injuries cannot be compared with numbers where the subjects lack any such control over their own futures.
- Risks are unevenly distributed across society, so that the poor and disadvantaged are subjected to relatively high technological risks. Moreover, the people judging what is an acceptable risk are not the same as those exposed to them. In some cases, it is also possible to externalize the risk (see Box 7.6).

The difficulty is that making decisions in this context involves very different kinds of rationality:

[E]conomic or absolute rationality, which requires narrow, quantitative and precise goals; bounded rationality, which emphasizes the limits on our thinking capacities and our inability to often achieve or even seek absolute rationality; and social and cultural rationality, which emphasizes diversity and social bonding. (Perrow 1999: 323)

Box 7.6 Externalizing risk

A study of the US petrochemicals industry found that widespread use of subcontract firms meant that the core firms could externalize the risk of accidents to subcontractors and their workers (Kochan et al. 1994; Rebitzer 1995). Training of subcontract workers was poor and was not seen as a responsibility of the core firms. Workers were also asked about their accident history. For those directly hired by the core firms, safety training cut accidents but this was not the case for contract workers, suggesting that these workers were relatively underprepared. This study did not, however, look directly at reporting, though it is consistent with Collinson's evidence (Box 7.4). Reported accident rates may not be accurate measures of the dangers of an occupation.

Cost-benefit analysis deals with the first. Risk assessment begins to move to the second. But what is also needed is the third. Social rationality recognizes the importance of *threat*: perceived risk shaped by a lack of control, inequitable distribution of hazards, high catastrophic potential, and a sense that technological fixes are not the solution (Perrow 1999: 328). Based on his analysis of the potential for catastrophe and the costs of alternatives, Perrow concludes that some technologies (primarily nuclear ones) should be abandoned, others (e.g. DNA research and experiments) restricted, and yet others subjected to more specific and narrow improvements.

Two observations are needed at this point. First, there is a danger of elevating social rationality above the others and implying that the kinds of false fears discussed at the start of this chapter have a reasonable basis. Yet it may be that some elements of dread are based on incorrect assumptions. It is true that 'expert', absolute, rationality needs to be questioned, for it too makes assumptions, and indeed has to do so since it is trying to estimate future probabilities of untried technologies, which are inherently unknowable. But the point is to encourage critical engagement between different ways of thinking. Indeed, the same people may use the different rationalities at different times. Perrow's analysis suggests that 'experts' engage in social rationality when they develop shared 'scientific' models. As we have already argued, even scientific judgements involve social rationality, and this will be especially the case where large profits are involved (e.g. biotechnology). A debate between different rationalities, and using the rationalities to throw different light on a given problem, is called for, not an argument that one rationality is superior to another.

Second, Perrow offers specific recommendations. These should not be seen as emerging directly from his social science. The science shows what the issues are and that safety concerns, for example, to do with dams, chemicals plants, and nuclear power are importantly different. What conclusions are drawn is a political question. A proponent of nuclear power might read the analysis as leading *inexorably* to the conclusion that nuclear systems should be abandoned, and hence dismiss the analysis as special pleading. Yet the underlying analysis remains central to how such a person might improve his thinking about risk and safety. The conclusion cannot follow from the analysis, though the analysis certainly shows that there are issues to be addressed.

The nature of the political argument about technologies needs clarifying. An early view said that technologies have necessary social effects. It was rapidly criticized for neglecting the choices that can be made, for example the very different social arrangements that grew up in different countries around assembly-line technologies. Yet choice is not limitless,

and technologies are not simply neutral. Winner (1999) has identified two ways in which technologies are political. The first is where a given technology is chosen with a social aim in view, a reasonably well-attested case being forms of automation intended to reduce the power of skilled workers. The second theme is that some technologies are inherently political, in that they entail a particular form of social organization. The strongest form of this argument is that something like a nuclear power plant has patterns of authority built into it. A weaker argument is that there are strong tendencies towards one set of social relations and not another. Nuclear power in our view has strong implications for social organization, and that fact constrains the sorts of politics that can be pursued around it.

As this chapter has proceeded, we have moved from approaches based on 'safe' systems to a more critical account that locates injuries in the pressures of the production process. The most developed 'political economy' of injuries is provided by Nichols (1997), whose analysis may be introduced via his criticisms of Perrow and of Turner. In his view, the systems approach has the following limitations:

- It neglects normal, routine injuries in favour of large disasters.
- In doing so, the focus on particular systems may be relevant to their designers, but the analysis cannot address the major fluctuations of injury rates across time and space.
- It treats systems as equivalent, for example nuclear plants and universities, and it says nothing about one all-embracing system, that of capitalism (pp. 4-5, 11, 90-1).

In relation to our present purposes, the first two points are important but may be left to one side. It is true that small, routine events cause more injuries than do dramatic disasters. But we are interested in disasters for the clear light that they throw on organizational functioning, and not because they are to be equated with injuries. On the third point, Perrow's initial work did tend to abstract 'systems' from the social reality in which they are embedded, though his analysis of Bhopal would be consistent with the approach advocated by Nichols.

That approach treats injuries as one outcome of a work organization driven by profit. Workers are exposed to unsafe conditions when the drive for profit is particularly intense, when the work process is hazardous (mines or chemicals works) or prone to contradictory demands, and when countervailing forces are weak. A good example, which goes beyond the categories of normal accidents, is an accident on the Midland Railway in England at Aisgill in 1913 (Howell 1991). A train hit the rear of

another after running through a signal at danger. The case became a cause célèbre when the driver was prosecuted and imprisoned, being released after widespread protests. Some aspects are consistent with NAT: railways are complex and closely coupled systems. But the degree to which the Midland Railway had these properties was worsened by its political and economic behaviour. The first train stopped because of the poor quality of coal it was using, and the driver of the second train was not watching for the signal because poor maintenance required him to check the lubrication of his engine so that he could not watch for the signal. These features in turn reflected financial pressures within the company to cut costs. In short, the system had the features that it did because of its political as well as its technological characteristics.

A neat way of capturing this, applicable to many issues other than accidents, is the question posed by W. G. Baldamus, 'what determines the determinants?' (quoted in Nichols 1997: 90). That is, specific causes can be located in a wider structure of influences.

A key point is that safety does not always suffer where profits are high. There is no direct trade-off. Nichols (1997: 104) gives several examples when the relationship may be positive:

- There is the threat of a major explosion, so that safety is taken seriously.
- Injuries are so common that production could be disrupted.
- Injuries are a worker-organization issue.
- An injury record leads to recruitment difficulties or customer discontent or the threat of state regulation.

Other examples could readily be given. A political economy view does not argue that safety is necessarily neglected, but sees the tension between safety and profits as one shaping how an organization will perform.

The tension between safety and profit is an example of the contradictions that run through this book: they reflect opposing *principles* that have to be managed, and in particular circumstances pressure for profit can be so intense that accidents are very likely. But how the principles work through in practice will reflect many other factors. It is not, moreover, the case that attention to safety is a cost necessarily to be minimized: a firm needs to attract and retain workers, and safety is important to business reputation and ultimately to the viability of the operation. The tension between safety and profit is a matter of degree, and the relationship will be different in different organizations. For example, firms pursuing a business strategy of innovation and 'high performance' would find their claims undermined if they took a callous view of safety. A considered

approach to safety would reasonably be expected to be part of the strategy. Nonetheless, the contradiction has not disappeared.

In the UK an improved injury rate in the 1970s was reversed in the early 1980s, reflecting changes in numbers of hours worked but also a weaker ability of trade unions to challenge managerial control (Nichols 1997). This last influence illustrates the third bullet point above, and also the fourth (the threat of powerful state regulation was reduced at this time).

Numerous case studies have revealed the operation of the politics of safety. Grunberg (1986) showed that in a UK plant then owned by Chrysler a weakening of labour's position was associated with more work effort, while in a French plant a stronger position of labour at plant and national level reduced work intensity. As might be expected, direct links with injury rates were less apparent, for injuries reflect many contingencies, but there were trends in the expected direction. Novek et al. (1990) report similar evidence from the Canadian meat-packing industry, showing that mechanization of tasks and efforts to intensify labour led to reduced levels of safety. By contrast, it appears from Hall's (1993) analysis of a mine, also in Canada, that a new work process increased the impact of injuries on the whole production process and therefore encouraged a new and more systematic approach to safety. Wokutch and VanSandt (2000) argue in relation to Japan that safety has a high priority because the just-in-time lean production system cannot afford inefficiencies, and injuries are an important source of inefficiency. They argue that injury records in large firms are better than in similar US ones, a marked change from the past when injury rates were up to five times the US level.

Alongside such trends, there is evidence that in some countries the link between safety and profit remains inverse. Nichols reports his work on coal mines in Turkey, showing that productivity and injuries tend to move together because work is labour intensive. Sass (2000) notes that in Taiwan the injury rate is, despite health and safety protections, between five and ten times that in Japan or the USA. He also highlights a group of about 300,000 migrants (out of a work force of about 9 million) who work in the informal sector and lack legal protections.

We have come a long way from the disciplined safety systems of US aircraft carriers. The politics of safety continue to be fought out around the globe.

Conclusions

Accidents are a good illustration of the risk (or, better, the uncertainty or ambiguity) society. Processes for the individualization of risk include the

putting of responsibility onto subcontractors and pressures on operators to disguise accidents. Wherever there is a blame culture there are likely to be such processes. Whether or not these processes are distinctively new, and hence whether they illustrate a risk society in the sense identified by Beck, is a different question. In our view, it is probably unanswerable: how do we know whether uncertainty has increased compared to some time in the past?

Several points are worth bearing in mind when thinking about the risk society. First, comparison of the present with a supposed stable or 'Fordist' past is unhelpful, since such an era was a relatively short one and was less coherent than might appear. The present may look uniquely risky only by comparison with a misleading picture of the past. Second, there is some evidence that long-term stability exists. Studies in the US textiles industry and the UK railways show remarkable stability of organizational processes over a period of approximately 100 years (Edwards and Whitston 1994). These sectors are, however, ones of limited technical and organizational change, so that when we look across society as a whole there may be more change, as the risk society thesis would suggest.

The key is to understand the organizational dynamics at work. What are the processes generating uncertainty and how are they distributed across society? To the extent, for example, that there is more subcontract working then we can expect an externalization of uncertainty. (See the discussion of status and contract in Chapter 6.) Yet there may also be limits on the process, for example tighter monitoring of contract compliance and legal regulation on the wearing of safety equipment. The 'risk society' is an element of social organization or one tendency, not a description of concrete social reality.

Lessons from the cold war, outlined above, have wider application. The particular context of two superpowers each capable of launching a devastating attack on the other has disappeared. But the growing number of nations with nuclear capability and the rise of regional tensions, as between India and Pakistan, suggests that the threat has not gone away. The issues of complex and tightly coupled systems, in particular problems of handling ambiguous data in times of tension, remain present. One of the key arguments of HROT is that high reliability organizations are capable of learning through analysis of experience and through structured training exercises. But new organizations will lack this experience: any learning in one organization will not generalize across all relevant organizations. Moreover, high reliability cases are not typical of all organizations, so that systematic attention to procedures and a strong safety culture are likely to be conspicuous by their absence.

In short, it is not the case that technological change has wrought unalloyed progress, but nor is it true that its results have been generally disadvantageous. Unexpected and contradictory implications often ensue. The broader implication in thinking about organizations is: some organizations may stand as exemplars of 'good practice', but learning from them is never easy. There are many ways in which people in other organizations can find out that they are different so that the lessons are inapplicable. The point to consider is which aspects of practice are transferable, and which are not.

It is also useful to distinguish between the economic and the organizational context. The economic context establishes the degree of pressure to meet financial targets and may help some people to sustain arguments that certain safety procedures are too expensive. Yet the context does not wholly determine outcomes, and the arguments are open to challenge. Such questioning can be immediate or longer-term. In the immediate situation, it may be that cost pressures require completion of a project in a limited time, and that in practice this leads to the cutting of corners. But there may be other ways of organizing the work which make it safer and no less efficient. In the longer term the budget pressures can be challenged by finding genuinely better ways of organizing a task or by relieving the intensity of the pressures.

The organizational context covers the extent to which there is a 'culture of secrecy' or a tendency for work groups to take unsafe working practices for granted. These matters are pursued in Chapter 8. For now we can say that the salience of the culture will vary across organizations. The more a process is complex and tightly coupled, the more important it is that the people involved with it think about how they manage it and whether the purely social processes involved (e.g. are orders from above unquestioned, can people challenge authority?) heighten the dangers involved.

We have seen in this chapter, then, that safety is not a purely technical issue. The social organization of technical systems is a critical dimension. In thinking about why disasters happen, and what an individual in an organization can do, the contestability of views is central. Economic, bounded, and social rationality can be used as different perspectives to understand how we respond to uncertainty and risk. We can ask what kind of uncertainties exist in given situations and how they are managed. To what extent is there a safety culture, or how far do people feel that they cannot question existing assumptions? What can be done to change ways of thinking? Using the ideas in this chapter and in Chapter 8, it is possible to understand the causes of failures in organizations and the warning signs, as well as the ways in which they may be heeded.

Is Decision-making a Rational Process?

Chapter 7 focused on the particularly sharp issue of disasters and injuries and identified the social processes involved in them. We now consider decision-making of a less disastrous kind. Following the logic of Chapter 7, we begin with processes of social group behaviour, which are often analysed through a social psychological perspective that tends to see organizations as having shared purposes or as having 'politics' of a relatively minor kind. We then place this perspective in a wider critical analysis of the politics of organizations. As with Chapter 7, it is useful to take as examples cases where things go wrong. As we will see, however, many of the same processes underlie 'success' as well as 'failure'. We need to grasp how failure occurs but then aim to identify what is generic in the politics of organizations and what is particular to failure. With this approach in place, we can then ask how organizations might learn from experiences of success and failure and what organizational learning might mean.

A useful way to organize thinking here is the classic image of 'a garbage can model of organizational choice' (Cohen et al. 1972). What the authors mean by a garbage can is a choice opportunity into which problems and solutions are dumped. It exists because individuals and groups have their own projects and views on the ways in which an organization should develop. They are likely to have pet solutions that they put into practice when a new issue arises. Decision-making is often about the interplay of pre-existing solutions and recipes, which are looking for problems to which they may be an answer.

Social Group Processes

Normal Accident Theory discussed in Chapter 7 focuses on the properties of systems that produce 'accidents' and, if the technology is sufficiently dangerous, disasters. Yet how do we make links between organizational systems, which are relatively large-scale and which exist independently of people, and smaller-scale social group processes?

A study by Edmondson (1996) attempts an answer. It looks at errors in drug prescriptions, noting another study that found a huge range of