

Technological Accidents: Learning from Disaster

by

Robin COWAN

MERIT, University of Maastricht

Emmanuelle FAUCHART

CNAM, Laboratoire d'Econométrie

Dominique FORAY

IMRI, University of Paris - Dauphine

and

Philip GUNBY

Department of Economics

University of Canterbury

21 September, 2000

Financial support from the program “Risques collectifs et situations de crise” of CNRS is gratefully acknowledged.

1. Introduction

In 1979 a partial melt-down of the reactor core at Three Mile Island was the worst nuclear power accident to that time. Between June 1985 and January 1987 six people suffered massive radiation overdoses during treatment using the Therac-25 machine. In the summer of 1991, ten million subscribers in the United States lost telephone service for durations ranging from ten minutes to six hours. In 1996, after ten years of development and a \$500 million investment, the first Ariane 5 rocket exploded shortly after takeoff. In each of these cases, technological failure occurred. Loss of life, loss of huge investments, or loss of consumer services, caused disasters in a broad sense. In each case, the source of the disaster was uncovered and steps were taken to ensure that it did not re-occur. The disaster was followed by learning. The purpose of this paper is to investigate that learning, and to understand the economic aspects of “learning from disaster”.

From the point of view of the economics of knowledge, a disaster, like many other phenomena, is an opportunity to produce knowledge, to reduce the probability of future disasters, and to improve the performance of technological systems in general.¹ The opportunity lies in the fact that a disaster can be seen simply as an unplanned experiment producing data about the system concerned. In principle, these experiments were always possible – the European Space agency could deliberately have exploded a rocket – but the costs (broadly defined) of performing such an experiment were extremely high. Ex ante, expected costs exceeded expected benefits, and the experiment was not run. But ex post a disaster has the effect of lowering the cost of the experiment to zero – the experiment has been performed. The issue is then to profit from the experiment, and thus simply to maximize cognitive, technological, and organizational benefits.

Few empirical studies explicitly address the economics of technological disasters as learning opportunities. Studies of disasters tend to focus on what, or how much was learned, and they do show that learning tends to follow a disaster. For example, the gas spill in Bhopal and Three Mile Island have both been examined to see the consequent changes in operation, organization and reliability. The Bhopal chemical disaster changed the way the chemical

¹ For recent attempts at qualifying learning as the production of new knowledge “while doing”, see for instance von Hippel and Tyre (1995), Pisano (1996) and Thomke, von Hippel and Franke (1998).

industry organizes chemical stocks and storage, safety standards and safety procedures for example.² Three Mile Island prompted changes in the US nuclear industry – in equipment, safety practices and regulation – resulting in increased reliability.³ But these examine particular cases, and the things learned from them rather than the general incentives and forces surrounding this sort of learning. In this paper we wish to turn the question on its head and ask not what does our learning tell us about Bhopal, but rather what does Bhopal tell us about learning.

While learning does typically occur following a technological disaster, it is not automatic. A collection of organizational conditions and incentive mechanisms come into play to determine the course it takes, the kinds of things learned, and the extent to which this learning is diffused beyond the immediate actors. The factors that affect the organizational, technological and institutional learning that follow technological disasters are the focus of this paper.

The paper is organized as follows. In the next section we discuss the nature and types of disasters. The following section presents an analysis of the types of knowledge that disasters can produce. An economic framework will then be presented which identifies key economic factors affecting the amount of learning that occurs. A detailed case study of the Therac-25 episode is presented and used to illuminate the theoretical material. Finally, we present a discussion of policy implications of the analysis.

2. Types of Disasters

Every disaster is unique, and the type of learning, the population of agents who benefit both directly and indirectly from it, as well as the economic value of it will all be specific to each event. Nonetheless, it seems important to discuss *types* of disasters, as different types will invoke different kinds of learning, and will similarly present different incentives for actors. We begin by addressing the sources of disaster.

Technologies can often be better described as systems made up internally of constituent parts that are linked, and externally as a whole linked to other parts of its environment. In addition technological systems operate in some sort of state space which we can use to

² Fabiani and Theys (1986).

describe “what it is doing” and “what is going on around it”. Knowledge about this state space, and the probabilities of entering different parts of it act to define the types of disasters that can occur. We identify two distinct stylized types of disasters. While our discussion draws a stark distinction between types, this is obviously a caricature, as most disasters will have elements of each of our types.

Improbable events

When technological systems are being designed and constructed, the notion of the state space is used to test the system for performance characteristics, safety and reliability. But state spaces tend to be very large, particularly for complex systems comprising many interacting parts. To prevent every possible mishap in every part of the state space would be prohibitively expensive. Standard economic analysis dictates which possible accidents to prevent: those for which the expected loss exceeds the cost of preventative measures. If the mishap is unlikely (because the system is unlikely to enter the part of the space in which the mishap would occur) then regardless of the actual, realized costs if it does occur, the expected cost of the mishap will be small, and it is unprofitable to prevent it.⁴ For instance, consider the crash of flight TWA 800 in 1996. Boeing’s engineers clearly knew what causes explosions in the vicinity of jet fuel and that explosions were possible in some models of its aircraft. Indeed, results of the inquest showed that in the 1970s the Federal Aviation Authority (FAA) asked Boeing to make changes to the central fuel reservoir which, had they been made, would have prevented the explosion on TWA 800. Boeing resisted, arguing that these events were extremely improbable, and that to make these changes would involve very large costs for very small (in expected value) benefits. The FAA concurred. But TWA 800 did enter this improbable part of the state space, and the expected behaviour occurred.

³ David et al. (1996).

⁴ This is not the argument that private firms are cold blooded. Precisely the same reasoning applies when the analysis is done from the social point of view: preventing disasters uses social resources; incurring disasters uses social resources. It does not make social sense to prevent a disaster for which the resources used to prevent it exceed the resources used if it occurs, in expected value. “Resources” here can be defined as broadly as one wishes.

Ignorance

A second type of disaster is caused by ignorance.⁵ There are several things that can contribute to ignorance about how a system will behave in use. One is related to issues of state space. It is impossible to test every point in the space of a system, or even every region of a space, so testing and considerations about performance are based on notions of where the system is likely to reside. The less likely it is (thought) that the system will be in part of a space, the less important it is to examine how it will behave there. Economic considerations of *risk* typically demand that some parts of the space are not carefully examined. Thus, particularly for new, novel systems there will be areas of the state space in which behaviour is not known. If it enters those parts of the space, its behaviour is unpredictable, not in the sense of being erratic, but rather in the sense of being unknown *ex ante*.

Other forms of ignorance can be, more simply, just ignorance. There is economic *uncertainty*. Consider the case of the de Havilland Comet jet crashes that occurred in 1954. After two years of safe operation a Comet crashed after take-off in Rome. Thirty-five people died. Flights were briefly suspended, but a second plane crashed shortly after they resumed. An intensive investigation eventually came to the conclusion that the fault lay with metal fatigue from the high speeds and high altitudes, conditions which had previously been unknown to aeronautical engineers, and which they had thus simply not considered. Stanley (1986, p. 54) commenting on this disaster reports that “At the end of the war de Havilland had been sufficiently courageous to venture into the unknown and design and build the world’s first jet aircraft. The two accidents had been the result of factors beyond the limit of contemporary knowledge...” Similarly in the case of the Silver Bridge collapse in Ohio in 1967: it took four years of research and inquiry to discover that the cause of the collapse was metal fatigue. Both disasters resulted from ignorance about the performance of the technologies because at least some parts of them were new and untried. Both, we should note, produced knowledge instrumental in creating safer technologies.

There is a distinction between these two types of ignorance. In the first, the general principles underlying the technology in question are known, but some of the specific

⁵ This type of disaster would seem to be intimately related to the concept of bounded rationality, whether through the costs of acquiring and processing information, or through limitations on the cognitive abilities of human beings. Conlisk (1996) provides an excellent review of the area of bounded rationality in economics.

knowledge relating to implementing the technologies is unknown. That is, the overall form of the space is known, although specific parts of it are unexplored, and thus behaviour there is not known. In the second type of ignorance some of the general principles underlying the technology are unknown. Substantial portions of the relevant state space are not known to exist, and thus performance there is not considered.

The Issue of Human Error

The discussion above focuses on technology performance and whether it is known or anticipated in the design of the technology. There is another issue, though, having to do with the use of a technology, namely the role of humans. It is common to make a distinction between the physical technology, the human operator, and the relevant organizational structure in analyzing the causes of a technological disaster. Where we see this most commonly is in reports following accidents which ascribe the accident to “operator-error”.⁶ Consider the case of the Three Mile Island disaster in 1979.⁷ The initial problem was the failure in the pumps supplying water to the steam generators. By itself, this failure was not a major problem. What turned this into a disaster was the reaction of the control room operators. They made several mistakes that compounded the effects of the failed pumps and led to the full-scale crisis.

Many human errors have essentially no untoward consequences. Some are simply not important, and cause no pathological behaviour. Many errors that would cause pathological behaviour are anticipated and guarded against in the design. Others, unanticipated, are counteracted simply by general error correction or safety mechanisms that are built into the system.

But human action will cause problems if it takes the system to unanticipated places and generates behaviour outside the range of the safety mechanisms. This we can incorporate into our first category of disasters – an improbable event. An extreme illustration would be a pilot who deliberately crashes a plane, or a terrorist who puts a bomb in the hold. Human error can be involved in the second type of disaster. Operators can take a system to a place

⁶ This is a very common explanation of airplane crashes.

⁷ See Leveson (1995, pp. 619-639).

that designers did not really know existed. We argue below that this is the cause of one of the sets of Therac-25 incidents. This would fall into our second type, above.

We should also note, though, that disasters attributed to human error can have deeper causal factors. That is, operator error can stem from design or operational features of the technology or the organization running it. Again consider Three Mile Island. Operators did make mistakes in responding to the initial problems, but the Kemeny Commission pointed to design elements that contributed to the mistakes themselves: the control room and information systems were badly designed. Within minutes of the first alarm going off, over 100 other alarms were tripped, overwhelming the operators' ability to gather and process information. Key alarms were not grouped together or distinguished from less important alarms, and some were visually obscured. Some instruments stopped giving readings after a certain magnitude so operators lost valuable diagnostic information. Further, some data that would have been useful were not available as designers had not anticipated that they would be needed and had thus not installed the necessary instruments. None of the operators on duty at the time had knowledge about nuclear engineering, none was a college graduate, and none was trained to handle complex reactor emergencies. From these few added details alone, it should be clear that the errors made by the operators reflected things like poor design of the physical technology (the human-machine interface in this case) and problems with the organizational structure (inadequate training and a poor mix of human capital of the people on duty), rather than simply incompetence of the humans involved. Thus, human error in this case has to do with ignorance in design rather than with operator mistakes *per se*.

3. On the Nature of the Knowledge Learned and the Learning Process

Any experiment produces knowledge which can be classified across two dimensions. The type of knowledge produced, defined in this way, has different properties and different value as it is diffused. In addition, different types of knowledge tend to be produced at different phases in the response to a disaster.

Types of Knowledge

Knowledge can be local or systemic; it can be specific or generic. This categorization is encapsulated in Table 1. Local, specific knowledge pertains to some small part of a single

technical system. For instance, if a pump fails, local, specific knowledge would be about how to make a pump that will not fail. Systemic, specific knowledge would be about how the pump interacts with different parts of the system, referring to a particular technological system, but not necessarily generalizing beyond it. Here, the knowledge is not so much about a particular part of the system, but rather about how its parts interact to generate system behaviour. Generic, local knowledge has been abstracted and made general so that it will apply to technological systems that involve a similar part or more generally sub-system. The knowledge learned from a pump failure in a specific system may be generalized to pumps functioning in any system. Finally, systemic generic knowledge is about how system components interact, but applies beyond one technology. Knowledge, for example, about interactions between humans and machines and the interfaces that mediate them in the Three Mile Island system applies equally to other systems that need complex operator input. This type of knowledge has a large public good element, given that it applies to many technologies.

Table 1: Types of Knowledge

	Local	Systemic
Specific	One isolated part of one technological system.	Many interactive parts of one technological system.
Generic	One isolated part relevant to many technological systems.	Many interactive parts relevant to many technological systems.

Finally, some of the knowledge gained from a disaster may not be related specifically to the technology in question. For instance, the radioactive cloud emitted by Chernobyl was followed by many cases of radiation illness, leukemia in children, and a radioactively poisoned natural environment. It is possible to generate knowledge about each of these phenomena in unprecedented ways. But this knowledge has nothing to do with the generation of electricity using nuclear energy per se. It is in fact a by-product of the disaster, and thus we refer to it as by-product knowledge (see for instance Schelling, 1995). As with knowledge relating to the technology itself, by-product knowledge can be specific or generic, local or systemic.

What is Learned?

Two issues affect the nature of what is learned following a disaster. First, there is a natural temporal progression of learning following the event. Second, whether the system is mature or immature can play a role, particularly in the generality of what is learned.

Learning Dynamics

Following any disaster there is a response: to find the problem and control its consequences. Here learning is local and specific. Agents search for “the faulty part” in order to implement damage control and restore the system. But even if the failure is attributable to a single part, often the event reveals systemic weaknesses, which could contribute to this or other potential failures. Three Mile Island provides a good example of this. Learning about this will be local, but is also systemic in that it is about the interactions of different parts of the technology and how these create system behaviour. Improving the performance of the technology as a system is the goal in this phase. Finally, and especially if the disaster has gained public attention, we see lessons learned about generic aspects of the technology. Better understanding of general principles can ensure that similar technologies do not suffer the same types of problems. Here the goal is to generalize both the local and systemic knowledge beyond the technology involved to a broader context. It is in this phase that potentially large public value is created through what initially was only a disaster. Naturally, as we have described it this sequence is idealized, and how much learning eventually takes place remains dependent on many factors.

Mature and Immature Systems

The nature and scope of learning can change depending on whether the disaster occurs in a mature or an immature system. In mature technological systems, it is relatively straightforward to generalize what is learned beyond the particular installation involved, at least to other installations of the same technology. One feature of mature systems is considerable standardization across installations. Dominant designs exist and variety has been dramatically reduced so what is learned from or about one installation is likely to be applicable to many.⁸ We should add a cautionary note here, though. The presence of a

⁸ For discussions of this process see the literature on technological standardization. For example, see Arthur (1989), David and Greenstein (1990), Cowan (1991), Choi (1996), and Tassej (2000).

dominant design often implies the presence of a dominant problem-solving heuristic. This can create a barrier to solution if a problem is not well-addressed in the dominant heuristic. Thus while mature systems offer the opportunity for significant generalization beyond the installation, they can also make learning difficult if the problem lies outside the scope of the dominant heuristic. A second feature of mature systems is that they tend to be associated with institutions and organizations which, as part of their structures, have in place procedures designed for learning and knowledge diffusion. Codification and storage of knowledge contribute to rapid diffusion and an extensive ability to broaden what is learned through relating it to general, existing archives.⁹ Further, diffusion is facilitated because the users have a common body of knowledge (related to the existence of a common problem-solving heuristic) and thus similar absorptive capacities.¹⁰

In younger systems the ability to generalize is circumscribed. When a technology is young, it typically exists in several variants, each having idiosyncracies peculiar to its environment. There is no dominant design, as users and producers experiment with different features: how to value them; improve them or add new ones. In this case, the knowledge produced may not be immediately applicable to other instances of the technology, it may be extremely potent in selecting or creating the dominant design. A disaster in one variant can have a dramatic effect on people's perceptions of it. These can be well- or ill-founded. A disaster is often thought to show that this variant is unreliable, (and dangerous) even if all variants are in fact prone to the same event. Disasters are events that occur with some non-zero, non-unit probability. If all variants are susceptible to this event, it is not necessarily the case that the one *most* susceptible will be the first to suffer it. But the first to suffer it is very likely to be *perceived* as the most susceptible. This raises attention to secondary effects: disasters feed into technology choices among young, developing technologies. They can thus be very important in the evolution of a system, and will almost certainly tilt the system towards or away from certain technologies.¹¹

⁹ Cowan and Foray (1997) discuss the economics of codification and how codification relates to the diffusion of knowledge.

¹⁰ The presence of mature associations and institutions can also block learning if they seek first to protect their members. See Mueller (1997, Chapters 13 and 16) for a general discussion of the related literature.

¹¹ For a discussion of centralized learning and technology selection see Cowan (1991); for more general models of incomplete information and imitative behaviour see, for example, Bikchandani et al. (1998) or Young (1993).

In older systems, because of standardization and emergence of dominant design, what is learned is likely to be generalizable. However, because by definition an older technology has accumulated considerable operating experience, learning will tend to be more significant when disasters are experiments with unusual parts of the state space, especially parts of the space thought to be sufficiently improbable that little study has been made of them.

In younger technologies by contrast, it may be difficult to generalize learning beyond the specific system involved due to the lack of standardization and presence of many variants. Scope for learning may be high, though, regardless of the type of disaster, simply because little is known and little experience has been accumulated, so any experience is likely to inform.

4. Factors that Affect Learning

How much learning will follow a disaster? To answer this question we look at factors that have the potential to affect learning and the manner in which they do so. From the social point of view there are several critical processes: the production of knowledge or information; some central agglomeration of information; publication both of symptoms of the disaster and of what is learned from it.

Visibility

Some disasters are highly visible, such as the explosion of the Challenger Space Shuttle or the Ariane 5. Others, such as Therac-25 or the new air traffic control system, are not. A visible disaster casts a general, public pall on the reputation of the organization involved. Public scrutiny occurs and disavowal of the products of the firm or the projects of the agency may follow. This creates strong incentives to make a virtue of necessity, by turning the disaster into a dramatic learning opportunity. Simply to fix the problem is the bare minimum. Ensuring that it does not recur, and simultaneously preventing other, different disasters becomes the best possible response. Only in this way can public confidence be restored.¹² When a disaster is not visible, the incentives of the technology's operator or

¹² *Johnson and Johnson* was able to save Tylenol with this approach. Poisoned tablets were inserted into bottles of Tylenol and returned to the shelves of a pharmacy in Chicago. When this news broke,

producer are to satisfy those immediately affected. This involves simply fixing the problem and assuring this user that it will not recur. Beyond that, all else equal, incentives are weak. Learning about a disaster demands costly resources, and furthermore may cause damage to the reputation of the responsible organizations. Incentives may rationally invoke a minimalist response.¹³

Personnel

When a disaster occurs due to ignorance, it is seldom obvious what information is needed to diagnose and fix the problem. People involved simply do not know in advance what they are looking for and diagnosis often proceeds through recognizing empirical regularities (or irregularities) and connections among them. The ability to make these connections is central. Highly-skilled personnel are obviously a vital input to rapid and effective learning, and two qualities seem very likely to be desirable here: an advanced, and relatively general education, (whether formal or informal), and an approach to the problem that is not tightly bound to a detailed model of how the technology works. Obvious “experts” are those involved in the design, construction and operation of the technology. But in order to design, build or operate a technology it is vital to have a good idea of its workings, which implies having a detailed model of how it functions. But being steeped in such a model can be a hindrance to diagnosis particularly, especially if the problems have occurred in a part of the operating space that is unexplored, or is “outside the bounds of reasonable performance”. This follows from the common fact that models are designed to describe not the entire world but only a small part of it, and even a model of a single technology will describe its behaviour only under circumscribed conditions. If the technology operates outside those conditions, the very natural tendency simply to extend the model’s domain to include these

Johnson and Johnson recalled all bottles of the product (evoking predictions of a re-branding) and introduced “tamper-proof packaging”. Tylenol was saved as a brand name, and many products are now tamper-proof.

¹³ The implication of costly information or limited information processing abilities as shown by Cowan (1991) and Bikchandani et al. (1998) is that one piece of “bad news” which may only have a moderate impact on the reputation of a firm could lead to drastic consequences in terms of lost sales. On the other hand, the passage of time enables a firm to build up an “installed base” of customers and information about the firm and its products, which makes such a drastic impact from one “bad” event less likely. This feature of the world naturally leads firms not to divulge information about bad events. Furthermore, there is a significant cost to monitoring any one firm, the outcome of which usually has a strong public good aspect to it. This leads to a low probability of a firm being found out to have not divulged information, further increasing the incentive not to divulge the existence of a

conditions can mis-lead rather than lead towards a solution. We see this very clearly below in the Therac-25 case.

Professional and Regulatory Organizations

Some disasters encompass several different but related incidents, separated either spatially, temporally or both. Information relating to a disaster is dispersed, making it difficult to establish patterns in it. As a result, the presence of organizations such as professional bodies (for example, the Institute of Electrical and Electronics Engineers in the United States or the Union des Industries Chimiques in France) and government agencies (for example, the US Federal Communications Commission or the Direction de la Sûreté des Installations Nucléaires in France) can act to amass information, allowing general patterns to be seen, and thus promoting learning.¹⁴ Further, a central body can act as a clearing house not only for information but also for the knowledge that is created. In this function it can be central in the generalizing aspect of learning from disaster.

Of course it is also possible that this type of agent could suppress information and thus retard learning.¹⁵ For instance, both the United States General Accounting Office and the United States House of Representatives found that the Federal Aviation Agency (FAA) acted to suppress information relating to problems with the modernization of the United States air traffic control system, which proved a technological failure – the actors involved in the project misinterpreted the potential of the technology.¹⁶ For a professional body, this sort of incentive exists when the (unsafe) performance of a technology puts the value of the human capital of its members at risk. A similar incentive will exist if publicity stands to expose deficiencies in the quality assurance of the organization. For a government agency,

bad event.

¹⁴ Kirby (1988) reports that 24 of the 64 trade associations she surveyed had formal information exchange programmes. She also examines the conditions under which such associations are likely to lead to information sharing, although in a different context from that discussed in this paper.

¹⁵ This double-sided nature of such organizations is analogous to the case studied by Leland (1979), who looks at how industry associations set minimum quality standards when asymmetric information between buyers and sellers is present. The association can improve the welfare of consumers by imposing a collective minimum quality standard, but it can also use the standard to coordinate behaviour among firms in an attempt to extract surplus from consumers.

¹⁶ United States House of Representatives (1997, pp. 6-7) and General Accounting Office (1996, p. 24 and p. 26).

perceptions of its ability to assure quality are an issue, as is the possibility that a disaster will be interpreted as agency incompetence, as was the case with the FAA.

Reputation and Human Capital

In any disaster, individuals and organizations have assets whose value could be affected.¹⁷ Some parties will suffer a decrease in their asset values if the disaster is made public because those assets are tied to the technology involved. For instance, when it emerged that a piece of software from DSC Corporation contained a bug responsible for a series of major outages of the United States telephone system, the firm's share price fell from \$11.38 to \$4.81, on the back of a large fall in sales. Naturally, any agent has incentives to defend the value of his assets, and often the obvious immediate defence is to suppress information that will destroy value.¹⁸ Suppressing information clearly impedes learning.

In other cases though, agents' assets will decline in value if the information is not made public. For instance, when the Ariane 5 exploded, the European Space Agency (ESA) had every incentive to learn about what caused the disaster and to correct the faults as quickly as possible. In this situation, its reputation capital had already been adversely affected by the explosion, but by showing that it was only a teething problem, and more importantly that the ESA was on top of the situation, it could repair some of the damage. In the case of the Ariane 5, ESA's handling of the entire episode, including the investigation into the disaster and release of information that resulted from it, was seen as a model response.¹⁹

5. Case Study: Therac-25

To illustrate further this theoretical structure we present a study of the Therac-25 radiation treatment machine.²⁰ The case, which concerns radiation overdoses of patients during the

¹⁷ The human capital of the people involved, reputation capital, knowledge capital or physical capital, of individuals or organizations involved are all examples.

¹⁸ On the general issue of reputation capital see Shapiro (1983), and Milgrom and Roberts (1992, pp. 259-269).

¹⁹ The Inquiry Board, complete with terms of reference was established almost immediately after the disaster and released all of its findings at a press conference, as well as publishing them on the internet, only one month after beginning the investigation. The ESA made many press releases about the disaster and investigation during and after the investigation.

²⁰ Detailed treatment of the Therac-25 disasters can be found in Leveson (1995) and Leveson and Turner (1993). Additional information can be found in *Computer*, (1993), 26(9), p. 109; *Computer* 26(10),

treatment of their cancers, is useful for understanding factors that affect learning from disaster because it involves a series of incidents that occurred at different places and at different times. Moreover, this case illustrates different aspects of learning, from generating the information about a failure to the production of new knowledge. It is particularly interesting in that two types of accidents, having different causes, occurred. Notably, the response of the producer of the machine changed during the period in which the accidents were occurring. Its initial responses severely curtailed the learning, whereas its later responses were conducive to it.

Background

In 1982, Atomic Energy Canada Limited (AECL) introduced a computer-controlled radiation therapy machine, the Therac-25. This superseded the Therac-20, a mechanically-controlled machine. Over the next two years it was used thousands of times without incident.²¹ But between June 1985 and January 1987 six people suffered massive overdoses of radiation while undergoing treatment with the Therac-25. Three people died, and three others suffered major physical injuries. These were the first deaths attributed to treatment with a therapeutic radiation machine.²² The faults at the core of these accidents were eventually corrected, and the machines have operated without incident since. Further, the general performance of the machine was improved as a result of these experiences with it. More importantly, however, this event created large amounts of both local and general knowledge about embedded software, software re-use, and quality assurance practices in software engineering. All of these go beyond the Therac-25 machine, applying to the entire software industry and in fact to safety critical systems in general. Indeed, the Therac-25 episode is considered one of the paradigm cases from which software engineers learn about “good practice” in North America. The episode also illustrates the role of the interaction between industry and regulatory bodies in learning from disasters. The US Food and Drug Administration (FDA) and the Canadian Radiation Protection Bureau (CRPB) were both

pp. 4-5); *Communications of the ACM* (1990), 33(12), p. 138; Joyce (1987); Plummer (1986); and Thompson (1987).

²¹ By 1986 there were 1144 radiotherapy facilities in the United States employing about 2,000 radiation therapy machines in total. These machines were used to treat about 600,000 patients each year. If the average of 306 patients per year used the Therac-25 machines then it would have treated between 3,366 patients and 13,464 patients, depending precisely when the 11 machines were sold. See Karzmark et al. (1993, Appendix B).

involved and played important roles. Finally, the episode can also be used to illustrate how the incentives of different agents affect the learning that takes place after a technology fails.

The Therac-25 is a linear beam accelerator used to treat cancerous tumours, and unlike other radiation therapy machines of that era it could be programmed to fire both accelerated electrons and X-ray photons.²³ The Therac-25 was based on its two direct predecessors, the Therac-6 and the Therac-20. Relative to the Therac-20, the Therac-25 incorporated several innovations that made the machine more flexible and able to treat a wider variety of tumours. In addition to these hardware innovations, it introduced innovations to the control system, largely in the machine's embedded software. Both the Therac-6 and -20 included software, but neither was dependent on it for control. In the Therac-25 version, though, many of the hardware mechanisms of the Therac-6 and the Therac-20 were replaced with software functions. Positioning the patient, setting type and amount of radiation, checking patient position, machine position and settings, shutting down the machine in case of malfunction or bad settings were all now software-controlled. It is important to note that this software was not written from scratch. Following the common belief that software re-use not only speeds up development but also improves reliability, much of the control software developed for the earlier generation machines was re-used. Software re-use was thought to be good engineering practice, since old software has been tested, not only in development, but also in its application. This was thought to improve its reliability.²⁴

AECL performed a safety analysis on the machine in March 1983, but appears only to have examined the hardware, and failure rates were given for hardware alone.²⁵ The AECL report on this analysis appears to assume that software does not degrade, and that computer errors are caused by failures in hardware or by interference from background radiation.

²² Joyce (1987, p. 6).

²³ Of the 2,000 radiotherapy machines employed in the United States in 1986, 1,200 were microwave linacs, of which the Therac-25 was an example. The trend in the United States at the time was the replacement of the Cobalt-60 machines by microwave linacs. See Karzmark et al. (1993, p. 287 and Appendix B).

²⁴ See Yourdon (1993, p. 35 and Chapter 9) and Hatton (1997, p. 51).

²⁵ Leveson (1995, p. 520).

The Disaster

The Therac-25 disaster includes six separate incidents. In each case, a patient received a massive radiation overdose, but the machine operator was unaware that this had happened, (and in some cases insisted that it could not have happened), and the medical staff did not diagnose it until some time after it had taken place.²⁶ When AECL was notified about the possible problems with the machines, in the early cases it was unable to produce doses to match the circumstances of the accidents. Once the cause had been found, however, the events were indeed reproducible. In our schema, the technical cause of the accidents was that the machine entered a part of the state space that was not considered in its design. In essence, experienced operators made a mistake in data entry, but corrected it quickly enough that the machine, having logged the original input, failed to log the correction. Normally mistakes would not have been a problem, as the machine was designed to catch errors, but in this case the sub-routine which set the parameters operated on the incorrect input, while the checking sub-routine, called slightly later, operated on the corrected input. This could only happen with an experienced operator who was able to make corrections quickly: between the time the setting was made and the time the checking was done by the software. It seems likely that this scenario never entered the minds of the designers as a possibility. Thus the case is a disaster caused by ignorance.

1. The first incident took place in Georgia in June 1985. A patient received a massive overdose during the course of treatment. When the patient asked about the cause of his intense, unusual pain, he was assured by the technician that the machine could not produce incorrect radiation dosages. The technician was of the opinion that the machine had operated normally. The resident physicist was somewhat more skeptical, however, and thought a radiation overdose seemed a likely explanation of the patient's symptoms. The event was reported neither to AECL nor to the FDA nor the CPRB. It was considered a fluke by the user.

2. The second incident took place in Hamilton, Ontario, the following month, again involving a massive overdose. Both vendor and regulators were notified, and AECL investigated. It was unable to locate the source of the overdose, but did discover some

²⁶ A radiation treatment typically involves in the neighbourhood of 200 rads. The accidents involved doses in excess of 15,000 rads. It is estimated that a dose of 500 rads to the entire body is fatal 50% of

problems with a mechanical part unrelated to the incident. Here, local, specific learning took place regarding the mechanical control of the turntable. It applied to all installations of the Therac-25, and so was generic in that restricted sense. But it did not have to do with the overdose incidents. Again, the accident was treated as a one-off fluke, and the conceptual model of how the machine worked, used both by AECL engineers and the operators, remained unchanged.

3. In December 1985 in Yakima, Washington, a third overdose occurred. AECL was notified, but denied that such a thing was inherent in the machine, and the event was again treated as a fluke. No learning occurred.

4. In March 1986, in Tyler, Texas, the fourth overdose took place. Again, AECL was notified, but its technicians were again unable to reproduce the result, and were puzzled by it. Their investigation indicated that the machine was functioning normally and should have delivered the right dose. Again their conclusion was, "One-off." Interestingly, when asked by the hospital, AECL responded that it knew of no other similar incidents. If AECL is treated as a monolith this response seems false. The patient in the Marietta, Georgia, incident instigated a lawsuit in which AECL was named. The corporation should have been informed before the first Tyler accident. If AECL is not treated as a monolith, however, it is possible that the technician stating that he knew of no other incidents could have been speaking ingenuously. Whatever the internal workings of AECL over this matter, it is true that up to this point each user was of the opinion that his experiences were unique to him. As a result, there was no pooling of experience, information or expertise in an attempt either to keep informed or specifically to understand the events.

5. The fifth incident changed this, as it occurred in the same facility, with the same operator, only 21 days later. Through the efforts of the resident physicist, the immediate cause of the overdose was found. His suspicions had been raised by the previous incident, and he had made inquiries of his colleagues at other institutions. He was willing to entertain the idea that the machine was producing this effect in a deterministic way, and had no prejudices regarding what the source could or couldn't be (or if he did, he was willing to overlook them). Interestingly, his approach to the problem seems to have been highly empirical – he asked the operator to repeat her actions as accurately as she could. By working with the

the time.

operator in this way he was able to reproduce the event. The absence of a strongly held theory of the machine's operating conditions, environment or user interactions, permitted him to take the machine into a region of its operating space that was not considered by others investigating the events.

Two things are key here. First is the agglomeration of information. With two events occurring in a single venue, and under the "intellectual authority" of one person, here the resident physicist, enough information was assembled in one place to convince some actor that the events should be linked and jointly investigated to determine the cause. Two occurrences of roughly the same event made the "fluke" explanation too improbable to be credible. There must be a deterministic cause. Prior to the second accident in Tyler, accidents had occurred at geographically distant locations, and information had not been agglomerated, certainly among the users. Thus, each user thought his experience was unique. Whether or not AECL had agglomerated the information internally is unclear. Its response to the users' experiences was to confirm their uniqueness. Whether this was deliberate obfuscation on the part of AECL, or whether internal communications within the firm were not what they might have been is a matter of speculation. With the second event at Tyler, however, information agglomeration did take place, and the cause was quickly found. The second key in this episode was that the investigator, here again the physicist, was not tied rigidly to beliefs about the machine's performance. He did not have strong views on state spaces, neither what they looked like nor the likelihood of visiting different regions of them. Those who develop complex systems must, in the course of development, form precisely such strong beliefs. If they did not, development would be impossible. Thus the presence of an "outsider" was probably crucial. The results of the investigation were reported to AECL and eventually to the FDA. Furthermore, the incident was picked up by the media and entered the public realm. The fact that the matter gained public prominence and also the fact that the information provided in the report to the FDA was highly detailed regarding the causes of the events meant that it clearly had to act. It did so, forcing AECL to accept that a problem existed and that a solution must be found and publicized. At this point we see an interesting change in the behaviour of AECL. After users of the Therac-25 had been informed about the problems editing inputs, and an immediate, short term fix had been disseminated (it involved disabling one of the editing keys on the keyboard – removing the key cap and putting insulating electrical tape on the contacts!), AECL

continued to work on a more permanent solution. The events were now public however, and, seemingly in response to the problems with the Therac-25, a users group had formed.

6. When the next mishap occurred, for a second time in Yakima, in January 1987, all Therac-25 users were immediately informed, and a solution was quickly found. Again it was a software fault. The accident was of a different type than those in Tyler, and a different part of the software was the cause, but nonetheless, the source was quickly isolated and a solution quickly found. Response to the second incident in Yakima was different from responses to previous incidents. It was immediately publicly acknowledged, which created very big incentives for a convincing solution to be found. The existence of the user group, and the presence of FDA and CRPB scrutiny made it impossible to keep the diffusion of information circumscribed. The credibility of the machine, and of AECL, was at stake. Second, having discovered that the software was fallible, the model and problem-solving heuristic in which AECL engineers had been entrenched, namely that the software was robust and the problems were more likely to occur in hardware, had been lost. Search for a solution was not restricted to particular parts of the system (that is to say to hardware components), but took place throughout the technology complex (and indeed, given the discoveries in Tyler, probably focussed on software, in contrast to previous investigations which had severely down-played that part of the technology). This broadening of scope in the search for the cause was central in the speed of solution, and could only exist after the engineers had changed their basic understanding of how different parts of the technology interacted, and, perhaps more importantly, on what had been learned about the robustness of the software that was inherited from the previous generations of the machine. Both of these things – the change in views about interaction, and the reliability of re-used code – have since been generalized into software engineering textbooks.

Factors that Affected Learning in the Therac-25 Disaster

It is clear that several types of learning occurred as a result of the Therac-25 disaster, as summarized in Table 2.

Table 2: Summary of the Therac-25 Disaster

Place	Date	Those Notified	Outcome with Respect to the Therac-25	Type of Learning

Place	Date	Those Notified	Outcome with Respect to the Therac-25	Type of Learning
1. Marietta, Georgia	25 June, 1985	Lawsuit filed against AECL by patient.	No changes	None regarding the technology
2. Hamilton, Ontario	26 July, 1985	AECL, CRPB, FDA	Changes to microswitch	Local ²⁷ & specific
3. Yakima, Washington	December, 1985	AECL	No changes	None
4. Tyler, Texas	21 March, 1986	AECL	No changes	None
5. Tyler, Texas	11 April, 1986	AECL, State of Texas Health Dept, FDA, Users	Major changes to computer hardware & software, computer-user interface, and manuals. Changes to software testing practices. User group formed.	Initially local & specific. Subsequently, systemic & specific. Eventually systemic and generic.
6. Yakima, Washington	17 January, 1987	AECL, FDA, Users	Detailed instructions about how to avoid fault. Added to software changes proposed after Tyler incident.	Local & specific

AECL did eventually learn what was causing the incidents and successfully repaired them. Other learning also occurred, relating to the software industry, the design and production of safety critical devices, especially those that relate to the medical industry, and to regulatory bodies and how they approach their tasks, especially in a world of rapid technological change.²⁸

²⁷ This could be classified either as local or generic: it is local to the Therac-25, but generic to all installations of it.

²⁸ For instance, up until the Therac-25, software re-use had been seen as a “cure” for continuing problems with software defects. This myth was shattered by this case because the code containing the defects actually re-used code from earlier versions of the Therac machines.

In the case of the Therac-25 there were several factors at work affecting the amount and rapidity of learning.

Regulatory Bodies

The first factor, acting to increase the amount of learning, was the presence of the FDA. The primary objective of the FDA is to ensure that medical devices and pharmaceutical products are safe. In demanding that AECL make the machine safe, the FDA forced the firm to learn more about how the machine worked. Further, the FDA ensured that users not involved in the accidents received information about the machine, its problems, and the remedies being developed and applied. An interesting side point here concerns the importance of regulatory details. In the case of Tyler, the cancer center did not contact the FDA itself since users of medical devices were not required to do so; only the producers were obliged to report incidents. The FDA was in fact notified by the State of Texas Health Department, which had been notified by the cancer center – under Texas law users must report incidents to the State authority. The two sets of laws, while similar in principle, had quite different effects because of differences in certain key details.²⁹

Personnel

Another important factor increasing the amount of learning was the presence of the professional employees of the users. The professional employees of the users, in this case the physicists, can take credit for the localized learning that occurred. The physicists were able to replicate the faults on at least two occasions and in doing so both provided the key information regarding where the fault lay, and provided data that led directly to improvements in the Therac-25. Most significant here was that these data challenged the conceptual model used by the AECL engineers and forced them to re-think not only details of how the machine worked (what happens when the machine is run by an experienced, fast, operator) but also more generally where faults could lie. Having been forced out of their conceptual model, the AECL technicians were able to address not only the original but also subsequent problems. In addition, the professional employees of the users were instrumental in forming the Therac-25 users group which resulted in widespread diffusion of information.

²⁹ The US Federal law has since been changed.

Public Visibility

The fifth incident, at Tyler, Texas, showed the importance of public visibility. In the first four incidents, awareness was private, and thus there was little pressure on AECL, the FDA, the CRPB, the hospital physicists, or anyone else for that matter, to pursue them, especially as a group. This changed with the fifth incident, which was reported by the media. Publicity changes incentives considerably and pressure to find out what is going on can be severe. The key point here is that the initial dispersion of the information about problems with Therac-25 hindered the learning process. The fifth incident acted as a catalyst for information agglomeration which then triggered the learning process. In this respect, the Therac-25 case differs from the Challenger explosion or the collapse of a bridge under metal fatigue. In cases like the Therac-25 the mechanism or event that creates an opportunity for agglomerating information is crucial. It creates more general awareness that there are problems with the technology, and it creates a much better environment for learning about and from it.

Legal Environment

The ability of injured patients to sue for damages may have adversely affected learning. In most cases the victims of the accidents launched lawsuits against the parties involved, including AECL.³⁰ If AECL were to provide information showing that the Therac-25 had defects this could have negative consequences in any lawsuit, regardless of the fact that this information would contribute to learning from the accidents and to making a safer medical device. This may help explain the reluctance of AECL, especially its quality assurance manager, to divulge information about problems.³¹

³⁰ Joyce (1987) provides details of the legal implications of software defects and of the lawsuits stemming from the Therac-25 disaster.

³¹ Indeed, Orton (1995, p. 675) says in a section on radiation therapy accidents that, "It is vitally important that we learn from these experiences but, unfortunately, very few have found their way into the referenceable literature, probably because of malpractice fears or legal restrictions. It is hard enough that these errors are made in the first place, but the real tragedy is that our ability to learn from them is hampered by our litigious society." In other words, Orton suggests that the potential for being sued is a factor that greatly retards learning from accidents because most of the benefit from the information accrues to society, but all of the cost, mainly in the form of expected litigation costs, is ultimately borne by a relatively small number of individuals.

The Technology and the Environment

Another factor that retarded learning concerned AECL's economic environment. AECL had a new and unproven technology and it was trying to build sales. It had a natural incentive to suppress negative information about its product.³² As highlighted by Bikhchandani et al. (1998), when information is scarce and costly, agents use each others' actions to guide their own. Here, even a minor piece of negative information can lead to large adverse consequences for a person or organization. As a result, situations in which information is gathered by observing the actions of others conduce to retarding the production and diffusion of information relating to a disaster. That this could have been a factor is shown by the fact that AECL renamed its radiation therapy subsidiary after the disaster and then finally sold it; AECL is no longer in the radiation therapy industry.

In a competitive industry one might expect that the competition of AECL would have publicized the problems with the machine.³³ This process may have been inhibited by the fact that the source of the problem lay in the software. Without access to the source code, which AECL refused to release, it was not possible to investigate the software directly. This feature of the technology reduced dramatically the ability of "outsiders" to search for the causes of the disaster. It demonstrates the importance of transparency in technology more generally in learning from a disaster.

Human Capital Tying

A final key factor affecting learning from the Therac-25 disaster was the relationship between the wealth of people and organizations, the technologies involved, and the disaster. Hospital or clinic physicists were committed to radiation therapy technology and not to the Therac-25. Their main interest was to protect the reputation of radiation therapy. Thus they had strong incentives to show that the problem was not with the generic technology but rather with a particular implementation of it. They had a natural incentive to learn about the

³² Even though users were replacing their Cobalt-60 machines with microwave linacs, and the Therac-25 was supposed to offer powerful new features and be very cost-effective, by 1985 it had in fact only garnered 0.6 percent of the market for radiation therapy machines and 0.9 percent of the sales of microwave linacs in the United States.

³³ There were in fact thirteen producers of radiation therapy machines in the world in 1985: AECL, ATC, and Varian, in the United States; BBC, CGR-MeV, Scanitronix, and Siemens, in Western Europe; the Russian government; the Chinese government; and Mitsubishi, NEC, and Toshiba, in Japan. See Karzmank et al. (1993, p. 290).

incidents to show that it was Therac-25 at fault and not radiation therapy in general. The physicists did uncover the faults and were instrumental in forming the user groups. Unlike the clinic physicists, AECL engineers and technicians were tied directly to the Therac-25 itself. Their natural interest was to protect the reputation of the machine and consequently their goal was to show that it was not the source of the problem. The general reluctance to disclose information to users is consistent with this interpretation. Interesting is how that particular incentive can, when coupled with publicity of the events, explain the change in their behavioural attitude from “there is no problem” to “there is a problem but we can fix it and in so doing improve the machine”.

6. Discussion

The non-rivalrous and extendible nature of knowledge implies that it is a quasi-public good. As such, it will in general be under-supplied by a market (Arrow, 1962). This suggests that there may be a role for intervention to increase the production of knowledge in general. In the context of learning from disasters, the knowledge with the strongest public good features is the generic knowledge, since it is most widely useful. Nonetheless, because of the way inventive mechanisms and institutions are involved in the production even of local and specific knowledge, there may be a role for policy to facilitate and speed up learning of this type of knowledge following a disaster.

Information Agglomeration and Diffusion

When a disaster consists of several dispersed events, as in the Therac-25 case, information agglomeration can be a strong factor contributing to the speed and extent of learning. But for this to be effective, there must be rapid distribution of information among the installations and to the central coordinator, who, in the Therac case, was the vendor of the machine, namely AECL. The key here lies in the incentives both of users and of the producers. Are the incentives to reveal or conceal information? From the point of view of the vendor, who often plays the key role both in collecting and in distributing information, incentives to conceal information increase if the vendor has a weak market position, since negative information creates a big risk that market share will be lost to competitors, particularly if the turnover of the technological artefacts is rapid. Similarly, incentives to conceal are relatively strong if information is asymmetric – if the users are dispersed and

report individually to the producer. In this case, a bandwagon away from that vendor (or technology) cannot form since information is centrally controlled and not diffused to potential buyers.

However, in an intensely competitive market, rival firms will monitor each other and incline to publishing each other's problems. As Eads (1980) suggests, diffusion of information can take place by mutual denunciation of rivals, providing that the problems do not cast a pall over the entire industry. But if there exists a credible substitute firm, with a product that is insulated from that sort of spillover, information will be diffused very quickly. In this case, the firm suffering the disaster has very strong incentives to gather information and publish its findings regarding causes and future prevention, and possibly how what it has learned has created improvements to the product. More generally, a competitive industry is likely to enhance the diffusion of information about a disaster and the urge to respond to it, if negative information about competitors is a competitive weapon.

From the Therac-25 case we see that professional associations, user groups and regulatory bodies can play a role in collecting and agglomerating the information needed to diagnose and learn about failures. We should also note that the disaster with the Therac-25 machine did not involve the machine's producers in the events themselves. The events involved only the users directly, and their input was vital in solving the problem and fostering the more general learning that followed.³⁴ This makes Therac-25 different from cases like Challenger. What is made clear by the case is that for information diffusion and agglomeration to take place agents must know they have useful information; they must have incentives to diffuse or collect it; and there must be mechanisms or structures through which this can happen.

Institutions and structures

A relatively common institutional form that conduces to rapid action is the industry association. Very often, providing a forum for information diffusion is one of the key motives to form these associations. This will particularly be the case when the industry feels under threat, either from regulators, from other substitute industries, or from public opinion. Here the industry as a whole has very strong incentives to "keep things under

³⁴ See von Hippel (1988) for accounts of the user's role in technological advance.

control”; to solve problems quickly and very publicly, advertising what and how the industry has learned following this unplanned event.

Commissions of inquiry serve a similar function, being temporary institutions set up precisely to gather and diffuse information. This structure is completely under policy control and typically functions only in respect of a single incident. This single-incident aspect suggests that relative to a well-functioning industry association it may be less effective since it has less institutional background or history on which to draw.

Finally, public regulation can force firms to respond quickly and openly in the face of problems, particularly if their actions affect many other agents. The airline industry is a good case in point here, as is the nuclear industry. However, the legal structure can also act negatively in this regard. If tort law is liberal in its interpretation of damage, and generous in its calculation of damages, firms have strong incentives not to admit fault when accidents occur, and this can create an obvious brake on learning about a disaster.

Production and Diffusion of New Knowledge

Here there are two aspects. Above all, for an experiment to be profitable in terms of knowledge production, the information or facts that it produces in principle must be captured to be analyzed and diagnosed. Instrumentation of a system, which aims to capture facts produced by unexpected events constitutes an important aspect of system design. The example that comes to mind immediately are the black boxes and voice recorders installed in airplanes, and the general principle here is the recording of data, states of the system or interactions with operators while the technology operates, even under normal conditions. But of course disasters occur precisely when the technology operates outside normal conditions, so the instruments must be prepared for extreme values. The second aspect here is relative only to the contents of the learning. To profit from the “experiment” implies discovering elements that had previously been ignored, either consciously or through ignorance. But further, the goal is to develop these discoveries from local knowledge into systemic knowledge. Here “systemic” is considered broadly to include questions of management, maintenance and control, technology and organization. The conjecture here is that an organization that has experienced a technological failure maximizes its learning from the disaster the larger the scope of knowledge production. The temptation should be high in a first instance just to fix the problem. If this problem is localized, only local

knowledge will be produced after the disaster. Organizational incentives are required for learning to extend to more systemic knowledge. The accident at Three Mile Island was in the first instance caused by a pump failure. It could have been fixed, and repetitions prevented, simply by improving that pump. But TMI offered the opportunity to see how other aspects of the system performed under these unusual circumstances. From this much more systemic, and later generic knowledge was created which prompted a significant improvement in control-room design and layout.

7. General Conclusions

This paper provides a framework with which to analyze learning from technological failures from an economic point of view. Our purpose has been to examine technological disasters as (unplanned) experiments that open opportunities for learning to take place. In this respect learning is the production of new knowledge in the course of the use of the technology. What is learned from a disaster then, is knowledge that has not been produced during the conception of the technology, either because it was deemed too costly or simply through not knowing that this knowledge could be produced. Thus we have distinguished between two types of disasters. If disasters caused by improbable events fit well the traditional economic framework that implies that costs of avoiding them are simply too high, disasters caused by ignorance are far more challenging for economic theory since the firm suddenly faces an unknown region of the state space, revealing its bounded rationality.³⁵

While the net benefits of creating certain types of knowledge increase dramatically as a result of a disaster, learning after a disaster is not a spontaneous process. As the case study has emphasized, many factors affect this type of learning. Most have to do with the institutional mechanisms that create the incentives for the actors involved.

The immediate environment of the firm or organization affected by the disaster may have conflicting effects on willingness to learn. Information about a disaster might prove competitively adverse, as publicity about problems or failures can affect users' willingness to trust the technology. At the same time, an active, public response can also act as good publicity. What the case has shown is that conflicting immediate effects can to some degree be controlled when the right institutional mechanisms exist. Compulsory implementation of

data collecting instrumentation, or compulsory problem-reporting to a regulatory agency, for instance, can counteract the influence of adverse factors like the use of information as a competitive weapon.

³⁵ We can see here the classic distinction between risk and uncertainty.

References

- Arrow, K. (1962). "The Economic Implications of Learning By Doing". *Review of Economic Studies*. 29: 155-173.
- Arthur, B. (1989). "Competing Technologies, Increasing Returns, and Lock-In by Historical Events." *Economic Journal*. 99(394):116-131.
- Bikchandani, S., D. Hirshleifer, and I. Welch. (1998). "Learning from the Behaviour of Others: Conformity, Fads, and Informational Cascades." *Journal of Economic Perspectives*. 12(3):151-170.
- Choi, J. (1996). "Standardization and Experimentation: Ex Ante vs. Ex Post Standardisation." *European Journal of Political Economy*. 12(2):273-290.
- Communications of the ACM* (1990), 33(12).
- Computer*, (1993), 26(9).
- Computer* (1993) 26(10).
- Conlisk, J. (1996). "Why Bounded Rationality?" *Journal of Economic Literature*. 34(2):669-700.
- Cowan, R. (1991). "Tortoises and Hares: Choices Among Technologies of Unknown Merit." *Economic Journal*. 101(407):801-814.
- Cowan, R. and D. Foray. (1997). "The Economics of Codification and the Diffusion of Knowledge." *Industrial and Corporate Change*. 6(3):595-622.
- David, P. and S. Greenstein. (1990). "The Economics of Compatibility Standards: An Introduction to Recent Research." *Economics of Innovation and New Technology*. 1(1-2):3-42.
- David, P.A., R. Maude-Griffin, and G. Rothwell. (1996). "Learning By Accident?: Reductions in the Risk of Unplanned Outages in US Nuclear Plants after Three Mile Island." *Journal of Risk and Uncertainty*. 13:175-198.
- Eads, G.C. (1980). "Regulation and Technical Change: Some Largely Unexplored Influences". *American Economic Review*. 70(2):50-54.
- Fabiani, J.L. and J. Theys. (1986). *La société vulnérable*. Presses de l'ENS.
- General Accounting Office. (1996). *Aviation Acquisition: A Comprehensive Strategy is Needed for Cultural Change at FAA*. Report to the Chairman, Subcommittee on Transportation and Related Agencies, Committee on Appropriations, House of Representatives, August, Washington D.C.
- Hatton, L. (1997). "Software Failures: Follies and Fallacies." *IEE Review*. 43:49-52.
- Joyce, E. (1987). "Software Bugs: a Matter of Life and Liability." *Datamation*. 33:88-92.
- Karzmark, C., C. Nunan, and E. Tanabe. (1993). *Medical Electron Accelerators*. New-York: McGraw-Hill.
- Kirby, A. (1988). "Trade Associations as Information Exchange Mechanisms." *Rand Journal of Economics*. 19(1):138-146.
- Leland, H. (1979). "Quacks, Lemons, and Licensing: A Theory of Minimum Quality Assurance Standards." *Journal of Political Economy*. 87:1328-1346.
- Leveson, N. (1995). *Safeware: System Safety and Computers*. Reading, Mass.: Addison-Wesley.

- Leveson, N. and C. Turner. (1993). "An Investigation of the Therac-25 Accidents." *Computer*. 26(7):18-41.
- Mueller, D. (1997). *Public Choice II*. Cambridge: Cambridge University Press.
- Milgrom, P. and J. Roberts. (1992). *Economics, Organization and Management*. Englewood Cliffs: Prentice Hall.
- Orton, C. (1995). "Uses of Therapeutic X Rays in Medicine." *Health Physics*. 69(5):662-676.
- Pisano, G. (1996). "Learning-Before-Doing in the Development of New Process Technology." *Research Policy*. 25(7):1097-1119.
- Plummer, W. (1986). "A Computer Glitch Turns Miracle Machine into Monster for Three Cancer Patients." *People Weekly*. 26:48-50.
- Schelling, T. (1995). "Research by Accident". WP 95-40. IIASA.
- Shapiro, C. (1983). "Premiums for High Quality Products as Returns to Reputation." *Quarterly Journal of Economics*. 98(4):659-679.
- Stanley, S. (1986). *Air Disasters*. London: Ian Allan.
- Tassey, G. (2000). "Standardization in Technology-Based Markets." *Research Policy*. 29(4-5):587-602.
- Thomke, S., E. von Hippel, and R. Franke. (1998). "Modes of Experimentation: an Innovation Process - and Competitive - Variable." *Research Policy*. 27(3):317-334.
- Thompson, R. (1987). "Faulty Therapy Machines Cause Radiation Overdose." *FDA Consumer*. 21(10):37-38.
- United States House of Representatives. (1997). *Allegations of Cost Overruns and Delays in the FAA's Wide Area Augmentation System (WAAS)*. Report of the Subcommittee on Aviation, 1 October, Washington D.C.
- von Hippel, E. (1988). *Sources of Innovation*. Oxford: Oxford University Press.
- von Hippel, E. and M. Tyre. (1995). "How Learning By Doing is Done: Problem Identification in Novel Process Equipment." *Research Policy*. 24(1):1-12.
- Young, A. (1993). "Invention and Bounded Learning By Doing". *Journal of Political Economy*. 101(3):443-472.
- Yourdon, E. (1993). *Decline and Fall of the American Programmer*. Englewood Cliffs, New Jersey: Yourdon Press.

Appendix: Therac-25 Disaster Timeline

1985

- 3 June Marietta, Georgia, overdose.
Later in the month, Tim Still calls AECL and asks if overdose by Therac-25 is possible.
- 26 July Hamilton, Ontario, Canada, overdose. AECL notified and determines microswitch failure was the cause.
- September AECL makes changes to microswitch. Independent consultant recommends potentiometer on the turntable.
- 16 September AECL sends letter to users of Therac-25 claiming an improvement in safety of 10,000,000%.
- October Georgia patient files lawsuit against AECL and hospital.
- 8 November Letter from CRPB to AECL asking for additional hardware interlocks and software changes.
- 13 November Lawsuit filed against the Kennestone Regional Oncology Center, Marietta, AECL, and a servicing company, by the patient overdosed on 3 June.
- December Yakima, Washington, overdose.

1986

- January Attorney for Hamilton clinic requests potentiometer be installed on turntable.
- 31 January Letter from Yakima to AECL reporting overdose possibility.
- 24 February Letter from AECL to Yakima saying overdose was impossible and no other accidents had occurred.
- 21 March Tyler, Texas, overdose. AECL notified. AECL claims overdose impossible and no other accidents had occurred previously. AECL suggests hospital might have an electrical problem.
- 7 April Tyler machine put back into service after no electrical problem could be found.
- 11 April Second Tyler overdose. AECL again notified. Software problem found.
- 15 April AECL files accident report with FDA.
- 2 May FDA declares Therac-25 defective. Asks for CAP and proper re-notification of Therac-25 users.
- 13 June First version of CAP sent to FDA.
- 23 July FDA responds and asks for more information.
- August First Therac-25 users group meeting.
- 19 August AECL sends letter to FDA claiming that in March it had received notice of a lawsuit filed by the patient at Marietta.

26 September AECL sends FDA additional information.
 30 October FDA requests more information.
 12 November AECL submits revision of CAP.
 December Therac-20 users notified of a software bug.
 11 December FDA requests further changes to CAP.
 22 December AECL submits second revisions of CAP.

1987

17 January Second overdose at Yakima.
 26 January AECL sends FDA its revised plan.
 February Hamilton clinic investigates first accident and concludes there was an overdose.
 3 February AECL announces changes to Therac-25.
 6 February FDA contacts Canada's Health and Welfare and advises that the FDA recommends that all Therac-25s be shut down until permanent modifications are made. Canadian authorities concur and agree to co-ordinate with FDA.
 10 February FDA sends notice of adverse findings to AECL, declaring Therac-25 defective under US law and asking AECL to notify customers that it should not be used for routine therapy. Health Protection Branch of Canada takes the same action. This lasts until August 1987.
 March Second Therac-25 user group meeting.
 5 March AECL sends third revision of CAP to FDA.
 9 April FDA responds to CAP and asks for additional information.
 1 May AECL sends fourth revision of CAP to FDA.
 26 May FDA approves CAP subject to final testing and safety analysis.
 5 June AECL sends final test plan and draft safety analysis to FDA.
 July Third Therac-25 users group meeting.
 21 July Fifth and final version of CAP sent to FDA.

1988

29 January Interim safety analysis report issued.
 3 November Final safety analysis report issued.

Source: Leveson and Turner (1993).