

Learning from Incidents

DAVID L. COOKE¹

Haskayne School of Business, University of Calgary
Author's Address: 1008 Beverley Blvd. SW, Calgary, AB T2V 2C5
Tel. (403) 255-3878 Fax: (403) 255-0820 E-mail: cookedl@shaw.ca

Abstract

Many disasters have occurred because organizations have ignored the warning signs of pre-cursor incidents or have failed to learn from the lessons of the past. Risk is inherent in many high technology systems, but society views the benefits of continuing to operate these systems as outweighing the cost of the occasional disaster. Must we continue to live with disasters? Normal accident theory sees accidents as the unwanted but inevitable output of complex systems, while high reliability theory sees accidents as preventable by certain characteristics of the organization. This paper proposes that an incident learning system can provide a bridge between these two theories. By learning from the incidents that inevitably occur in a complex system, an organization can reduce risk and minimize loss. Thus, an organization with an effective incident learning system sustains a process of continuous improvement that allows it to become a high reliability organization over time. Incident learning theory suggests that implementing a system to encourage reporting of *more* incidents will drive a cycle of continuous organizational improvement that will reduce incident severity and reduce risk of disaster.

1. Introduction

On January 28, 1986 seven crew members of the *Challenger* were killed when the space shuttle exploded just over a minute after take-off. The Report of the Presidential Commission on the Space Shuttle *Challenger* Accident (1986) concluded that neither Thiokol, the seal designer, nor NASA “responded adequately to internal warnings about the faulty seal design. ...A well structured and managed system emphasizing safety would have flagged the rising doubts about the Solid Rocket Booster joint seal.” There had been nine prior incidents of O-ring seal failure.

On May 9, 1992 an explosion in the Westray mine at Plymouth, Nova Scotia killed 26 miners. There were many incidents leading up to the disaster that could have claimed lives but instead ended up as production losses or “near-misses.” Because of the many warning signs, Richard (1996) called Westray a “predictable path to disaster.”

In May 1996, ValuJet Flight 592 exploded and crashed into a Florida swamp killing all 110 people on board. Langewiesche (1998) reports that by early 1996 the US Federal Aviation

¹ David L. Cooke is a doctoral student in the Haskayne School of Business at the University of Calgary. Formerly, he was Director of Safety and Emergency Planning with NOVA Chemicals Corporation. He holds an MBA from the University of Calgary and a B.Sc. in chemical engineering from the University of Birmingham, England.

Authority was concerned “about the disproportionate number of infractions committed by ValuJet and the string of small bang-ups it had had.”

On June 22, 1997 at a Shell Chemical Company plant in Deer Park, Texas, the drive shaft blew out of a check valve causing the release of a large quantity of flammable gas. The resulting explosion and fire caused extensive damage and several workers suffered minor injuries. Fortunately, no one was killed. The EPA and OSHA (1998) investigation noted that there had been several prior incidents involving a similar mode of failure of this particular check valve at this and other Shell facilities, but the lessons learned from these prior incidents were not adequately identified, shared and implemented.

On Sept 11, 2001 terrorist attacks in New York, Pennsylvania and Washington killed over 3000 people. Coleen Rowley’s memo to FBI Director Robert Mueller and her testimony to the Senate Judiciary Committee suggest that the FBI failed to connect together the “various incidents and reports worldwide” and the Minneapolis and Phoenix Division warnings of “Al Qaeda operatives in flight schools seeking flight training for terrorist purposes.” FBI computers were “too antiquated” to allow anything more than one-word text searches so that “field agents could not have searched the FBI’s computer system for ‘airline training’ or ‘flight schools’ to determine if there was additional information on other terrorism suspects.” (Rowley (2002), Seper (2002)). There had been an earlier terrorist attack on the World Trade Center in 1993.

All of these disasters have one thing in common. That is the inability of the organization in question to effectively synthesize and share the information from separate “pre-cursor” incidents with the relevant people across the organization so that appropriate action could be taken to reduce the risk of disaster. In this paper I will use the term *incident* to mean an unexpected and unwanted event that represents a deviation from normal system behavior and which may or may not result in a loss. The commonly used term *accident* is an incident in which a non-trivial loss occurs, and a *disaster* is a very serious incident involving loss of life and/or extensive property damage.

None of the organizations in these examples had an effective capability to learn from the pre-cursor incidents. In other words, they lacked an effective incident learning system. Without such a system, the pre-cursor incidents are only visible with the benefit of hindsight. An *incident learning system* is the collection of organizational capabilities that enable the organization to extract useful information from incidents of all kinds and to use this information to improve organizational performance over time. In the context of the “learning organization” described by Senge (1990), it is just one of a number of possible management systems that enable the organization to learn, adapt and grow. Implementing an incident learning system is one way to operationalize and manage “organizational learning cycles” as conceived by Kim (1994).

There is no doubt that the inquiry reports into the aforementioned disasters make or will make recommendations to ensure that a similar disaster will never happen again. I doubt if there has ever been an inquiry into a disaster that did not make such recommendations. Yet

why *do* similar disasters happen again and again? Why, as Kletz (1993) points out, do organizations have no memory and accidents recur? This paper submits that some organizations fail to learn from pre-cursor incidents because they lack an effective incident learning system. Such a system serves as the organizational memory of incidents. Sadly, the *Columbia* space shuttle disaster occurred as this paper was being written, suggesting that NASA has not learned the lessons of the past. Indeed, preliminary information from the incident investigation suggests that a contributing factor to the *Columbia* disaster may have been damage to the thermal protection tiles during launch, which had certainly happened many times in the past. As in the *Challenger* disaster, the warning signals from precursor incidents were not dealt with effectively.

This paper explores the concept of an incident learning system, which is a broader concept than that of the “near-miss management system” described by Phimister et al (2001), and is more like the continuous improvement cycle described by Repenning and Sterman (2001). The purpose of the paper is to present a model incident learning system, to show how incident learning acts as a continuous improvement process to achieve organizational success, and to show how incident learning theory can provide a bridge between normal accident theory and high reliability theory. The model deals with long-term system behavior and does not deal with frequent interruptions or being overcome by events, as modeled by Rudolph and Repenning (2002). Furthermore, it is concerned only with overall organizational response and not with the response of individual agents.

The paper is organized as follows. The next section reviews the previous literature on organizational accident theory. Section 3 introduces the incident learning system model, examines each element of the system and shows how it behaves as a continuous improvement process. Section 4 discusses the dynamics of the incident learning system and explains the paradox of how many incidents can be better than few. Section 5 includes some suggestions for how some barriers to incident learning might be overcome in practice, and the final section makes some closing remarks and summarizes some propositions that can be drawn from the theory that has been presented.

2. Normal Accidents and High Reliability Organizations

It is proposed that incident learning system theory may provide a bridge between normal accident theory and high reliability theory, which are seemingly at odds with one another as discussed by Rijpma (1997). The theory of incident learning relies on the observation first made by Turner (1978), and updated by Turner and Pidgeon (1997), that disasters have long incubation periods during which warning signals (or incidents) are not detected or are ignored. This suggests that a system for incident learning could provide the antenna that an organization needs to detect these early warning signs and provide a process for putting preventative measures in place. Thus, while the occurrence of incidents may be normal, an organization with an effective incident learning system can respond to these incidents to prevent serious accidents from occurring in future. Through this process of continuous improvement, the organization may evolve into a high reliability organization over time.

The foundations of normal accident theory were laid by Perrow (1984) and consolidated by Sagan (1993). The theory holds that accidents are a normal consequence of the interactive complexity and close coupling of an organizational system. The measure of “interactive complexity” is the number of ways in which components of the system can interact. It represents the number of variables in the system, the number of relationships between the variables and the number of feedback loops through which the variables interact. Typically, interactive complexity increases with the technology incorporated into the system. The measure of “close coupling” is the speed at which a change in one variable cascades through the system to cause changes in other system variables. Close coupling represents tightness in the process, which is influenced by such things as component redundancy, resource buffers/slack, and process flexibility. The idea behind normal accident theory is that some of the system responses to change are unforeseen, are causes of incidents, and can potentially lead to catastrophes. Using the analogy of safety defenses being like slices of Swiss cheese (Reason, 1997), normal accident theory would say that no matter how high you stack the slices it is inevitable that organizational juggling will cause all the holes to line up eventually and the defenses will be breached.

High reliability theory is a competing organizational theory of accidents that parallels the “zero defects” school of thought in quality management. Its proponents such as La Porte and Consolini (1991), Weick and Sutcliffe (2001) and Roberts and Bea (2001) believe that while accidents may be normal, serious ones can be prevented by implementing certain organizational practices and processes. Roberts and Bea suggest that high reliability organizations “aggressively seek to know what they don’t know ... balance efficiency with reliability ... and communicate the big picture to everyone.” Weick and Sutcliffe suggest that high reliability organizations implement business processes to instill “mindfulness” qualities into the organization, which include preoccupation with failure, reluctance to simplify, sensitivity to operations, commitment to resilience, and deference to expertise.

Sagan (1993) distils high reliability theory down to four essential elements for success: high management priority on safety and reliability; redundancy and backup for people and equipment; decentralized organization with a strong culture and commitment to training; and organizational learning through trial and error, supported by anticipation and simulation. However, from the perspective of normal accident theory, he argues that there are many reasons why organizational learning will be restricted. These include ambiguity about incident causation, the politicized environments in which incident investigation takes place, the human tendency to cover up mistakes, and the secrecy both within and between competing organizations. Thus, Sagan’s view would be similar to that of Kletz, namely that accidents will recur despite organizational experience with similar incidents in the past.

Closely related to high reliability theory is the literature associated with the psychology of human error and accident prevention by creating a “safety culture” or “safety climate.” Pidgeon (1997) defines a safety culture as “the set of assumptions, and their associated practices, which permits beliefs about danger and safety to be constructed.” Reason (1997), Geller (2000), and Petersen (1996) provide advice and guidance to organizations striving to improve their safety culture. They advocate less emphasis on finding

engineering solutions to organizational accidents and moving more towards finding human/social solutions. Dorner (1996) uses dynamic system simulations or “micro-worlds” to show that when people make decisions with all the best intentions, disasters can happen if the decision-makers have poor knowledge and understanding of system behavior. Recognizing that human error does and will occur is the first step towards building a high reliability organization. In contrast, Perrow (1999) describes the achievement of a safety culture as being a sincere goal, but nevertheless a fantasy.

There have been many case studies but few empirical tests of these various theories. For example, Roberts (1990) describes some of the strategies employed by nuclear powered aircraft carrier organizations to avoid some of the antecedents to normal accidents found in case studies by Perrow (1984) and Shrivastava (1986). Vaughan (1996) describes in great detail the culture of production and the normalization of deviance in the NASA organization that led to the *Challenger* disaster.

Hofmann and Stetzer (1998) found that both the safety climate and safety communication significantly influenced the attribution of causes of industrial accidents. In a negative safety climate with poor safety communication, workers were less willing to attribute cause to the actions of a fellow worker.

Wolf (2001) carried out a test of normal accident theory by characterizing the interactive complexity of 36 oil refineries based on the number of unit processes, number of nodes connecting these processes, number of process parameters at each node, and the number of possible states for each of these parameters. His results supported normal accident theory in that the more complex refineries tended to have more frequent spills of hazardous materials. However, safety performance measured by Total Case Incident Rate (TCIR)² was not significantly correlated with interactive complexity, although those refineries classified as “more complex” that had high spills also had significantly higher TCIR. Although Wolf’s results provide some support for normal accident theory they do not necessarily reject the high reliability theory since no variable for “reliability” was included in the study. Furthermore, some of the less complex refineries reported no spills and provided no injury and illness data, which raises questions about the quality of reporting among the various refineries. A refinery with a poor incident learning system would tend to report fewer incidents even though many more unreported incidents may have occurred.

Both normal accident theory and high reliability theory have been criticized for taking an overly objectivist, positivist view of reality. Gephart (1984) argues that disasters are inherently political phenomena with many divergent views of reality. Focusing on two case studies of environmental disasters, he makes the disturbing assertion that “The occurrence of a disaster is a political accomplishment.” Gephart’s political sense-making view of accidents can be used by critics of both normal accident theory and high reliability theory. Critics of normal accident theory could argue that “unavoidable accidents” are political

² Defined by the US Occupational Safety and Health Administration, this is the number of recordable injuries and illnesses per 200,000 hours worked.

constructs of those not wanting to take responsibility for accident prevention. Critics of high reliability theory might say that the whole notion of a high reliability organization is a political construct created for the purpose of allaying public fears about risk. For example, critics would argue that the many near-misses in the nuclear weapons industry cited by Sagan (1993) and the many deaths from friendly fire and the near-misses reported by Thorne (2003) and Regan (2003) suggest that not all US military organizations are as reliable as the aircraft carrier studied by Roberts (1990).

3. Incident Learning System Concepts

Learning from incidents is not a new concept (see Bird and Germain, 1986 and Carroll, 1998, for example), but it has not been fully explored as a system for continuous improvement in organizational performance over time. Rudolph and Repenning (2002) describe a “disaster dynamics” model that provides insight into the role that a stream of events or frequent interruptions can play in causing disaster by “information overload,” but they were not concerned with incident learning. The time period for their dynamic simulation was minutes rather than the months and years involved in incident learning. However, their model does provide a very relevant warning that an incident learning system will collapse if it becomes overloaded with incidents.

I will use the term *system* to mean an interconnected set of related variables in which a change in one variable has the potential to affect all other variables in the system over time. The variables in business systems will be such things as people, resources, work processes, information, etc. A system will normally have a goal or purpose, even if the goal is simply to adapt to change (as in many biological systems) or to restore equilibrium (as in many physical or chemical systems), and may have a multitude of goals and purposes, some of which may conflict. For example, it is clear that a business system has several purposes including the transformation of input resources into the more valuable outputs representing the products and services of the business, but the goal of “low cost” may conflict with “employee satisfaction” and so compromises must be reached in the quest to optimize system performance.

To help understand why accidents happen, it is useful to introduce the concept of a *risk system*. This is the system that generates incidents. It is inseparable from the business system that generates the useful outputs of the organization, but we can gain valuable insights from thinking of them as distinct systems. Although incidents are actually unwanted outputs of the business system, it is instructive to view them as outputs of the risk system. The risk system may be hidden from view, but its outputs are real enough.

Just as we would apply quality management principles to control the quality of products and services from the business system, so we must apply similar principles to control the “quality” of incidents from the risk system. In fact, it would be equally valid to consider incidents to be “quality problems” or to consider quality problems to be “incidents.” The same principles of monitoring and control will apply. Organizations should employ an *incident learning system* to identify and analyze incidents so as to correct deficiencies in the

risk system in the same way as they employ a *quality management system* to deal with quality problems and improve the business system. Figure 1 shows how feedback from the quality management and incident learning systems improves business performance.




Figure 1: Continuous Improvement of Business Performance

Ideally, the two systems should be integrated under the umbrella of total quality management (TQM). The integration of safety and quality management has been proposed by several authors including Roughton (1993), Lischeid (1994), Petersen (1994), Curtis (1995), Karuppan et al (1996), Manzella (1997), Weinstein (1998), and Ahmed (2001). Integrating an incident learning system with quality management is one way to operationalize the “systemic quality management” model described by Kim (1994). This integrated view is shown in Figure 2, in which the quality management system focuses on *prevention* of incidents and the incident learning system provides for business system improvement based on incidents that occur.




Figure 2: An Integrated View of Total Quality Management Includes Safety

Thus, the incident learning system provides a risk control process for the business. Its components include identification and response, reporting, investigation, identifying causal structure, making recommendations, communicating and recalling incident learnings, and implementing corrective actions. Effective work processes for all of these components

must be in place for the system as a whole to operate well. It should also be evident that an incident learning system will operate most effectively when a safety management system has already been put in place and avoidable risks have been addressed. Implementation will be less effective in the absence of other safety and quality management systems.

The first component of the incident learning system is **identification and response**. Phimister et al (2001) discuss the importance of identification, without which incident learning is not possible. Unless the organization is sensitized to learning from incidents, deviations from normal behavior will go unnoticed or be accepted as “normal deviation” as at NASA. However, the *threshold* for detection is important or else the organization will be bombarded with a deluge of incidents which may exceed its capacity to investigate. Phimister et al do not include a response component in their near-miss management system, perhaps because a “miss” by their definition does not require an immediate response. However, even in the case of a near-miss, there should be an immediate response to correct any unsafe conditions resulting from the incident, to provide first-aid response in the case of a minor injury, or to clean up a small spill. The more comprehensive incident learning system proposed here must deal with incidents ranging in severity from near-misses to major accidents. Thus, response is required to first assess the situation to determine whether or not emergency response is required or to immediately implement corrective actions that will stabilize the situation. For example, a small fire could escalate into a major incident unless immediate action is taken to extinguish it. Only when the incident is properly controlled and the situation is stable should data be gathered about the incident.

The next component of incident learning is **reporting**. As the Center for Chemical Process Safety (1989) and other references on incident investigation point out, an incident cannot be investigated unless it is reported. Furthermore, the fraction of incidents reported is dependent on the personal commitment to safety of the workers who observe or are involved in the incidents. As discussed in Cooke (2003), management creates the safety climate and so personal commitment to safety of the workers is strongly influenced by management’s commitment to safety. Management can show their commitment to safety by creating a climate in which incident reporting is rewarded instead of punished. Part of the “reward” should be to make the reporting process as easy as possible and to include the reporter in the investigation process if he/she so desires. The local manager should communicate the initial incident report according to pre-determined guidelines that relate the nature and severity of the incident to the appropriate distribution list. Pre-determined guidelines can also be used to describe the *nature* of the incident (safety/health, environment, production/quality losses, and other types of risks) and its *severity* (e.g. near-miss, minor incident, serious incident, major incident, etc).

Incident **investigation** is the most well known component of the incident learning system involving examination of the site, interviewing witnesses, gathering and evaluating all available data to establish the sequence of events and determine exactly what happened. An investigation *team* will be more effective than a single investigator. Detailed elements of the incident investigation process can be found in sources such as Bird and Germain (1986), Center for Chemical Process Safety (1989), or National Safety Council (1995).

Many sources say that the purpose of incident investigation is to determine the basic or root causes of the incident. However, since there may be no single “root cause,” efforts are better directed towards **identifying causal structure**. This should be viewed as a separate process step, which reduces the investigation team’s temptation to leap to a conclusion before all relevant data has been gathered and evaluated. Most safety text books refer to determining the one or two root or basic causes of the incident, and indeed the widely-used International Loss Control Institute approach (Bird and Germain, 1986) reduces all possible root causes to a comprehensive look-up table. However, this linear reductionistic approach may not work for complex system accidents in which there may be a multiplicity of “root causes.” The desire to find a single root cause was observed by Carroll (1995), who called it “root cause seduction.” Carroll (1998) describes the various techniques of root cause analysis he observed at two companies, and again calls for a different approach to that of finding a single root cause. For complex incidents, one possible approach is to supplement the techniques of root cause analysis with systems thinking. Causal loop diagrams, which would capture the causal structure of the incident in a network of interlinked causes and effects, would be an appropriate way to represent complex causation. Incident investigation teams could be trained so that they could construct a system model of the causal relationships. Cooke (2002) provides a case study which can be used to teach the causal modeling approach for analyzing incidents in complex systems.

The work of the investigation team is usually completed by issuance of an incident report detailing the findings and recommendations. Distribution of this report is one way to **communicate incident learnings**. Another way is to capture the important learnings in an abstract or executive summary which can be distributed electronically within the company and, if appropriate, externally to members of the industry association. The distribution list can vary depending on the severity of the incident. In the chemical industry, the Center for Chemical Process Safety has played an important role in disseminating learnings from member company incidents. Although I agree with Perrow (1999) that “decentralized units are better able to handle the continual stream of small failures, forestalling the widespread multiple failures,” the lesson to be learned from the Shell Deer Park disaster and the 9/11 terrorist attacks is that an effective communication mechanism is needed to synthesize the information from the many small failures into organizational knowledge that can prevent a much larger failure.

Part of the learning process is to **recall previous incidents** and to visualize possible failure modes that have not yet occurred, but which previous incidents have suggested might be possible. Bird and Germain (1986) provide details of the method. Incident recall and visualization can be done individually, through an interview process or in groups. Group processes are particularly valuable for stimulating ideas and reinforcing learning.

Of course it is important to **implement corrective actions** and follow up on all recommendations made by the investigation team. Effective action and follow-up to completion of all outstanding recommendations will test the quality of management’s planning and control systems. This is particularly true for actions to eliminate systemic

causes of incidents, which may span the organization and involve many people in different geographic locations. In the modern information age, computer databases can be used effectively for tracking recommendations and following them up to completion. Processes outside of the incident learning system, such as audits and inspections, are also useful in checking that corrective actions have been successfully implemented. Completion of safety-related improvements can also be built into the management and employee compensation systems.

Figure 3 summarizes the above discussion by showing how the components of the incident learning system complete the continuous improvement cycle for the business system. The diagram also shows the important external link by which learnings at the local level are shared with other locations and businesses within the same company, and with other companies through industry associations.




Figure 3: The Incident Learning System Closes the Feedback Loop

In Figure 3, it is easy to see that if any arrow is missing then the feedback loop is broken and learning from that incident does not occur. The arrow does not even have to be missing for learning not to occur. It may be that the quality of information passed from one stage of the process to the next is poor or diminishing. Considering this effect over hundreds of incidents leads to an aggregate model in which each “arrow” has a weight equal to the percentage of information flowing back to the business system. This is similar to the concept of “stage efficiencies” discussed by Phimister et al (2001).

We may also hypothesize that if the cycle shown in Figure 3 does not keep turning then the organization may get complacent or suffer from over-confidence in the safety of its operations. To keep the circle turning smoothly there must also be a balance between the number of incidents being reported and the organization’s capacity to investigate and learn from the incidents. These considerations suggest it is important to establish an appropriate threshold for incident reporting and assign appropriate resources for incident investigation.

This completes the discussion of incident learning system concepts. The next section will expand this discussion by describing how the incident learning system dynamically interacts with the business system.

4. Incident Learning System Dynamics

The Westray mine disaster will be used as a case study to illustrate the dynamics of the incident learning system. Westray was a “typical” mine disaster in which 26 miners were killed in an explosion below ground on May 9, 1992, following a number of serious incidents each of which could have led to disaster but instead translated into production loss until the “big one” arrived. The organizational dynamics of the Westray mine system have been described by Hynes and Prasad (1997) and Cooke (2003). At the Westray mine, a disaster was inevitable despite it being a linear, loosely-coupled system, which normal accident theory would suggest should not be prone to disaster. In theory, the behavior of a linear system is more easily predictable and loose-coupling means that slack in resources should allow for system recovery. However, in a system such as Westray, where there is total disregard for safety in the face of production pressure, incidents will occur at an increasing rate until disaster strikes. To prevent a disaster, management must impose on the system some other limits to growth in the incident rate. One such limit, discussed by Cooke (2003), would be a management policy under which safety would take priority over production. To prevent such a policy from becoming just another “safety first” slogan, management actions and behaviors must continually reinforce a safety culture by demonstrating management commitment to safety. Geller (1996), Krause (1996), and Petersen (1996) have written extensively about what it takes to create a safety culture in the workplace. Another way to improve system performance, and indeed to help create a safety culture, is to implement an incident learning system as discussed in the previous section. Here, we will propose a model for the dynamics of the incident learning system and show how this model modifies the behavior of the Westray system.

If an organization counters the culture of production by successfully implementing methods and processes to create a safety climate in the workplace, then the linkage between safety and production will be broken, or at least severely weakened, and the incident rate will be controlled over time as was shown by Cooke (2003). This means that the fewer the incidents, the lower the risk of disaster. But how can we reconcile such a conclusion with the concept of learning from incidents? If we have more incidents, don't we learn more? The solution to this paradox is to recognize that every incident has a different *severity*, or magnitude of loss. Thus, the loss arising from incidents is a product of two random variables, the number of incidents and the severity of each incident. If N is a random variable representing the number of incidents in a given period, and S_i is a random variable representing the severity of an incident i , then it can be shown (Ross, 2000) that:

$$\text{Expected loss in the period} = E \left[\sum_{i=1}^N S_i \right] = E[N]E[S]$$

$$\text{Variance of loss in the period} = \text{Var}\left(\sum_{i=1}^N S_i\right) = E[N]\text{Var}(S) + (E[S])^2 \text{Var}(N)$$

From this we conclude that an objective of minimizing losses can be achieved by minimizing *severity* per incident and not necessarily by minimizing the number of incidents. If the expected severity is close to zero, then the expected loss will be close to zero no matter how many incidents there are. Thus, the objective of the incident learning system is not to minimize the number of incidents but to minimize the severity of the incidents. Indeed, in a study of nursing units, Edmondson (1996) found that higher performing units reported *more* errors than did lower performing units.

In the previous section we discussed the importance of encouraging incident reporting. As safety performance improves, reporting can be maintained by lowering the severity threshold for reporting to include incidents for which the actual loss is zero (so-called “near-misses” or “near-hits”). A parallel can be drawn with quality control in terms of where you set the control limits. Indeed Gothard and Wixson (1994) describe how Buffalo General Hospital integrate control charting with incident investigations to eliminate special cause variation. Consistent with the previous discussion about incident severity, they found that implementing an incident reporting system resulted in a 20% increase in incidents being reported without any change in actual loss experience.

The severity of an incident is often measured in financial terms because cost/benefit analysis is a common approach to management decision-making and financial measures allow losses from many different types of incidents (safety, production, quality, health, environmental etc) to be aggregated on the same basis. This cost analysis approach is advocated by leading safety bodies such as the National Safety Council. However, other appropriate units of measure for severity can be used. For example, a manufacturing facility might draw attention to losses by expressing them in terms of the amount of production needed to generate an equivalent financial contribution.

A general system dynamics model for an incident learning system is shown in Figure 4. This model is consistent with that shown in Figure 3 except that “identification and response” are not modeled separately from “reporting.” The model makes several simplifying assumptions, for example it assumes no restriction on resource availability for incident investigations, and no effects associated with size of organization. To illustrate its use and examine its behavior, this general model was applied to the particular case of the Westray mine model described by Cooke (2003). Variable names are referred to in the text in *Headline Style Italics* and the variables calculated by the Westray model are shown in Figure 4 in **bold**. A listing of the incident learning system model equations is given in the appendix and details of the equations for *Incident Rate* and other Westray model variables are given in Cooke (2003).




Figure 4: Incident Learning System Dynamic Model

The generic model shown in Figure 4 operates as a first order negative feedback system. The system response (based on the equations listed in the appendix) is shown in Figure 5. This behavior adequately represents the desired reference mode behavior, whereby the initial value of *Lessons Learned* declines exponentially as a result of *Organizational Memory Loss* with no new learning taking place because the incident rate is zero until week 50. The variable *Severity of Incidents* rises exponentially until week 50, which could be interpreted as representing the potential for disaster as there are no incidents in this period. After week 50, an *Incident Rate* of 1 per week is sufficient to restore *Lessons Learned* and reduce *Severity of Incidents* to a low level.




Figure 5: Incident Learning System Response to a Step Change in Incident Rate

The model will now be described in more detail. *Lessons Learned* accumulates from *Incident Learning*, and these lessons can be conceptualized as keeping the memory of incidents alive as long as *Incident Learning* is equal to or greater than *Organizational Memory Loss*. The more *Lessons Learned*, the more *Corrective Actions* are taken, and the lower the *Severity of Incidents*. These relationships follow from Kletz (1993) observation that accidents recur in the absence of organizational memory. Kletz also says that preventing serious incidents is often not a difficult or technically insurmountable problem but rather “they occur because we do not use the knowledge we already have.” Through the variables *Fraction of Incidents to Which Learning is Applied* and *Effectiveness of Corrective Actions*, the model captures the idea that knowledge must be effectively applied as well as acquired.

The structure of each equation in the model is chosen to be as simple as possible whilst still reflecting reference mode behavior. For example, we expect that the *Severity of Incidents* will decline the more *Corrective Actions* are taken, and model this by a simple inverse relationship. We also expect organizations to take time to implement *Corrective Actions*, so we use a smooth function to create a first-order delay between *Corrective Actions* and *Severity of Incidents*. Although we might expect *Corrective Actions* to reduce both the *Severity of Incidents* and the *Incident Rate*, and certainly this is true indirectly, we make the simplifying assumption that *Corrective Actions* have a direct effect only on severity. *Corrective Actions* have an indirect effect on *Incident Rate* through *Losses* and their effect on the *Commitment to Safety* variables, as shown in Figure 6.

To adapt the generic model in Figure 4 for use in the Westray model, *Losses* were expressed in terms of loss of mine capacity and the severity of each incident was interpreted as the fractional capacity loss per incident. In the original Westray model the *Fractional Losses from Incidents* was set at a constant value. In the modified Westray model with incident learning, the *Fractional Losses from Incidents* becomes a variable that is equal to *Severity of Incidents* multiplied by the *Incident Rate*. Thus, for the same incident rate, higher incident severity means higher losses.

Consistent with the stage efficiency concepts of Phimister et al (2001), the model contains some “efficiency-type” variables such as the *Fraction of Incidents Reported*. Some of these efficiency-type variables will be dynamic. For example, the *Fraction of Incidents Reported* will depend on the workers’ *Personal Commitment to Safety* and the *Fraction of Incidents to Which Learning is Applied* will depend on *Management Commitment to Safety*.

An important driver of success in an incident learning system will be the commitment to safety of the people working within the system. If the organization does not adequately respond to the pressure that a widening *Performance Gap* creates, then the wheels of the system will not turn. It is the worker’s *Personal Commitment to Safety* that causes more incident reporting, which leads to higher *Quality of Incident Investigations* and more *Incident Learning*, both individually and organizationally. Similarly, it is *Management Commitment to Safety* that causes *Corrective Actions* to be taken. The “commitment to

safety” concept is a central component of the Westray model and the interactions between this commitment to safety model and the incident learning model are shown in a causal loop diagram in Figure 6. Referring to Figure 6, a higher quality of information coming from the incident learning system increases *Management Commitment to Safety*, causing managers to act proactively to address unsafe conditions, and as management is seen to “walk the talk” (do what they say they are going to do) there is a cascade effect onto *Personal Commitment to Safety* causing people to work more safely, engage in less risky behaviors, and report more incidents.




Figure 6: Interaction between Incident Learning and Commitment to Safety

The equation for the variable *Quality of Incident Investigations* is the most complicated equation in the model. It serves two purposes. The first is to recognize that higher quality of information is generated if more incidents are reported and if employees participate in the incident investigations. This is modeled by the *Quality of Incident Investigations* being proportional to the *Fraction of Incidents Reported* and to *Employee Participation*. The second purpose is to translate the “pressure” created by the Performance Gap into resources that will sustain a high quality investigation process. This is modeled by an exponential function with a single parameter called *Organizational Response to Losses*. Based upon experience, we would expect that *Quality of Incident Investigations* would be an increasing concave-down function of the *Performance Gap*, but otherwise the shape of this function is not known. The chosen function captures this behavior and possible trajectories of this function, depending on the value of the *Organizational Response to Losses* parameter, are shown in Figure 7.

From the way that *Performance Gap* is defined, it can be seen that *Losses* would be equal to *Acceptable Losses* when *Performance Gap* = 1. Thus, the larger the *Organizational Response to Losses* parameter the more closely the *Performance Gap* will approach 1. There will always be an “offset” (i.e. *Performance Gap* > 1) to drive the system and maintain learning.




Figure 7: Possible Effects of Performance Gap on Quality of Incident Investigations

The next set of simulation results show the effect of adding the incident learning system model to the Westray mine model. Figure 8 summarizes the behavior of the Westray mine system with high losses from incidents, which were originally reported in Cooke (2003).



Figure 8: Westray System Response without Incident Learning

Without incident learning, as seen in Figure 8, *Management Commitment to Safety* declines as priority is diverted to maintain production. The *Incident Rate* accelerates as the effectiveness of safety management declines. Although *Losses* fall initially because of the initial commitment to safety, this only serves to “fool” management into thinking the situation is improving. Incidents accelerate and commitment to safety declines in a vicious circle until an explosion becomes inevitable.

Figure 9 shows that the Westray system response is quite different when the incident learning system model is incorporated. *Losses* fall sharply, and stay at a fraction of their starting value, while the *Incident Rate* falls and *Management Commitment to Safety* increases.



Figure 9: Westray System Response with Incident Learning

The results in Figure 9 were generated with the parameter values set as shown in the appendix. How sensitive are the results to variation of incident learning system parameters? Overall, the results show that Westray model performance is always better with an incident learning system, even if parameter values are set to create a weak system. In other words, a weak incident learning system is better than no system at all. To answer this question we can vary either the *Employee Participation* parameter, which changes the strength of the feedback from the Westray model to the incident learning system, or we can vary the *Organizational Response to Losses* parameter, which changes the strength of the feedback around the incident learning system loop. The results for varying the *Employee Participation* parameter are shown in Figure 10. Varying the *Employee Participation* parameter from 0.1 to 1 and varying the *Organizational Response to Losses* parameter from 0.03 to 1 give a similar response, and so these results can be interpreted more generally as the sensitivity the system’s response to the strength and quality of the organization’s ability to learn from incidents.

In summary, the stronger the organizational response to unacceptable losses the more incidents are reported, the more corrective actions taken and the lower the severity of the incidents. The nature of the organizational response to incidents may explain why incident learning is successful in some organizations, which may go on to become “high reliability organizations,” and not in others, which go on to become organizations in which recurring accidents are accepted as normal.




Figure 10: Sensitivity of Westray System Response to Changes in the Strength of the Incident Learning System (*Employee Participation* parameter range 0.1 to 1)

A strong organizational response to incidents should not be interpreted as management making heavy-handed safety prescriptions or empty exhortations to put safety first. Pidgeon (1997) is right to raise a red flag about the politics of safety, and the risk that “safety culture” may become merely rhetoric that management hides behind. A strong organizational response means putting a high priority on learning from incidents, auditing and continuously improving the incident learning system and the other management systems that regulate quality, safety, health and environmental performance. Learning from incidents is work in progress that will never be complete. It is this striving for continuous improvement in the various management systems, not talking about safety culture, which will create a “high reliability” organization over time.

5. Overcoming Barriers to Incident Learning

The models of the incident learning system shown in Figures 3 and 4 could be made much more complicated without adding a lot of value. Their strengths and weaknesses could also be critiqued, as could any of the assumptions on which these models are based. However, the true value of these models is to highlight the incident learning system as a continuous improvement process, to generate discussion and to suggest ways in which the learning process can be strengthened and improved. The following suggestions for strengthening the incident learning system and may go some way towards addressing the barriers to organizational learning identified by Sagan (1993), Rijpma (1997) and Pidgeon (1997):

- Management should not be discouraged by those who dismiss the creation of a safety culture as a myth or fantasy. Simulations by Cooke (2003) show that it can take several years to build a safety culture. Implementation of an incident learning system in which people are dealt with fairly, safety is openly discussed, and corrective actions are implemented in a cross-functional team environment will go a long way towards demonstrating that management can “walk the talk.”
- Organizations should put their focus on determining the causal structure of the incident in the context of the overall business system, rather than on finding the “root cause.” Obviously the extent of this analysis will depend on the severity of the incident. Incident investigation teams should be trained and equipped with the tools of system thinking as set forth by Sterman (2000), using case studies like that provided by Cooke (2002). Ragan and Carder (1994) provide some guidance on the systems approach to safety. A systems approach will help to combat confusion created by an ambiguity of causes and reduce the risk of improper cause attribution.
- Organizations should implement a reward system that encourages reporting of incidents and implementation of corrective actions. These are the two steps in the process that open and close the incident learning cycle respectively and so they need to be done well for business system improvement to occur. Compensation systems that reward specific safety targets such as “zero spills” or “no lost time accidents” will only serve to discourage reporting and drive incidents underground. No one wants to be the worker who is injured on the job in the week when his colleagues were expecting to get a bonus for one million man-hours without a recordable injury. Eliminating the blame game is

difficult, but it can be done by following the fourteen steps laid out for quality improvement by Deming (1989) and adapting them to safety. Viewing incidents as learning opportunities is a management approach that is similar to the “just-in-time” operations management philosophy of operating with lower inventory so as to uncover operational problems, which can be fixed once they have been revealed (Lieberman and Demeester, 1999).

- The importance of an incident learning system for strengthening risk communications cannot be over-emphasized. Managers and supervisors should discuss learnings from incidents at every opportunity. For example, the first item of business on the agenda of a weekly maintenance team meeting could be to review the learnings from incidents reported in the previous week. Conversely, communicating learnings from incidents both internally and externally will validate and strengthen the incident learning system itself. Grabowski and Roberts (1997) discuss the importance of communication processes for the reduction of risk in large-scale (multi-organizational) systems. Although they don’t specifically mention learning from incidents, such communications are exactly the kind of risk-mitigating communications that are needed for reducing risk in large-scale systems. Incident learning systems operating across industry sectors have proven possible as long as the contributors are assured anonymity and freedom from prosecution, which can be difficult in some legal environments. Nevertheless, examples of successful industry incident sharing mechanisms can be found in the airline, chemical, and nuclear industries.




Figure 11: External Benchmarking Closes another Feedback Loop

- The external loop for shared learning shown in Figure 3 should be closed by a benchmarking process that analyses best practices and adapts these practices to improve the business system. This concept is shown in Figure 11. Sharing the learnings externally encourages other organizations not only to learn from the incident itself but

to share what they have learned with the original organization at which the incident occurred. Other organizations may contact the original organization at which the incident occurred to say, “We have experienced that particular problem and we are willing to share our solution with you.”

- Organizations should track and aggregate losses from incidents and include this metric among their key performance indicators. Reduction in losses provides the economic justification for investment in the incident learning system. However, management should be careful in how loss information is communicated. Although communicating and giving visible status to loss metrics will help to raise awareness of losses, there is a risk that it will discourage reporting if people perceive the organization’s objective is to simply reduce loss. Therefore, the communication should be structured so as to reinforce the importance of incident learning. The communication goal should be to celebrate the number of incidents reported and success in completing corrective actions.
- Organizations should maximize employee participation in the incident learning system to improve learning and reduce risk of complacency. As discussed by Gonzalez and Sawicka (2003), learning and risk perception play an important role in compliance with safety and security procedures. Employee participation in the incident learning process will not only improve the effectiveness of the incident learning system but will also improve the participant’s perception of workplace risks by challenging their existing mental models of safety. The dynamic aspects of risk perception have been explored by Sawicka and Gonzalez (2003), who show how an “out of sight, out of mind” mentality can increase the risk of disaster. An incident learning system can mitigate this “deconditioning” process by creating a higher awareness of risk. Employee involvement in incident investigations, and other interactions with the incident learning system, helps to keep the risks “in sight and in mind.”
- Management should use climate surveys and other feedback tools to measure organizational commitment to incident learning. The survey and feedback information can be useful in designing and implementing policies and procedures to encourage a proactive learning response to increased losses. It could also be used by the CEO and Board of Directors to monitor the “safety pulse” of the organization.
- Since an organization may experience thousands of low severity incidents a year, there must be an easy-to-use database for capturing the lessons learned. Corporate intranets are increasingly being used for storing and sharing this information. However, storage in databases is not enough. There must also be a management of change process for upgrading operating procedures, maintenance procedures, engineering standards etc based upon learning from incidents. For further information on management of change see Center for Chemical Process Safety (1989).
- Finally, many organizations will find opportunities to integrate quality, safety, environment, health and other risks into a single comprehensive incident learning system. A company that has already implemented a quality management system can extend the definition of quality to include deviations from safety, health and environmental norms. Similarly, a company with a strong operational risk management

system that wishes to implement total quality management can introduce quality management concepts into the overall risk management framework.

In summary, implementing an incident learning system equips an organization with a management process for operational improvement, but this process faces a number of barriers to effective implementation and ultimate success. However, these barriers should not discourage organizations from implementing an incident learning system and improving it over time. In particular, organizations should be wary of falling into what Repenning and Sterman (2001) call a *capability trap*, in which organizations fail to allocate sufficient resources to process improvement and then have to work harder and harder to sustain even the same level of performance in the face of declining business system capability. Further research is needed to provide economic models and sound business cases for investment in incident learning systems.

6. Concluding Remarks

This paper suggests that an incident learning system can bridge the gap between normal accident theory and high reliability theory. Although accidents may be “normal,” disaster is not an inevitable consequence of complex socio-technical systems. Since incidents of varying severity *are* normal, a system must be put in place to control the severity of these incidents. Without such a system the incident rate and severity will not be controlled and only *then* is a disaster predictable. This conclusion rests on the assumption that a disaster is not a spontaneous event, but a consequence of a chain of events or pre-cursor incidents that could have been detected by an effective incident learning system, thereby breaking the chain. Given the importance of disaster-prevention to the continuing survival of a socio-technical organization, this paper argues that an incident learning system should be just as central to the organization’s mission as the production or service delivery system that constitutes the main focus of the organization. Besides operationalizing a process for organizational learning from incidents, an incident learning system also helps to reduce the risk of organizational complacency or what Melara et al (2003) call “victims of their own success.”

There is little or no evidence of effective incident learning systems being in operation at any of the organizations experiencing the disasters cited in the introduction to this paper. For example, if NASA had learned from the previous O-ring failure incidents then a decision rule to not launch in low ambient temperature conditions would have reduced the risk of failure. Unfortunately, the absence of something does not demonstrate that its presence would make a difference. The fact that there is little or no empirical evidence in the literature showing whether or not an incident learning system makes a difference provides ample scope for further research. In particular, more industry studies like that of Wolf (2001) should be conducted but the measurement of organizational characteristics should include not only complexity and coupling but also the effectiveness of the organizations’ incident learning systems. Complexity moderated by effectiveness of the incident learning system should be a better predictor of safety performance than complexity alone.

To support this argument, consider the incident learning systems operated by the airlines and the various airline regulatory authorities. Although “normal accidents” still occur, imagine the carnage that would result if *no* incident learning took place. Similar comments apply to the contribution of incident learning systems to safety in the petrochemical and nuclear industries. Perhaps the focus on normal accidents in complex, tightly-coupled “high risk” systems such as petrochemicals and nuclear power has obscured the fact that their safety performance per hour of exposure is often better than that of linear, loosely-coupled, “low risk” systems such as construction or agriculture. Certainly, one is more likely to find an effective incident learning system at a petrochemical plant than at a farm or construction site.

If this theory of incident learning systems is valid, then we would expect to find that the following propositions will hold true:

1. Given the same degree of socio-technical interactive complexity and coupling, organizations having more effective incident learning systems should have better safety, quality, environmental and economic performance.
2. All socio-technical organizations will have incidents or “normal” accidents, but the presence of an effective incident learning system will mitigate the risk of disaster and lead to performance that may be interpreted as “high reliability.”
3. No organization can claim to be “high reliability” unless it can demonstrate that a large number of incidents are reported and dealt with through the organization’s incident learning system.
4. In an organization with an effective incident learning system, the number of incidents reported may increase initially but the average severity of the incidents reported will drop over time.

There are many obstacles to implementation of an effective incident learning system in an organization. Not least of which is the political/legal climate which seeks to apportion blame and file lawsuits when major incidents occur. This climate may lead to in-house counsel discouraging senior management from implementing an incident learning system because of the many “smoking guns” that incident reports could represent. This political/legal climate needs to change for long-term systemic improvement in accident prevention to occur. The change needed is for society to recognize that an effective incident learning system should be present in any organization that undertakes risky activities. Then if a disaster does strike, one focus of inquiry should be to assess the effectiveness of the incident learning system that was in place at the time. This assessment can be used as an indication of safety commitment and due diligence on the part of management. Organizations not having an effective incident learning system in place should be dealt with more harshly by society than those that do.

One last point to emphasize is that although the incident learning system has been the subject of this paper, it is only one of several management systems that organizations should put in place for loss control. Other systems include planned inspections, job/task

analysis, management of change, emergency preparedness and response, and quality improvement programs.

Appendix: Incident Learning System Model Equations

Acceptable Losses = 300

~ Tonnes/(Week*Week)

~ *This parameter is set to a value that is reasonable for the Westray mine system (see Cooke (2003) for a full description of this model).*

Capacity = 5500

~ Tonnes/Week

~ *This parameter is actually calculated by the Westray model. For the stand-alone test of the incident learning system capacity is set to a constant value.*

Corrective Actions = Lessons Learned * Fraction of Incidents to Which Learning is Applied

~ Learning

~ *It is not enough to learn lessons from incidents. These lessons must be applied in the field to correct unsafe conditions and address risky behavior.*

Effectiveness of Corrective Actions = 1

~ Learning/Incident

~ *The model assumes that the corrective actions are 100% effective.*

Employee Participation = 1

~ Learning/Incident

~ *This parameter captures the idea that the quality of incident investigations will be poor unless employee participation is high.*

Fraction of Incidents Reported = $\min(\text{Personal Commitment to Safety}/100, 1)$

~ Dimensionless

~ *All incidents are reported when Personal Commitment is 100% or more.*

Personal Commitment to Safety is calculated by the Westray model, which was modified so that Personal Commitment responds to losses and not simply to just the incident rate.

Fraction of Incidents to Which Learning is Applied = $\min(\text{Management Commitment to Safety}/100, 1)$

~ Dimensionless

~ *This variable captures the idea that learning is applied more effectively, to correct problems revealed by incidents, if management commitment to safety is higher.*

Management Commitment to Safety is calculated by the Westray model, which was

modified so that Management Commitment responds to losses and not simply to just the incident rate. The fraction cannot be greater than 1.

Fraction of Lessons Forgotten = 0.1

~ 1/Week

~ *This parameter captures Kletz's observation that some organizations have no memory and accidents recur. For such organizations, the value of this parameter would be high.*

Fractional Losses from Incidents = $\min(\text{Incident Rate} * \text{Severity of Incidents}, 1)$

~ 1/Week

~ *By definition, losses are equal to the incident rate multiplied by the severity of the incidents. Incident Rate is calculated by the Westray model.*

Incident Investigation Time = 6

~ Weeks

~ *This parameter represents the time taken to investigate an incident. Implicit in the model is the assumption that there are no resource constraints for incident investigations.*

Incident Learning = $\text{smooth}(\text{Incident Rate} * \text{Quality of Incident Investigations}, \text{Incident Investigation Time})$

~ Learning/Week

~ *It is assumed that incident learning is simply proportional to the number of incidents and the quality of the incident investigations.*

Incident Rate = $\text{STEP}(1, 50)$

~ Incident/Week

~ *The step function was used to test the stand-alone incident learning system (ILS) model. When the ILS model was integrated with Westray, then the Incident Rate is calculated by the Westray model.*

Lessons Learned = $\text{INTEG}(\text{Incident Learning} - \text{Organizational Memory Loss}, 10)$

~ Learning

~ *This variable represents organizational learning, which accumulates from incident learning and is depleted by organizational memory loss. The stock of Lessons Learned is initialized with an arbitrary value of 10 "learning units."*

Losses = $\text{Fractional Losses from Incidents} * \text{Capacity}$

~ Tonnes/(Week*Week)

~ *In a generic incident learning system, losses are measured in whatever units are important to the business system. In the Westray mine model, losses are measured in terms of lost production capacity. Capacity is calculated by the Westray model.*

Organizational Memory Loss = $\text{Lessons Learned} * \text{Fraction of Lessons Forgotten}$

- ~ Learning/Week
- ~ *Organization memory loss is assumed to be a fraction of the lessons learned.*

Organizational Response to Losses = 1

- ~ Dimensionless
- ~ *This parameter changes the shape of the function that models how the Quality of Incident Investigations variable responds to organizational efforts to reduce losses.*

Performance Gap = Losses / Acceptable Losses

- ~ Dimensionless
- ~ *This variable normalizes Losses so that a Performance Gap of 1 represents an acceptable level of loss.*

Quality of Incident Investigations = (1 - exp(-Performance Gap * Organizational Response to Losses)) * Fraction of Incidents Reported * Employee Participation

- ~ Learning/Incident
- ~ *The function in brackets represents the organizational response to losses. See Figure 6 for a picture of how this function affects the Quality of Incident Investigations, which is also assumed to be proportional to the Fraction of Incidents Reported and the degree of Employee Participation.*

Severity of Incidents = smooth(1/(Corrective Actions * Effectiveness of Corrective Actions), Time to Implement Corrective Actions)

- ~ 1/Incident
- ~ *The severity of incidents is assumed to be inversely proportional to the corrective actions taken and to the effectiveness of those actions.*

Time to Implement Corrective Actions = 16

- ~ Weeks
- ~ *The time to implement corrective actions is assumed to be constant in the model, but of course will vary with the backlog in real life.*

References

- Ahmed, N.U. 2001. Incorporating environmental concerns into TQM. *Production and Inventory Management Journal* (1), 25-30.
- Bird, F.E. and Germain, G.L. 1986. *Practical Loss Control Leadership*, Loganville, GA: Institute Publishing, a Division of the International Loss Control Institute.
- Carroll, J.S. 1995. Incident reviews in high-hazard industries: sensemaking and learning

- under ambiguity and accountability. *Industrial and Environmental Crisis Quarterly* 9, 175-197.
- Carroll, J.S. 1998. Organizational learning activities in high hazard industries: The logics underlying self-analysis. *Journal of Management Studies* 35 (6), 699-717.
- Center for Chemical Process Safety 1989. *Guidelines for Technical Management of Chemical Process Safety*. Center for Chemical Process Safety of the American Institute of Chemical Engineers, New York.
- Cooke, D.L. 2002. Westray, a tragedy of errors. Unpublished Case Study
- Cooke, D.L. 2003. A system dynamics analysis of the Westray mine disaster. *System Dynamics Review* 19 (2)
- Curtis, S.L. 1995. Safety and total quality management. *Professional Safety* (January), 18-20.
- Deming, W.E. 1989. *Out of the Crisis*, Ninth Printing ed. Cambridge, Mass.: Massachusetts Institute of Technology, Center for Advanced Engineering Study.
- Dorner, D. 1996. *The Logic of Failure: Recognizing and Avoiding Error in Complex Situations*, Cambridge, MA: Perseus Books.
- Edmondson, A.C. 1996. Learning from mistakes is easier said than done: Group and organizational influences in the detection and correction of human error. *Journal of Applied Behavioral Science* 32, 5-28.
- EPA and OSHA (1998) EPA/OSHA Joint Chemical Accident Investigation Report - Shell Chemical Company, Deer Park, Texas. EPA 550-R-98-005, United States Environmental Protection Agency.
- Geller, E.S. 1996. *The Psychology of Safety : How to Improve Behaviors and Attitudes on the Job*.
- Geller, E.S. 2000. *The Psychology of Safety Handbook*. 2nd ed., Lewis Publishers Inc.,
- Gephart, R. 1984. Making sense of organizationally based environmental disasters. *Journal of Management* 10 (2), 205-225.
- Gonzalez, Jose J. and Sawicka, Agata 2003. The role of learning and risk perception in compliance. 21st International Conference of the System Dynamics Society, New York.
- Gothard, L. and Wixson, N.E. 1994. Charting a course for continuous quality improvement. *Risk Management* January, 27-33.

- Grabowski, M. and Roberts, K.H. 1997. Risk mitigation in large-scale systems: lessons from high reliability organizations. *California Management Review* 39 (4), 152-162.
- Hofmann, D.A. and Stetzer, A. 1998. The role of safety climate and communication in accident interpretation: implications for learning from negative events. *Academy of Management Journal* 41 (6), 644-657.
- Hynes, T. and Prasad, P. 1997. Patterns of 'mock bureaucracy' in mining disasters: an analysis of the Westray coal mine explosion. *Journal of Management Studies* 34 (4), 601-623.
- Karuppan, C.M., Karuppan, M. and Schaefer, I.M. 1996. Total safety management: a TQM approach to industrial safety. *Production and Inventory Management Journal* (2), 15-20.
- Kim, D.H. 1994. Managing Organizational Learning Cycles. In: Wardman, K., (Ed.) *Reflections on Creating Learning Organizations*, Waltham, MA: Pegasus Communications
- Kletz, T. 1993. *Lessons from disaster*, Gulf Pub. Co.
- Krause, T.R. 1996. *The Behavior-Based Safety Process: Managing Involvement for an Injury-Free Culture*. 2nd ed., John Wiley & Sons,
- La Porte, T.R. and Consolini, P.M. 1991. Working in practice but not in theory: Theoretical challenges of 'high reliability organizations'. *Journal of Public Administration Research and Theory* 1 (1), 19-47.
- Langewiesche, W. (1998) The Lessons of Valujet 592. *The Atlantic Monthly* (March):
- Lieberman, M.B. and Demeester, L. 1999. Inventory reduction and productivity growth: Linkages in the Japanese automotive industry. *Management Science* 45 (4), 466-485.
- Lischeid, W.E. 1994. TQM & safety: new buzz words or real understanding? *Professional Safety* (June), 31-36.
- Manzella, J.C. 1997. Achieving safety performance excellence through total quality management. *Professional Safety* (May), 26-28.
- Melara, Carlos, Sarriegi, Jose Mari, Gonzalez, Jose J., Cooke, David L., and Sawicka, Agata 2003. Modeling insider attacks to information systems. 21st International Conference of the System Dynamics Society, New York.
- National Safety Council 1995. *Accident Investigation*. 2nd ed., 51 NSC, Itasca, IL.

- Perrow, C. 1984. *Normal Accidents*, New York: Basic Books.
- Perrow, C. 1999. Organizing to reduce the vulnerabilities of complexity. *Journal of Contingencies and Crisis Management* 7 (3), 150-155.
- Petersen, D. 1994. Integrating safety into total quality management. *Professional Safety* (June), 28-30.
- Petersen, D. 1996. *Human Error Reduction and Safety Management*. 3rd ed., John Wiley & Sons,
- Phimister, J.R., Oktem, U., Kleindorfer, P.R. and Kunreuther, H. 2001. Near-miss management systems in the chemical process industry. Working Paper, Wharton Risk Management and Decision Processes Center
- Pidgeon, N. 1997. The limits to safety? Culture, politics, learning and man-made disasters. *Journal of Contingencies and Crisis Management* 5 (1), 1-14.
- Ragan, P.T. and Carder, B. 1994. Systems theory and safety. *Professional Safety* June, 22-27.
- Reason, J. 1997. *Managing the Risks of Organizational Accidents*. Ashgate Publishing, Aldershot, England.
- Regan, G. (Jan 23, 2003) U.S. friendly fire: they never could shoot straight. *The Globe and Mail* Globemedia Publishing Inc.
- Repenning, N.P. and Sterman, J.D. 2001. Nobody ever gets credit for fixing problems that never happened: creating and sustaining process improvement. *California Management Review* 43 (4), 64-88.
- Report of the Presidential Commission on the Space Shuttle *Challenger* Accident (1986) Washington, D.C.: U.S. Government Printing Office.
- Richard, K.P. 1996. *Report of the Westray Mine Public Inquiry*. Province of Nova Scotia.
- Rijpma, J.A. 1997. Complexity, tight-coupling and reliability: connecting normal accidents theory and high reliability theory. *Journal of Contingencies and Crisis Management* 5 (1), 15-23.
- Roberts, K.H. 1990. Some characteristics of one type of high reliability organization. *Organization Science* 1 (2), 160-176.
- Roberts, K.H. and Bea, R. 2001. Must accidents happen? Lessons from high-reliability organizations. *Academy of Management Executive* 15 (3), 70-79.
- Ross, S.M. 2000. *Introduction to Probability Models*. 7th ed., Academic Press, San

Diego, CA.

- Roughton, J. 1993. Integrating a TQM system into safety and health programs. *Professional Safety* (June), 32-37.
- Rowley, C. Letter to FBI Director Robert Mueller, May 21, 2002.
- Rudolph, J.W. and Repenning, N.P. 2002. Disaster dynamics: Understanding the role of quantity in organizational collapse. *Administrative Science Quarterly* 47 (1), 1-30.
- Sagan, S.D. 1993. *The Limits of Safety: Organizations, Accidents and Nuclear Weapons*. Princeton University Press, New Jersey.
- Sawicka, Agata and Gonzalez, Jose J. 2003. Choice under risk in IT-environments according to cumulative prospect theory. 21st International Conference of the System Dynamics Society, New York.
- Senge, P.M. 1990. *The Fifth Discipline: The Art and Practice of the Learning Organization*, New York: Doubleday.
- Seper, J. (Jun 7, 2002) Rowley: 'climate of fear' impeded agents. *The Washington Times*
- Shrivastava, P. 1986. *Bhopal*. Basic Books, New York.
- Sterman, J.D. 2000. *Business Dynamics - Systems Thinking and Modeling for a Complex World*, McGraw Hill.
- Thorne, S. (Jan 20, 2003) Troops experience three near-misses. *Calgary Herald* CanWest Publications Inc.
- Turner, B.A. 1978. *Man-Made Disasters*. Wykeham Science Press, London.
- Turner, B.A. 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, Chicago.
- Weick, K.E. and Sutcliffe, K.M. 2001. *Managing the Unexpected: Assuring High Performance in an Age of Complexity*. Jossey-Bass, San Francisco.
- Weinstein, M.B. 1998. Improving behavior-based safety through TQM. *Professional Safety* (January), 29-33.
- Wolf, F.G. 2001. Operationalizing and testing normal accident theory in petrochemical plants and refineries. *Production and Operations Management* 10 (3), 292-305.