Was Three Mile Island a 'Normal Accident'?

Andrew Hopkins*

Perrow's normal accident theory suggests that some major accidents are inevitable for technological reasons. An alternative approach explains major accidents as resulting from management failures, particularly in relation to the communication of information. This latter theory has been shown to be applicable to a wide variety of disasters. By contrast, Perrow's theory seems to be applicable to relatively few accidents, the exemplar case being the Three Mile Island nuclear power station accident in the U.S. in 1979. This article re-examines Three Mile Island. It shows that this was *not* a normal accident in Perrow's sense and is readily explicable in terms of management failures. The article also notes that Perrow's theory is motivated by a desire to shift blame away from front line operators and that the alternative approach does this equally well.

Introduction

In 1984, Charles Perrow published an analysis of the Three Mile Island nuclear power station accident in the US. His argument was that the Three Mile Island incident was a distinctive kind of accident, a 'normal' accident, and his theory has become known as Normal Accident Theory.

Perrow's book became an almost instant classic and has had a profound influence on social science research on major accidents. One indicator of its status is the fact that it has just been re-issued (1999), fifteen years after its original publication. Ironic testimony to the influence of normal accident theory is the emergence of a contrasting school of thought known as High Reliability Theory (e.g., La Porte and Consolini, 1991). A lively debate between these contrasting perspectives has ensued (Sagan, 1993; La Porte and Rochlin, 1994; Perrow, 1994).

This article is not intended as a contribution to the debate between these two perspectives. It critiques Normal Accident Theory from a different vantage point, namely, Turner's (1978) work on man-made disasters, to be outlined below. I have dealt elsewhere with some of the limitations of Normal Accident Theory (Hopkins, 1999a); the present article takes the critique a step further. But before we can begin the critique we need a clear statement of Perrow's argument.

Perrow on Normal Accidents

Perrow outlines what he means by a 'normal' accident in the following passage.

> It is termed normal because it is inherent in the characteristics of tightly coupled, complex systems and cannot be avoided. A tightly coupled system is highly interdependent; each part is linked to many other parts, so that a failure of one can rapidly affect the status of others. A malfunctioning part cannot be easily isolated either, because there is insufficient time to close it off or because its failure affects too many other parts, even if the failure does not happen rapidly.
>
> A normal accident occurs in a complex and tightly coupled system when there are unanticipated multiple failures in the equipment, design, or operator actions ... The crucial point about a normal accident is that unexpected, multiple failures occur. As a result, for some critical period of time, the nature of the accident is incomprehensible to those seeking to control it.
>
> In addition to being unforeseeable, incomprehensible and not amenable to knowledgeable intervention, the normal accident cannot be prevented because it is not possible to create faultless systems.
>
> ... [T]he [Three Mile Island] accident was unexpected, incomprehensible, uncontrollable and unavoidable; such accidents had occurred before in nuclear plants, and would occur again, *regardless of how well they were run* (1982: 174, 176, emphasis added).

This passage has been quoted at length so that there can be no doubt about Perrow's argument: certain technologies make major accidents inevitable, no matter how well managed an operation may be. This is an unashamedly technological determinist argument.

Turner: Warnings Plus Sloppy Management

Standing in sharp contrast to Perrow is Turner's analysis, first published in his book *Man-Made*

*Department of Sociology, Faculty of Arts, Australian National University, Canberra, 0200, Australia, E-mail: andrew.hopkins@anu.edu.au.

Disasters, in 1978, and augmented in a significant article in 1994. According to Turner, prior to large-scale accidents there are nearly always warning signs that are missed, overlooked or ignored and that, if acted on, would have averted the accident. The central question then is sociological rather than technological: 'what stops people acquiring and using appropriate advance warning information, so that large-scale accidents and disasters can be prevented?' (Turner and Pidgeon, 1997: 162). His answer, in its simplest form, is sloppy management (Turner, 1994). Sloppy management fails to put in place adequate information gathering systems. Moreover, sloppy management allows various processes to operate that may nullify warnings: the normalisation of deviance (Vaughan, 1996), group think (Janis, 1982), cultures of denial (Hopkins, 1999b) and so on (Turner, 1994: 218). Good management would have systems designed to override these tendencies and to highlight and respond to warning signs; the failure to establish such systems is a management failure.

Turner showed in his 1978 book that his analysis was apposite for a large number of disasters. Analyses of various disasters since that time confirm the applicability of his approach (Appleton, 1994; Hopkins, 1999b; Vaughan, 1996).

Reconciling the two perspectives?

Here, then, are two very different explanations of large-scale accidents, apparently in competition with each other. How shall we make sense of this situation? One possibility is to assume that they apply in different circumstances and are therefore not in competition. After all, not all systems are characterised by tight coupling and complexity. And in complex and tightly coupled systems, accidents may occur for different reasons. Perrow himself suggests that the scope of his explanation is limited in this way. The gas leak from a chemical plant that killed thousands at Bhopal in India, the fiery destruction of the space shuttle, *Challenger*, the Soviet nuclear reactor accident at Chernobyl from which people are still dying, the *Exxon Valdez* tanker oil spill in Alaska, an undoubted environmental disaster if not a human one – none of these is a normal or system accident, according to Perrow (1994: 218). They are no more than 'component failure accidents', which cannot be analysed in system terms. 'They are alarmingly banal examples of organisational elites not trying very hard', he says (Perrow, 1994: 218). Even Seveso, the chemical accident in Italy, which precipitated a new genre of regulation in Europe for dealing with major hazards, was merely a component failure accident in Perrow's (1984: 295) judgement. In other words, many of the most publicised disasters of our time are not explicable in terms of Normal Accident Theory. The theory is thus not as useful as might at first have been thought. But this does not invalidate it, since, presumably, there is still a class of major accidents caused inevitably by tight coupling and complexity. The paradigm case of such an accident is, presumably, the Three Mile Island incident.

This is the resolution that Turner and his associates seem to have accepted; namely that there are some major accidents to which Turner's theory does not apply and that are only explicable as 'normal accidents'. Thus Turner (1994: 218) writes

> As our technological systems become increasingly extensive and complex, the possibility grows of some accidents arising from the properties of the system as a whole, often as a result of unforeseen interactions which involve several organisations. To understand some disasters, therefore, it is necessary to recognise that some highly complex systems generate 'normal accidents'.

Pidgeon seems also to concede some validity to Perrow's argument, although he notes that 'the concepts of complexity and coupling have turned out to be difficult to use analytically' and that the account is overly deterministic (Turner and Pidgeon, 1997: 179).

But is this concession really necessary? Are there really cases of major accidents to which Turner's theory of ignored warnings does not apply – accidents that were inevitable for purely technological reasons?

In seeking to answer this question it is no use taking as yet unanalysed cases, even in high-tech industry, and showing that Turner's theory applies, since this does not dent the claim that at least some major accidents are normal in Perrow's sense. The most effective strategy is to focus on the exemplar case, Three Mile Island. Most commentators writing about major accidents have implicitly accepted Perrow's own account of this accident. What I propose to do here is to re-analyse this case and show that there was nothing technologically inevitable about the incident. I want to show that it conformed to a remarkable extent to the Turner model of ignored warnings and sloppy management.

The Issue of Multiple Failures

Three Mile Island was a water-cooled nuclear reactor. On 28 March, 1979, a chain of events began that resulted in the escape of the cooling water (a 'loss of coolant accident' in the jargon of the industry). The reactor came very close to a meltdown, but in the event no nuclear radiation escaped, and no appreciable consequence to the health of nearby residents occurred. However, the reactor was damaged beyond repair, clean up took more than a decade and cost $1 billion and no reactors have since been built in the US (Rees, 1994: 11).

There were no fewer than four quite separate malfunctions in the accident sequence, followed by a crucial error by control room operators (Perrow, 1984: 18). The issue of 'operator error' is of course controversial, and I shall return to it later. But the point to be stressed here is that these malfunctions did not constitute a series of failures following consequentially from a single original cause. Consequential failures are perfectly comprehensible and quite predictable. For example, poor roof design may lead to such an accumulation of snow that the load limits for the building are exceeded, causing it to collapse. In this situation the failure of a well-designed support structure is predictable given the initial roof design failure (hypothetical example inspired by Pidgeon, 1997: 2). The malfunctions at Three Mile Island were not consequential on some original failure in this way: they were independent of each other. No-one could have foreseen the way in which these four discrete failures interacted to produce the loss of coolant accident. The precise configuration of failures was unprecedented and unpredictable.

So far so good. But Perrow concludes from this that the accident was inevitable and not preventable. This is a logical error. It is not necessary to predict the entire accident sequence to avoid such an accident. The point is that had any one of the malfunctions not occurred, the loss of coolant accident would not have occurred. The accident sequence was highly susceptible to interruption, unlike accident sequences that involve consequential failures.

Perhaps the best way to make this point is to draw on Reason's (1997: 9) concept of an 'accident trajectory'. He argues that major technological hazards are managed by constructing a series of 'defences in depth'. Any one of these defences, if it works, will terminate the trajectory of a potential accident. However, defences often have 'holes' in them; Reason describes this as the Swiss cheese metaphor. If a potential accident trajectory passes through one of the holes in the first defence, it will encounter the second defence. For an accident to occur, the holes in all defences must line-up, that is, the defences in depth must all fail simultaneously.

According to the Swiss cheese model, if only one of the defences had been effective the accident would not have occurred. It makes sense, therefore, to ask of every defence that failed why it failed and whether its failure was predictable. If even one of these failures was predictable, then we can reasonably conclude that the accident itself was preventable. In short, it is not necessary to be able to predict the precise trajectory of an accident in order to be able to prevent it (for a related analysis, see Toft and Reynolds, 1994).

We could go further and argue that the simultaneous failure of all defences in depth is not evidence that accidents are inevitable, but rather that the system of defences is inadequate. One is reminded of Lady Bracknell's comment in *The Importance of Being Earnest*: 'To lose one parent may be regarded as a misfortune, to lose *both* seems like carelessness' (Wilde, 1980: 30).

The Failures

In order to be able to consider whether the malfunctions and failures were predictable, and therefore preventable, we will need to describe the process in more detail. Anyone who seeks to describe the Three Mile Island accident confronts the problem of how much detail needs to be conveyed to the reader and just how to convey those details judged to be necessary. Perrow (1984: 15), himself, struggled with this and in the end resolved it as follows:

> What I wish to convey is the interconnectedness of the system and the occasions for baffling interactions. This will be the most demanding technological account in the book, but even a general sense of the complexity will suffice if one wishes to merely follow the drama rather than the technical evolution of the accident.

For present purposes it is necessary that readers understand at least some of the detail, but I shall strive to present it as parsimoniously as possible.

Nuclear reactors generate enormous heat and at Three Mile Island this was removed from the reactor core by a primary, internal circuit of coolant. The heat from the coolant in this internal circuit was transferred via a heat exchanger to water in a secondary, external circuit, turning it into steam. This, in turn, drove steam turbines, located in the external circuit, which generated electricity. Two of the four malfunctions occurred in the external circuit and two in the internal circuit. I deal with these in turn.

1) Water used in the steam turbines must be pure, and the external circuit therefore contained a water purifier that needed regular maintenance. It was this maintenance work that triggered the accident sequence. A leaky valve set off a chain of events that caused the main pumps in the external circuit to close down. This loss of flow in the external circuit

meant that heat would no longer be removed from the internal circuit.

2) To overcome this danger, emergency pumps were supposed to come into action to maintain the flow of water in the external circuit. But they had been blocked off two days earlier for maintenance and, by mistake, the blocks had not been removed, rendering the pumps inoperative. The external circuit therefore ceased removing heat from the internal circuit.

   The failure of the external circuit would not have been critical if automatic safety devices in the internal circuit had functioned as intended. But they did not.

3) The heat was no longer being removed from the internal cooling circuit. As a result, the reactor core began to overheat and the reactor closed down, as it was designed to do. Predictably, however, it continued to emit 'decay heat'. Pressure in the internal circuit built up to the point where a relief valve in the circuit opened, as it was designed to do. Coolant escaped and the pressure returned to normal. But the value malfunctioned, failing to close properly, and coolant continued to escape.

4) Because of a design flaw, a light on the control panel indicated that the valve had closed even though it was open. Thus, operators did not know that they were experiencing a loss of coolant accident.

5) Because of the continuing loss of coolant, pressure in the internal circuit dropped to a level that meant that it would be unable to continue conducting heat away from the core. Unless the circuit could be re-pressurised, the core would melt and large amounts of radioactive material might be released. To avoid this danger, another automatic safety device kicked in: high pressure injection pumps came on, forcing additional water into the internal circuit to make up for the loss of coolant. Had the pumps been allowed to continue operating, the accident could still have been avoided, but operators, seeing the pressure rising and knowing that this could damage the system if it continued, manually throttled these pumps back. They were still unaware that a major loss of coolant was occurring and that coolant was not reaching the core at sufficiently high pressure. It was two hours and 20 minutes before operators discovered their error. They immediately reactivated the high pressure injection system and flooded the core with cold water. By this time major and irreversible damage had been done to the core, but by good fortune the situation was brought under control before there had been any release of radioactive material.

Perrow's claim is that this accident sequence was so complex and the information available to operators so flawed that there was no way they could have been expected to understand what was going on and react in an effective manner. Indeed, he provides rather more detail than has been presented here as to why operators were confused:

> 110 alarms were sounding; key indicators were inaccessible; repair-order tags covered the warning lights of nearby controls; the data printout on the computer was running behind (eventually by an hour an a half); key indicators malfunctioned; the room was filling with experts; several pieces of equipment were out of service or suddenly inoperative (1982: 180).

We can agree with Perrow that, given the situation the operators found themselves in, there was no way *they* could have avoided the accident. But it does not follow that the accident was unavoidable. We can legitimately ask: why did operators find themselves in this position? And more pointedly: could management have prevented this accident? This latter question is not seriously raised by Perrow. Had he done so he might have arrived at rather different conclusions about the avoidability of the Three Mile Island accident.

The Warnings

Although the particular sequence of events at Three Mile Island was unprecedented, *sections* of the event sequence had occurred previously. There had, in short, been warnings. Had these warnings been properly attended to, the Three Mile Island accident would not have occurred.

Consider first the initial failure of the pumps in the external circuit, triggered by maintenance work on the water purifier. During this work, a flow of water through a leaky valve resulted in a spurious electrical signal that automatically closed certain other valves. This interrupted the flow in the external circuit to the main pumps, which then shut down as they were designed to do in these circumstances.

The Report of the Office of the Chief Counsel assisting the Presidential Commission notes that 'virtually every detail of this ... sequence had been duplicated in an incident 17 months earlier on 19 October, 1977' (Gorinson, 1979: 157). Company investigators at the time recognised the potential for this accident to escalate and wrote a memorandum that 'recommended nine steps that should be acted on to preclude a recurrence'. This memorandum was summarised as it passed up the company hierarchy and the full significance of the event was not appreciated by senior management who therefore failed to provide a satisfactory response. The Chief

Counsel notes that had the company 'given careful follow-up attention to the 19 October, 1977 incident, the ... [purifier] might not have malfunctioned on 28 March, 1979 and the accident sequence might never have had a chance to begin' (Gorinson, 1979: 158).

Although the accident sequence began with failures in the external circuit, the malfunctions and failures in the *internal* coolant circuit (the circuit through the reactor core) were far more critical, as these failures ultimately allowed the loss of coolant to occur and the reactor to heat up dangerously. Consideration of these failures follows.

The reactor was manufactured and supplied to the utility company by another firm, Babcock and Wilcox, which retained a contractual responsibility for its product. The Presidential Commission found that there had been 11 previous relief valve failures in Babcock and Wilcox reactors, nine of them being failures in the open position. Each of these nine failures amounted to a loss of coolant accident that, fortunately, had been prevented from escalating. One of these failures had occurred at the Three Mile Island reactor a year before the major accident. Babcock and Wilcox was aware of these failures and in some cases recommended remedial action, but it did not systematically notify all its utility customers. Chief Counsel for the Commission was critical of the company for not recommending 'additional training of operators to ensure that they were aware of a likelihood of a relief valve failure and knew how to identify quickly the resulting small-break loss of coolant accident' and how to respond effectively (Gorinson, 1979: 151). It was also critical of the fact that Babcock and Wilcox had never recommended a modification of the value position indicator to ensure that it worked reliably (Gorinson, 1979: 151).

But far more disturbing than these repeated failures of the relief valve is that almost the entire sequence of events that occurred at Three Mile Island had occurred at another Babcock and Wilcox reactor owned by another utility company, Davis-Besse, 18 months earlier (Gorinson, 1979: 130). This event sequence involved:

- a total loss of circulation in the external circuit;
- a relief valve that failed in the open position in the internal circuit;
- a failure of the operators to recognise that a loss of coolant accident was occurring;
- the premature termination by operators of the high pressure injection of water to compensate for loss of coolant.

This was precisely what occurred at Three Mile Island. According to the Chief Counsel there were only two significant differences between the two accidents. First, operators at Davis-Besse realised after 20 minutes that the relief valve had stuck open, whereas at Three Mile Island this was not understood for two hours and 20 minutes. Second, Davis-Besse was operating at the time at 9 percent power while Three Mile Island was operating at 97 percent.

One month later operators at Davis-Besse again made the mistake of throttling back the high pressure injection in circumstances that might have proved disastrous (Gorinson, 1979: 130). Two engineers at Babcock and Wilcox were concerned about these events and one wrote a memo to his senior management in which, after describing the events, he made the following observation:

> Since these are accidents which require the continuous operation of the high pressure injection system, I wonder what guidance, if any, we should be giving to our customers on when they can safely shut the system down following an accident? (Gorinson, 1979: 130)

There was no response to this memo, so three months later the second engineer wrote a more strongly worded memorandum to his superiors entitled 'operator interruption of high pressure injection', in which he said in part:

> I believe it fortunate that Davis-Besse was at extremely low power Had this event occurred in a reactor at full power. ... it is possible, perhaps probable, that core uncovery and possible fuel damage would have resulted (Gorinson, 1979: 132).

Here was a clear warning of reactor meltdown. But senior management at Babcock and Wilcox did not attend to this warning and nothing was done to retrain reactor operators or to change their work practices in regard to the termination of high pressure injection. Chief Counsel concluded as follows:

> (The) Davis-Besse (incident) revealed that operators had been provided with inadequate procedures for termination of high pressure injection following a Davis-Besse-type loss of coolant accident. At Babcock and Wilcox employees sought to get the company to correct that error. Yet through neglect and bureaucratic mistakes that information was never conveyed to Babcock and Wilcox customers ... (Gorinson, 1979: 130).

It is true that Babcock and Wilcox periodically sent a bulletin to its utility customers and that the Davis-Besse incident was reported in one of these bulletins. But it was reported only as a loss of coolant accident caused by a failure of the relief valve. No mention was made of the fact that the high pressure injection system, meant to protect against the consequences of loss of coolant, had been terminated prematurely by operators. In the words of the Chief Counsel, the bulletin 'did not emphasise the real importance of the event – namely that operator termination of high pressure injection ... could result in core uncovery and fuel damage (i.e. meltdown)' (Gorinson, 179: 149)

Apart from the Davis-Besse incident, there was one other detailed warning. Two years before the Three Mile Island accident an engineer at another utility had carried out an analysis of the circumstances that could lead to such an accident. He predicted that operators were likely to terminate the high pressure injection system prematurely and explained precisely why this could be expected. The analysis was forwarded to Babcock and Wilcox, who effectively did nothing about it. According to Chief Counsel, 'by the end of May 1978 (10 months before Three Mile Island) Babcock and Wilcox knew that concern about operator interruption of high pressure injection ... had been expressed from three sources (only two mentioned here) ... Still nothing happened' (1979: 155).

Why Were Warnings Ignored?

It is clear that there were plenty of warnings of what might happen at Three Mile Island. Almost every bit of the accident sequence had occurred previously and the Davis-Besse incident was virtually a dress rehearsal for what happened at Three Mile Island. Several people had explicitly and persistently warned of the dangers. Why were all these warnings to no avail? The answer is again provided by the Chief Counsel: none of the companies concerned had adequate organisational procedures for attending to past experience.

Take Babcock and Wilcox first. This company had no individual or department responsible for analysing accidents occurring to its reactors. Moreover, the company received only abbreviated summaries of accidents occurring in its reactors rather than the full text prepared by the utility concerned. Again, Babcock and Wilcox did not incorporate accidents from other nuclear plants or even its own plants into training exercises it ran for plant operators (Gorinson, 1979: 125).

Consider, next, the utility that operated Three Mile Island, Met Ed. The only person responsible for reviewing the experience of other plants was the training manager. He told the President's Commission that he spent about one tenth of his time reading reports of other accidents. Nine months before the accident, Met Ed's parent company, General Public Utilities, expressed concern 'that Met Ed's internal system was not digesting information received about experiences at other plants' (Gorinson, 1979: 123). But there was no effective response by Met Ed to this concern and General Public Utilities did not push the point. The president of General Public Utilities concluded later that 'to me ... one of the most significant learnings of the whole accident is the degree to which the inadequacies of that experience feedback loop ... significantly contributed to making us and the plant vulnerable to this accident' (Gorinson, 1979: 125).

Part of the reason for this complacency was what has since been described at the fossil fuel mentality of the power generating industry (see Rees, 1994: 15). The utilities had traditionally produced their power using the well-established technologies of coal and oil. This was tried, true and safe, and management felt no need to attend to technical details of plant operation. They could afford to remove themselves from technical matters and focus their attention on marketing, cost cutting and profit maximising. Despite changed circumstances, this mentality persisted until the Three Mile Island accident. Utility management simply had not come to terms with the fact that they were now generating power using a relatively untried technology involving hazards of an unprecedented nature that required a far more diligent approach to safety. According to Rees (1994), this is now well understood in the industry, in part because it is obvious that another accident of the type which occurred at Three Mile Island would mean the end of the nuclear power industry in the United States.

The Applicability of Turner's Analysis

It is now clear that the accident at Three Mile Island conforms beautifully to Turner's account. The exemplar case of a normal accident turns out to be just another case of sloppy management. It is a story of ignored warnings, inadequate communication, and failure by management to focus its attention on safety. This was not an accident rendered inevitable by technological complexity and tight coupling; it was not an accident so different from other major accidents as to require a wholly new explanatory paradigm, the paradigm of the normal accident. Turner's earlier noted concession that some accidents may not be explicable in terms of sloppy management and may need to be understood as normal accidents, in Perrow's sense, now appears to be premature. I noted earlier that the category of accidents to which Normal Accident Theory applies is severely limited. The implication of the present discussion is that it may well be impossible to identify any major accidents to which the theory applies.

Perrow's Motivation

All this raises the question of why Perrow chose to analyse the Three Mile Island accident in the way he did. He was certainly aware of all the evidence of warnings and sloppy management that has been presented above and refers to it at various points in his book (1984: 20, 24, 48, 57). As he notes at one point:

> Time and again (in the story of Three Mile Island) warnings are ignored, unnecessary risks taken, sloppy work done, deception and downright lying practised (1984: 10).

Perrow appears at one point to play down the significance of warnings. They are ignored, he says, because 'signals are simply viewed as background 'noise' until their meaning is disclosed by an accident' (Perrow, 1982: 175). It is only with hindsight that they appear as warnings.

This account has an initial plausibility, but it is ultimately unsatisfactory. The key to understanding why, is to note Perrow's use of the passive voice in the above quotation that implicitly draws attention away from the agent of the act. Suppose we focus on the agent and ask: who viewed the signals as background noise? Clearly not the various engineers who issued the warnings about the possibility that high pressure injection might be terminated prematurely. To these people, the events to which they were reacting were not noise – they were deeply troubling and dramatic indications of just how a major accident might occur. And the memoranda they wrote were quite explicit and incapable of being dismissed as noise by anyone who read them. The failure of more senior management to respond stemmed not from an inability to distinguish noise from relevant signals but from the inadequacy of organisational procedures for attending to past experience.

But it is not this alleged problem of 'noise', which accounts for Perrow's decision to pass over the evidence of ignored warnings, as he explains in following passage.

> … warnings occurred well before TMI. A bureaucratic tale worthy of Franz Kafka came out of the investigation of TMI and the warnings, *which we shall forego telling* so we can stick with the villains of the piece, according to most reports: the hapless operators (Perrow, 1984: 24, emphasis added)

Perrow's ultimate purpose was to challenge the finding of the President's Commission that the major cause of the accident was operator error. His principal argument against the thesis of operator error is that the situation confronting the operators at Three Mile Island was so complex and opaque that they could not possibly be expected to understand what was happening or what actions they should take to deal with the problem. Normal Accident Theory was elaborated in order to demonstrate, theoretically, why this was the case. In complex, tightly coupled technological systems, he theorised, 'for some critical time period the nature of the accident is incomprehensible to those seeking to control it' (Perrow, 1982: 174). The operators, therefore, are not to blame.

From the perspective of the late 1990s Perrow's strategy for combating the theory of operator error as the cause of major accidents seems dated and unnecessary. The dominant accident causation models of today systematically down-play the significance of operator error. They emphasise management failures as the root cause of accidents and, in particular, the root cause of any operator error that may have contributed to an accident (Hale, Baram and Hovden, 1998: 3). It is noteworthy, too, that Reason's (1990) influential model of latent and active failures focuses attention on latent failures, which are necessarily a management responsibility (Chapter 7). These models provide an obvious explanation of the Three Mile Island accident; one that does not blame the operators.

But at the time Perrow conceived his theory of normal accidents, in 1979, explanations of accidents in terms of management system failures were not well known. (For instance, Turner's 1978 account, which stressed communication failure as the cause of accidents, was not referenced in Perrow's 1984 book.) In these circumstances, Perrow's theory of normal accidents provided an innovative and effective way to shift the blame for accidents away from frontline operators. Today's theories of management system failure achieve exactly the same effect. Perrow's theory is therefore no longer necessary from this point of view. It must stand on whatever other merits it may possess.

At this juncture, two points of clarification are in order. First, some readers of an earlier draft of this article have viewed this section as an attack on Perrow's motivation. That is not my intention. In fact I applaud his purpose. The common tendency to blame frontline operators for accidents is, in my view, unhelpful and unfair. My point is simply that Perrow's purpose is better achieved in other ways.

Second, some readers have objected to this account of Normal Accident Theory on the grounds that the theory is broader than I have allowed and encompasses questions of management and management failure. This is an important issue that I have canvassed in detail elsewhere (Hopkins, 1999a). It is not appropriate to repeat that discussion here. My conclusion, however, is that attempts to interpret Normal Accident Theory as including questions of management failure end up undermining the theory rather than enriching it.

Conclusion

So was Three Mile Island a normal accident? The accident certainly occurred in a complex and tightly coupled system. But that did not make it inevitable and unavoidable. It was sloppy management and failure to attend to warnings that allowed this accident to occur. It was in no way dissimilar in this respect to the many accidents about which Turner theorised in his book, *Man-Made Disasters*. Three Mile Island was *not* a normal accident in the specialised sense in which Perrow uses this term. It is doubtful if any accident is. If accidents are normal, what makes them so is sloppy management, not complexity and tight coupling. And sloppy management of major hazards is what gives us the potential for disaster.

This is not just a theoretical debate. There are practical consequences for the way we go about accident prevention. Normal accident theory suggests a *technological* approach: reduce complexity and coupling. The alternative approach is to make *organisational* changes designed to improve flows of information, decision-making processes and so on. It is this latter approach that is likely to be more useful for accident prevention.

References

Appleton, B. (1994), 'Piper Alpha', in Kletz, T. (Ed.), *Lessons from Disaster: How Organisations Have No Memory and Accidents Recur*, Institute of Chemical Engineers, London, pp. 174–184.

Gorinson, S. et al. (1979), *Report of the Office of Chief Counsel on the Role of the Managing Utility and its Suppliers*, Washington.

Hale, A., Baram, M. and Hovden, J. (1998), 'Perspectives on Safety Management', in Hale, A. and Baram, M. (Eds), *Safety Management: The Challenge of Change*, Pergamon, Oxford, pp. 1–18.

Hopkins, A. (1999a), 'The Limits of Normal Accident Theory', *Safety Science*, Volume 32, Number 2, pp. 93–102.

Hopkins, A. (1999b), *Managing Major Hazards: The Lessons of the Moura Mine Disaster*, Allen and Unwin, Sydney.

Janis, I. (1982), *Groupthink: Psychological Studies of Policy Decisions and Fiascoes*, Houghton Mifflin, Boston.

La Porte, T. and Consolini, P. (1991), 'Working in Practice But Not in Theory: Theoretical Challenges of 'High-Reliability Organisations'', *Journal of Public Administration Research and Theory*, Volume 1, Number 1, pp. 19–47.

La Porte, T. and Rochlin, G. (1994), 'A Rejoinder to Perrow', *Journal of Contingencies and Crisis Management*, Volume 2, Number 4, pp. 221–227.

Perrow, C. (1982), 'The President's Commission and the Normal Accident', in Sils, D., Wolf, C. and Shelanski, V. (Eds), *Accident at Three Mile Island: The Human Dimensions*, Westview, Boulder, pp. 173–184.

Perrow, C. (1984), *Normal Accidents: Living With High-Risk Technologies*, Basic Books, New York.

Perrow, C. (1994), 'The Limits of Safety: The Enhancement of a Theory of Accidents', *Journal of Contingencies and Crisis Management*, Volume 4, Number 2, pp. 212–220.

Perrow, C. (1999), *Normal Accidents: Living with High-Risk Technologies*, Princeton University Press, Princeton.

Pidgeon, N. (1997), 'The Limits to Safety? Culture, Politics, Learning and Man-Made Disasters', *Journal of Contingencies and Crisis Management*, Volume 5, Number 1, pp. 1–14.

Reason, J. (1990), *Human Error*, Cambridge University Press, Cambridge.

Reason, J. (1997), *Managing the Risks of Organisational Accidents*, Ashgate, Aldershot.

Rees, J. (1994), *Hostages of Each Other: The Transformation of Nuclear Safety Since Three Mile Island*, University of Chicago Press, Chicago.

Sagan, S. (1993), *The Limits of Safety*, Princeton University Press, Princeton.

Toft, B. and Reynolds, B. (1994), *Learning from Disasters*, Butterworth Heinemann, Oxford.

Turner, B. (1978), *Man-Made Disasters*, Wykeham, London.

Turner, B. (1994), 'Causes of Disaster: Sloppy Management', *British Journal of Management*, Volume 5, Number 3, pp. 215–219.

Turner, B. and Pidgeon, N. (1997), *Man-Made Disasters*, Butterworth Heinemann, Oxford.

Vaughan, D. (1996), *The Challenger Launch Decision: Risky Technology, Culture and Deviance at NASA*, University of Chicago Press, London.

Wilde, O. (1980), *The Importance of Being Earnest*, (Edited by R. Jackson), Ernest Benn, London.