Learning from incidents: from normal accidents to high reliability

David L. Cooke[a]* and Thomas R. Rohleder[a]

David L. Cooke is Adjunct Assistant Professor of Operations Management in the Haskayne School of Business and in the Department of Community Health Sciences at the University of Calgary. He has a PhD and MBA in Operations Management from the University of Calgary and a BSc. in Chemical Engineering from the University of Birmingham, England. Prior to his doctorate, Dr. Cooke enjoyed a long career in the chemical industry in the areas of process engineering, business and product development, technical services management and safety management. He was formerly the Director of Safety and Emergency Planning with NOVA Chemicals Corporation. His research interests include safety management, system dynamics, and OM/OR applications in health care.

Abstract

Many disasters have occurred because organizations have ignored the warning signs of precursor incidents or have failed to learn from the lessons of the past. Normal accident theory suggests that disasters are the unwanted, but inevitable output of complex socio-technical systems, while high-reliability theory sees disasters as preventable by certain characteristics or response systems of the organization. We develop an organizational response system called incident learning in which normal precursor incidents are used in a learning process to combat complacency and avoid disasters. We build a model of a safety and incident learning system and explore its dynamics. We use the model to motivate managers to implement incident learning systems as a way of moving safety performance from normal accidents to high reliability. The simulation model behavior provides useful insights for managers concerned with the design and operation of incident learning systems. Copyright © 2006 John Wiley & Sons, Ltd.

Syst. Dyn. Rev. **22**, 213–239, (2006)

Introduction

On January 28, 1986 seven crew members died when the space shuttle *Challenger* exploded just over a minute after take-off. The Report of the Presidential Commission on the Space Shuttle *Challenger* Incident (1986) concluded that neither NASA nor Thiokol, the seal designer, "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."

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 U.S. Federal Aviation Authority was concerned "about the disproportionate number of infractions committed by ValuJet and the string of small bang-ups it had had."

[a] Haskayne School of Business, University of Calgary, 2500 University Drive NW, Calgary, Alberta, Canada T2N1N4.

* Correspondence to: David L. Cooke. E-mail: dlcooke@ucalgary.ca

Received January 2006; Accepted May 2006

System Dynamics Review Vol. 22, No. 3, (Fall 2006): 213–239
DOI: 10.1002/sdr.338

Thomas R. Rohleder is Professor of Operations Management, Associate Dean (Research), and Director PhD and MBA Thesis Programs in the Haskayne School of Business at the University of Calgary. He has a PhD in Business Administration and a BSB in Finance from the University of Minnesota. Before obtaining his doctorate, Dr. Rohleder worked for ITT Commercial Finance Corp. and Northwestern National Life Insurance Co. in the areas of quantitative analysis and capital equipment financing. His research interests include continuous process improvement, quality management and control, service operations management, and data-driven modeling.

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. 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.

These disasters have at least one thing in common. That is the inability of the organization involved to effectively synthesize and share the information from separate "precursor" incidents with the relevant people across the organization so that appropriate action could be taken to reduce the risk of disaster. We define an *incident* as an unexpected or unwanted change from normal system behavior which causes or has the potential to cause 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 the previous examples appeared to have an effective capability to learn from the precursor incidents. Without an effective *incident learning system*, the precursor incidents are only visible with the benefit of the hindsight that comes from an accident. An incident learning system is the set of organizational capabilities that enable the organization to extract useful information from incidents of all kinds, particularly "near-misses," 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) and to provide an "organization-level perspective" on failure as suggested by Sitkin (1992).

Would all the accidents and disasters previously discussed have been prevented by the implementation of a formal incident learning system? Perhaps not: some accidents may be unpreventable in complex high-risk systems, even where incident learning is occurring. However, in all the cases we discussed there were significant precursor events that, with a formal proactive learning system, would require the attention and action of the organization in a way that apparently did not happen. As further support, Kletz (1993) reports several examples in the chemical industry of the same accident occurring multiple times in the same organization. In one case, a company realized it should build open-sided buildings to house compressors after an explosion at their plant. Several years later another explosion occurred at a sister facility with an enclosed compressor, killing four people. Meanwhile, at the original plant, the walls of a new compressor house had just been completed when management directed workers to tear them down. From all this evidence, we submit that it is not natural for organizations to learn from safety incidents. Even if *ad hoc* learning is occurring, it is not enough.

Our research shows how an incident learning system would work to help prevent disasters, accidents, and the associated losses. We suggest that such a system may help bridge the gap between the disparate theories of normal accidents and high reliability (Rijpma, 1997). The former theory asserts that accidents cannot be averted, while the latter suggests that accidents can be avoided by organizational attention to safety. An incident learning system provides an organizational mechanism to achieve high reliability. It uses the inevitable "normal" incidents that arise in any socio-technical system to reinforce commitment to safety, reduce incident severity and reduce the underlying unsafe conditions that lead to losses. To this end, we develop a dynamic model to show how an incident learning system could reduce an organization's chance for disaster.

Our model is intended to capture the dynamic, structural behavior of an organization that applies incident learning. The complete implementation of an incident learning system requires many complex managerial information, system, and human controls. As with all models, we made a number of simplifying assumptions that do not capture all aspects of a real system in practice. However, our primary purpose in this research is to provide a *theoretical* basis for incident learning systems and provide motivation for managers to consider their implementation. As Sterman and Wittenberg (1999) suggest, "theories of nonlinear, far from equilibrium systems . . . have great potential to illuminate evolutionary behaviour in social, economic, and other human systems."

The theories of normal accidents and high reliability

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 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 a set of holes to line up eventually and the defenses will be breached.

High-reliability theory is a competing organizational theory of accidents whose proponents such as La Porte and Consolini (1991), Roberts and Bea (2001), and Weick and Sutcliffe (2001) believe that, while accidents may be normal, serious ones can be prevented by implementing certain organizational practices. For example, Weick and Sutcliffe (2001) 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. From the perspective of normal accident theory, he argues that the organizational learning required for the success of high-reliability theory will be restricted for several reasons. 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, to promote the necessary learning, it seems clear that a formal organizational system for learning from incidents is required. The theory of incident learning relies on the observation made by Turner (1978) that disasters have long incubation periods during which warning signals (or incidents) are not detected or are ignored. 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 the future. Incident learning is not unlike the continuous improvement cycle described by Repenning and Sterman (2001). An organization effectively implementing a formal incident learning system may evolve into a high-reliability organization over time.

The theory of incident learning

To help understand why incidents happen, and why we need to learn from them, it is useful to introduce the concept of a *risk system*. As shown in Figure 1, it is inseparable from the business system that generates the useful outputs of the organization. However, 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.

Fig. 1. The business and risk systems

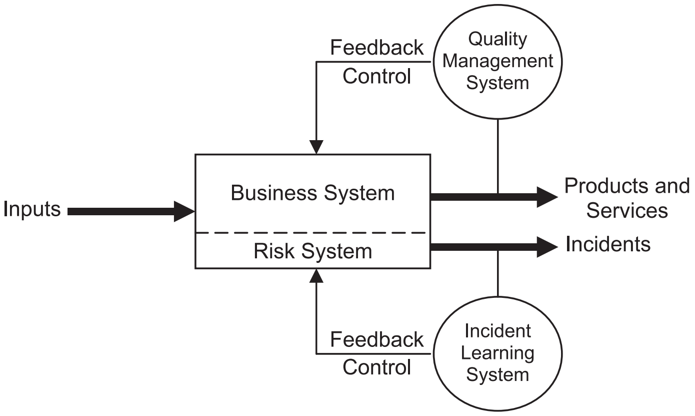

Just as we would apply quality management principles to control the quality of products and services from the business system, so we must apply similar principles to control the "quality" of incidents from the risk system. In fact, it would be equally valid to consider incidents to be "quality problems" 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.

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 learning, 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.

Components of the incident learning system

Learning from incidents is not an entirely new concept (e.g., Sitkin, 1992; Carroll, 1998), but it has not been fully explored as a system for long-term continuous improvement to organizational performance. To our knowledge it has not been considered in a formal dynamic model. Rudolph and Repenning

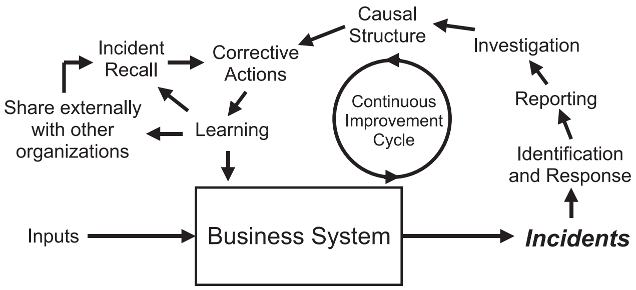

Fig. 2. The incident learning system

(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 relevant warning that an incident learning system will collapse if it becomes overloaded with incidents. To deal with the incident workload, dedicated resources and processes are required to ensure effective learning. As an example, the commercial airlines have these dedicated resources and, as Haunschild and Sullivan (2002) report, learning from incidents is indeed taking place in this industry.

To understand how learning can be facilitated, Figure 2 shows the fundamental components of an incident learning system. We will briefly describe each of these to help clarify how the system works.

Phimister *et al.* (2003) 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 (Vaughan, 1996). 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 next component of incident learning is *reporting*. As the Center for Chemical Process Safety (1989) points 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 or she desires.

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), and the National Safety Council (1995).

These sources suggest 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* (a system model of the causal relationships). 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 have been gathered and evaluated. The desire to find a single root cause was observed by Carroll (1995), who called it "root cause seduction."

Next, it is important to implement *corrective actions* and follow up on all recommendations made by the investigation team. This is particularly true for actions to eliminate systemic causes of incidents, which may span the organization and involve many people in different departments and locations. Processes outside of the incident learning system, such as management of change, audits and inspections, are useful in checking that corrective actions have been successfully implemented without introducing new risks. Completion of incident-related improvements can also be built into management and employee compensation systems.

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.

Finally, it is important to capture and communicate the *learning* from the incident, including the relative success or effectiveness of the corrective actions that were taken. This can be done by distributing a summary report by e-mail, website posting, or other means, and should be directed both locally and centrally. Although we 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 Shell Deer Park, Westray, and the other disasters 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.

Next, we use dynamic modeling to show how this incident learning system can help to reduce the chance for disaster and overall loss associated with organizational incidents.

Modeling the incident learning system

The system dynamics methodology, described in Sterman (2000), helps to focus organizational learning and understanding when dealing with complex systems. We employed this methodology to build a dynamic model of the incident learning system and then used this model to explore the various ideas discussed here. We explain and discuss the important elements of our model and the simulation results obtained, but focus primarily on the management insights to be gained from modeling.[1]

Balancing business goals and safety goals

Figure 3 represents a generic business system designed to achieve desired productive outputs. We use the term *Productivity* to designate a desired business objective; however, the construct is intended to represent any organizational goal (e.g., meeting a production goal, achieving a flight schedule, retaining patient flows while reducing costs). If the organization is not achieving its goal, pressure is increased to improve *Productivity*, a stock representing the current productive level in the organization. *Productivity Pressure* is relaxed as *Productivity* nears the goal; thus the feedback loop is balancing. In general, losses due to safety incidents lead to a loss of productive capacity, creating links between the safety system and the business system. Incidents that lead to accidents and disasters require time and resources to recover from, but even

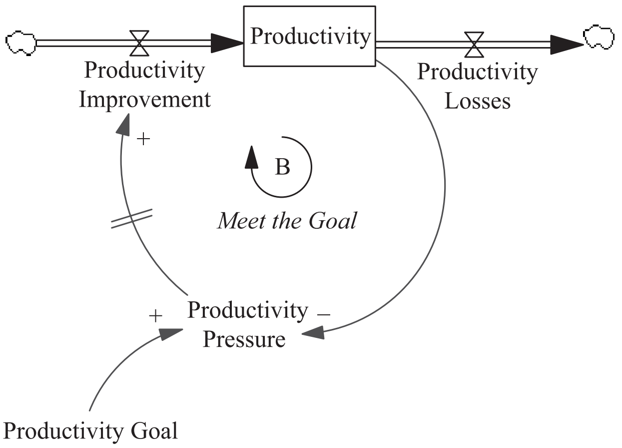

Fig. 3. The productive organizational system (The +'s and –'s represent the polarities of the causal links and read as follows: + means variables moving in the same direction, *ceteris paribus*; – means variables moving in opposite directions, *ceteris paribus*). The double hatches on lines represent system delays

Fig. 4. Balancing productivity and safety

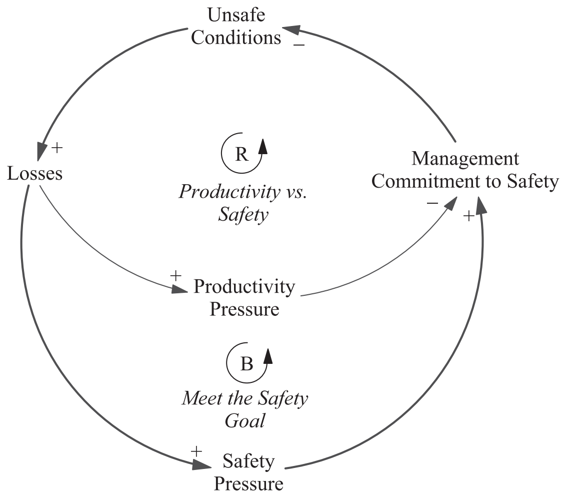

near-miss incidents will usually hurt productivity. For instance, if a nurse initially prepares the wrong medicine for a patient, but catches the mistake before administering it, the organization still incurs the cost of the incorrect preparation (Tucker, 2004). Such *Productivity Losses* will affect the level of *Productivity* and eventually increase *Productivity Pressure*. Figure 4 is a causal loop diagram showing the balance that must be maintained between productivity and safety. In the ideal world we have a balancing loop represented by the outer circle in Figure 4. As the number of *Unsafe Conditions* increase, *Losses* increase, causing more pressure on *Management Commitment to Safety*. This increased commitment translates into safety improvements that will eventually reduce *Unsafe Conditions* and restore safety performance. Unfortunately, in many real-world situations, short-term pressures to maintain the production schedule, keep the project on track, etc. may cause risky short cuts to be made or unsafe conditions to be tolerated. Management may not set out to deliberately compromise safety in order to reach productivity goals, but often this happens as productivity pressure increases (Boyd, 2001). For example, the underlying causes of the Westray mine disaster (Cooke, 2003) were the production pressures to satisfy contractual production targets, which caused management to ignore unsafe conditions in the mine, and the political pressure to maintain economic activity at the mine, which caused the safety inspectorate to overlook management's poor safety record. In Figure 4 this situation is represented by the inner circle in which increasing *Losses* lead to increasing *Productivity Pressure*, which causes *Management Commitment to Safety* to decrease as management's attention is diverted to productivity improvement issues. In feedback terms, this is a reinforcing loop because more productivity pressure leads to lower safety performance, more losses, and yet more productivity pressure. This "disaster spiral" can be averted only if management diverts its attention back to safety improvement before it is too late.

In Figure 4 the term "losses" can include production outages, quality problems, injuries to workers, environmental spills, etc. Safety losses lead to productivity losses due to the loss of morale, worker capacity (injuries, etc.), and myriad other causes. Productivity impacts capture management's attention in the form of reduced production, reduced services, or reduced profit. On the other hand, the relationships between safety losses and productivity may not be quite so visible to management. Further compounding the safety visibility problem is the common organizational practice of trouble-shooting or "fire-fighting" problems at the worker or operations team level (Repenning and Sterman, 2001; Tucker, 2004; Edmondson, 1996). Not only may management not get to hear about low-level problems for these reasons, but their systemic causes may be difficult to identify. Even if the operations team can identify the systemic cause of a safety problem, their ability to improve the system is often not under their control but is rather under the control of management.

The importance of incident reporting

Management relies on an incident reporting system to bring safety problems to their attention. Incident reporting is an important and necessary part of the incident learning system and, based on the discussion above, is modeled using the relationships shown in Figure 5.

The model shown in Figure 5 may be compared to that shown in Figure 2 by noting that, for simplicity, we have incorporated identification and response into *Reported Incidents* and incorporated the determination of causal structure

Fig. 5. The incident reporting system

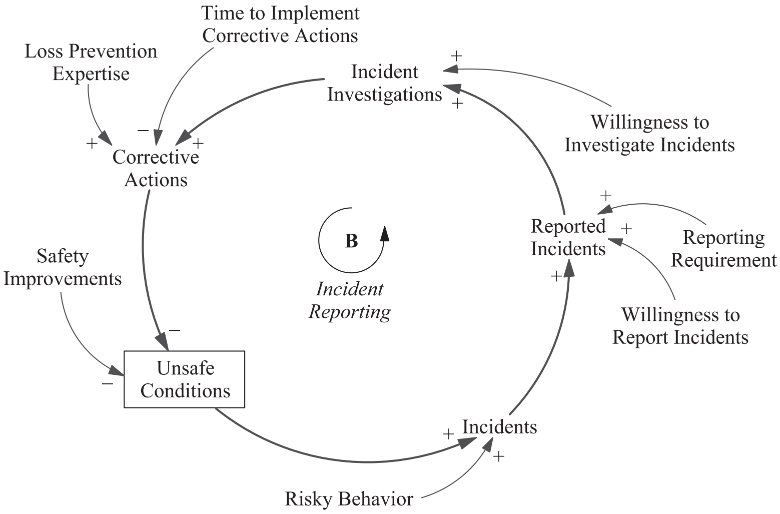

Fig. 6. The stock of unsafe conditions

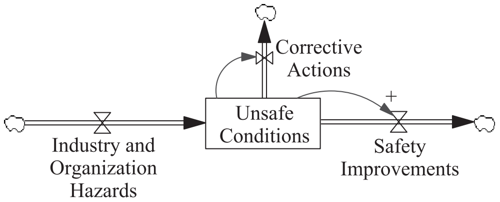

into *Incident Investigations*. The variable *Unsafe Conditions* is a stock that accumulates the latent failures that the risk system generates. Another way of looking at it is that unsafe conditions arise normally from the operation of a complex system. In our model, shown in Figure 6, we assume that *Industry and Organization Hazards* arise continuously from organizational change, new technologies, process hazards, etc., at a rate determined exogenously by the nature of the industry, to create a stock of *Unsafe Conditions*. The level of *Unsafe Conditions* can be reduced by *Corrective Actions* arising from *Incident Investigations* or by *Safety Improvements* arising proactively from *Management Commitment to Safety*. In the model, these relationships are represented by flows into and out of the stock of *Unsafe Conditions*.

Incidents arise from the interaction between *Unsafe Conditions* and the *Risky Behavior* of the workers in the system. Whether or not these incidents lead to corrective actions that improve the system (by reducing the number of unsafe conditions) will depend on the effectiveness of the feedback loop in Figure 5. How many of these incidents are reported to management will depend on the willingness of the workers involved and the *Reporting Requirement* set by management. If management creates a low severity threshold, many incidents will be reported. If management sets a high threshold, then only the more severe incidents will be reported. In our model, the workers' *Willingness to Report Incidents* depends on their relative commitment to safety and on the quality of the incident reporting system. If the commitment is low or the quality is poor, then few incidents will be reported. Similarly, management's *Willingness to Investigate Incidents* will depend on their own relative commitment to safety and on the availability of resources.

Finally, whether or not the investigated incident leads to corrective actions being taken will depend on the organization's *Loss Prevention Expertise*. For example, if the organization is not able to properly identify the systemic causes of an incident then corrective actions may not be effective; if the organization does not have a good process for managing change, then actions taken to correct one problem may introduce new problems that were not foreseen.

Safety leadership

Management acts as a safety leadership role model for employees. If management demonstrates a commitment to safety through their words and deeds then eventually

Fig. 7. Safety leadership

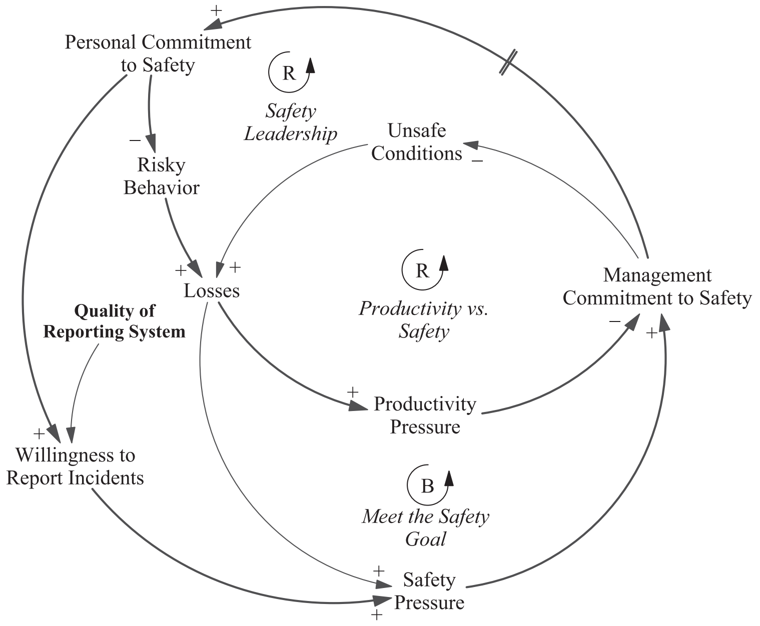

this will translate into a higher *Personal Commitment to Safety* on the part of employees. As shown in Figure 7, this role-modeling behavior helps to reinforce the balance between productivity and safety that management has struck.

To see this, consider what happens when *Management Commitment to Safety* goes up. After a delay, *Personal Commitment to Safety* also goes up. Via the outer feedback loop, this leads to a greater *Willingness to Report Incidents* and higher *Safety Pressure* to reinforce *Management Commitment to Safety*. Via the inner feedback loop, this leads to less *Risky Behavior*, causing lower *Losses* and less *Productivity Pressure*. Less productivity pressure allows management to maintain its focus on safety. The same type of positive reinforcement occurs if management commitment to safety is falling. However, research shows that incidents are typically *not* reported (for example, see Cullen *et al.*, 1995). Reasons for not reporting incidents include a fear of punishment, bureaucratic or confusing reporting requirements, or quite simply a desire not to interrupt the work flow.

We capture this effect in our model with the *Quality of Reporting System* parameter. If this parameter is equal to zero, then no incidents are reported, no safety improvements take place, and productivity eventually drops to zero. If this parameter is equal to unity, then all incidents are reported if the worker is willing to report them. More likely, this variable would have a value somewhere in between the two extremes. Figure 8 compares the model results for *Productivity* and *Management Commitment to Safety* for $R = 0$, 0.5 and 1.0, where R is the *Quality of Reporting System* parameter. H is the *Industry and*

Fig. 8. Effect of *Quality of Reporting System* parameter

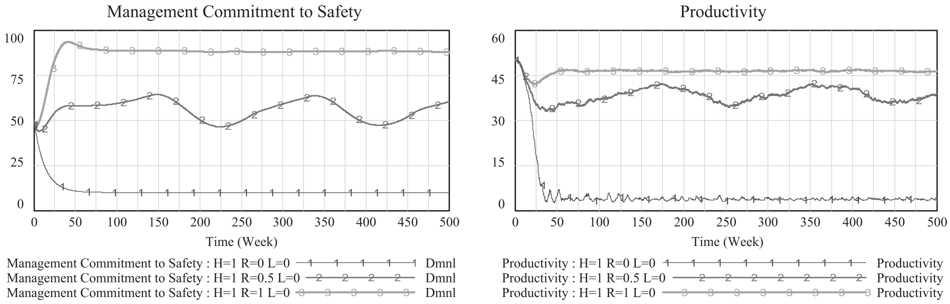

Organization Hazards parameter, set at 1 Condition/Week, and L is the *Organizational Ability to Learn* parameter, which will be discussed in the next section. The model behaves as we would expect, with the intermediate value of $R = 0.5$ demonstrating the cyclical behavior associated with safety campaigns. The extreme values of $R = 0$ and $R = 1$ represent "disaster potential" and "safety excellence," respectively.

Overall, our model shows that with an organization's desire to improve business performance safety results may worsen. Exhortations to improve safety may temporarily improve performance, but our model and the cases in the Introduction suggest that accidents and perhaps disasters will occur unless organizations are able to *sustain* their commitment to safety. The problem with the incident reporting model shown in Figure 5 is that the process needs to be driven by individuals' willingness to report incidents and management's willingness to investigate them. Otherwise, the balancing loop will settle down and management will become complacent when few incidents are reported, thinking that safety has improved when in fact it has not. To avoid this, we propose a sustained organizational effort to achieve incident learning. The goal of incident learning is to reinforce incident reporting by keeping the memory of past incidents alive.

In our model, the value of incident learning in an organization is accumulated into a stock of *Lessons Learned*. This stock is much like the organizational memory discussed by Anand *et al.* (1998) where the lessons learned are maintained via formal and informal mechanisms. The creation of this stock via organizational learning processes has been discussed by many authors, including Argote (1999). The stock and flow structure of this process is shown in Figure 9.

Lessons learned via the reporting–investigation–corrective actions cycle discussed above accumulate as a stock of lessons available to the organization to manage future incidents, and are "lost" at the rate at which an organization "forgets" these lessons. The memory loss is like the accidental forgetting discussed in de Holan *et al.* (2004) where memory decay erodes the knowledge (lessons) accumulated in the organization. *Incident Learning* is the rate at

Fig. 9. Stock and flow structure of incident learning

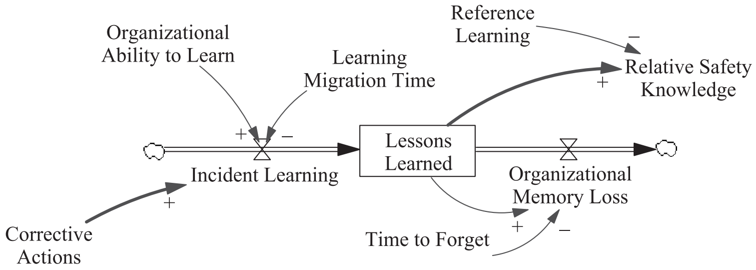

which lessons are learned and is driven by the product of the number of *Corrective Actions* and the *Organizational Ability to Learn*. Since it takes time to assess the effectiveness of corrective actions, and it takes time for the learning from the incidents to migrate across the organization, changes in *Incident Learning* are exponentially smoothed over a time defined by the *Learning Migration Time*. The net result of the learning process is that benefits from *Corrective Actions* are transformed into an increase in the *Relative Safety Knowledge* of the organization compared to what we call *Reference Learning* or the base level of safety knowledge. It is this relative safety knowledge or "memory of incidents" that can sustain management commitment to safety in the face of productivity pressure and, on the other hand, can overcome management complacency when things are going well. A paradox of incident learning is that incidents cause accidents and disasters, yet they are needed for learning to occur. The solution to this paradox is to recognize that every incident has a different severity or magnitude of loss. In fact, incidents that are near misses may have no loss, but the information from the incident can still be used to improve safety. Thus, incident learning reduces *Unsafe Conditions* by increasing *Relative Safety Knowledge*, which increases *Management Commitment to Safety*, motivates more *Safety Improvements*, and promotes organizational *Willingness to Investigate Incidents*. Over time, increasing *Management Commitment to Safety* also increases the *Personal Commitment to Safety* of individual workers, which supports their *Willingness to Report Incidents* and discourages their *Risky Behavior*. Increasing *Relative Safety Knowledge* can also reduce the severity of *future* incidents because of increased precautions, improved procedures, and safer designs that were implemented as a result of lessons learned from past incidents. Supporting this idea that information from incidents can be used in this way to improve safety, Edmondson (1996) found in a study of safety in nursing units that higher-performing units reported more errors (incidents) than did lower-performing units.

The incident learning system

The complete model of the incident learning system is shown in Figure 10. For simplicity of presentation, the "safety leadership" loop involving the

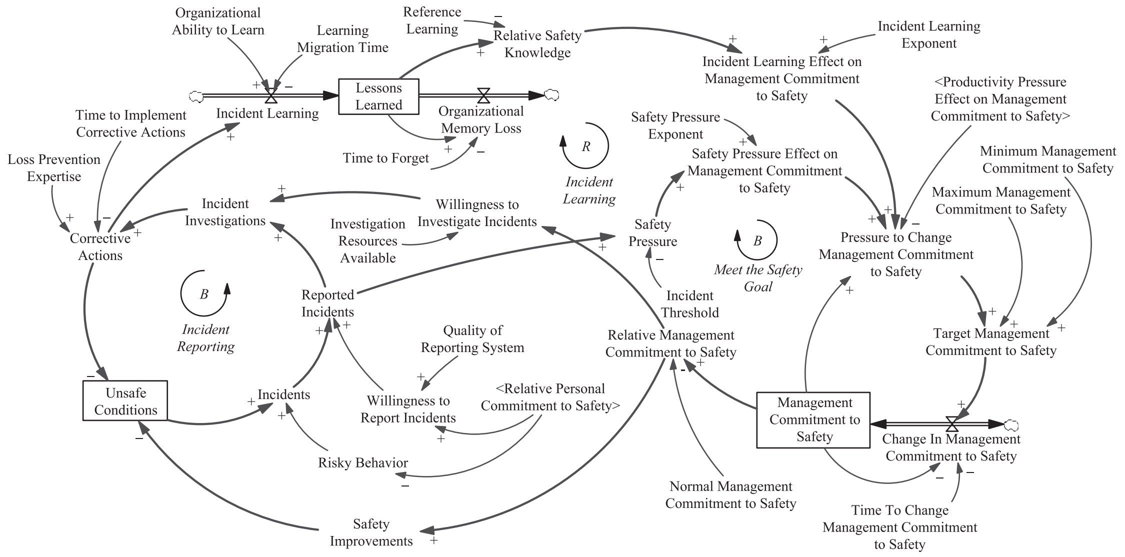

Fig. 10. The incident learning system

effect of *Management Commitment to Safety* on *Personal Commitment to Safety* has been omitted, but is shown by the dotted arrow. The dynamic structure for *Personal Commitment to Safety* is the same as for *Management Commitment to Safety*, and *Productivity Pressure* affects both of them. Note that the effect of incident learning and of various "pressures" on commitment to safety is modeled using a "learning curve" model of the form $C = C_0(P/P_0)^n$, where P is pressure or learning, C is commitment, and n is the learning exponent. The rationale for using a learning model is that safety and productivity pressures accumulate over time like experience does and that commitment (learning) is expected to change by a given percentage with each doubling of pressure (experience). See Sterman (2000) for several examples of learning curve models.

The "cycle" from *Lessons Learned* and back again is a reinforcing loop that represents incident learning. It follows two paths:

1. Through *Management Commitment to Safety* and *Willingness to Investigate Incidents.*
2. Through *Management Commitment to Safety*, *Personal Commitment to Safety*, and *Willingness to Report Incidents*.

In the long run, if organizational conditions are stable, this system leads to steady-state results centering on a consistent stock of lessons and level of productivity. From Figure 10, it is evident that every component in the cycle must be working: incidents must be identified, reported and investigated, corrective actions must be taken, and lessons learned from this process must be communicated and taught to the rest of the organization.

Table 1. Parameter values

Model parameter	Value
Productivity Evaluation Time	4 weeks
Time to Forget	52 weeks
All other System Delays	13 weeks
Quality of Reporting System	0.50
Minimum/Normal/Maximum Commitment to Safety	10/50/100%
Reference Learning	2 learning
All learning exponents	0.4
Loss Prevention Expertise	1 condition/incident
Investigation Resources Available	1
Industry and Organization Hazards	1 condition/week
Incident Threshold	0.25 incident/week
Productivity Goal	50%

Incident learning system dynamics

In this section we explore the dynamics of the incident learning system. We use the model parameter values in Table 1 unless otherwise stated. The stocks *Unsafe Conditions* and *Lessons Learned* were initialized with a value of zero, while *Productivity*, *Management Commitment to Safety* and *Personal Commitment to Safety* were initialized with a value of 50%.

Model validation

The model was validated by extreme condition testing, dimensional consistency, and assessment of whether model structure and parameter values are consistent with what is known about real-world safety systems and literature findings. The results for various extreme condition tests are shown in Table 2.

While helping to validate a model, extreme condition testing also helps to reveal the limitations of a model. For example, high and low values of the *Incident Threshold* produce the expected response given the simple structure of the model. However, in the real world, setting the *Incident Threshold* too low may indeed motivate more management commitment to safety improvements, but it could also motivate *less* incident reporting as individuals seek to "hide incidents" so as to meet management's goal and receive a reward.

Simulation results

To show the dynamics of the model we will start the base case simulation at time zero with a *Lessons Learned* level and *Unsafe Conditions* level of zero, and an "average" incident reporting system (*Quality of Reporting System* $R = 0.5$). In the second case, we add to this an incident learning system at time zero (*Organizational Ability to Learn* $L = 1$). The safety system response is

Table 2. Results for extreme condition tests

Model parameter	Test	Value	Test result[a]
Industry and Organization Hazards	Zero hazards	0	Commitment to safety drops to minimum but there are no losses, so no effect on productivity
Industry and Organization Hazards	Minimal hazards	0.01	Commitment to safety drops to minimum but losses cause productivity to drop from 50 to ~45
Industry and Organization Hazards	Maximum hazards	10	With $R = 1$ and $L = 0$, productivity falls to minimum. With incident learning, ($L > 0.1$) productivity approaches goal
Quality of Reporting System	No incidents reported	0	Even with full incident learning ($L = 1$), no learning takes place because no incidents are reported
Initial value for Commitment to Safety	Minimum	10	With $R = 1$ and $L = 0$, productivity falls to minimum. With incident learning, ($L > 0.1$) safety commitment recovers
Loss Prevention Expertise	None	0	No corrective actions are taken; greater amplitude swings in Management Commitment to Safety (exhortations to make safety improvements followed by complacency as safety improves)
Investigation Resources Available	None	0	Same result as above
Incident Threshold	Very high	10	Commitment to safety quickly drops to minimum values
	Very low	0.1	Commitment to safety quickly rises to maximum values

[a] Tests were done with the parameter values as in Table 1 and $L = 0$, unless specified. R is the *Quality of Reporting System* parameter and L is the *Organizational Ability to Learn* parameter.

shown in Figure 11. Incident reporting and corrective actions cause the rising unsafe conditions to be brought under control ($L = 0$ case), but the additional lessons learned enable unsafe conditions to be reduced ($L = 1$ case) and stable, non-cyclical, conditions to be restored.

Lessons Learned 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 *Safety Improvements* are made, and the lower the *Unsafe Conditions*. These relationships are consistent with the evidence provided by Kletz (1993) that accidents recur in the absence of organizational memory.

The other side of the story is the effect that incident learning has on *Productivity*. Figure 12 shows the comparison of *Safety Improvements* and

Fig. 11. Incident learning system response

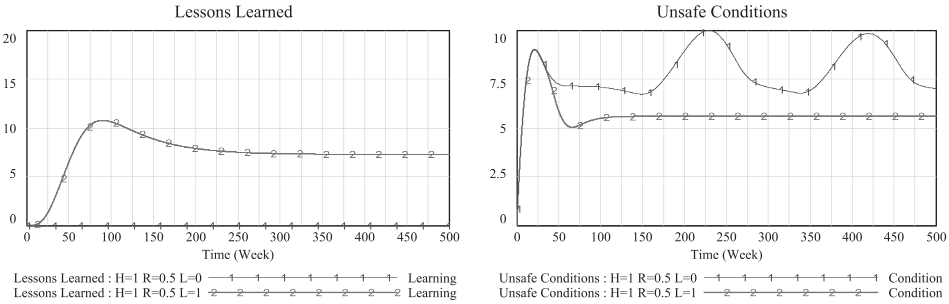

Fig. 12. Effect of incident learning on safety improvements and productivity

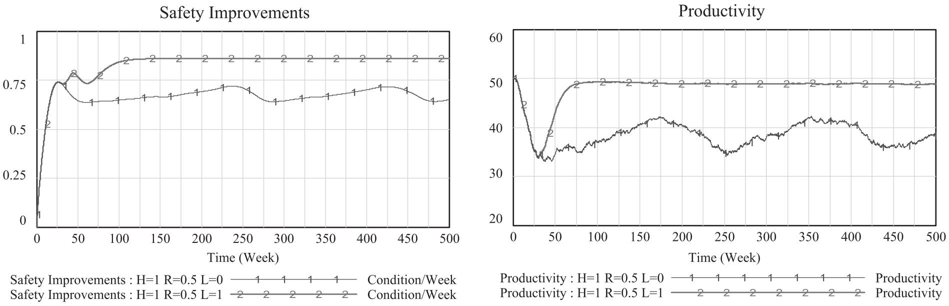

Fig. 13. Effect of step changes in productivity goal

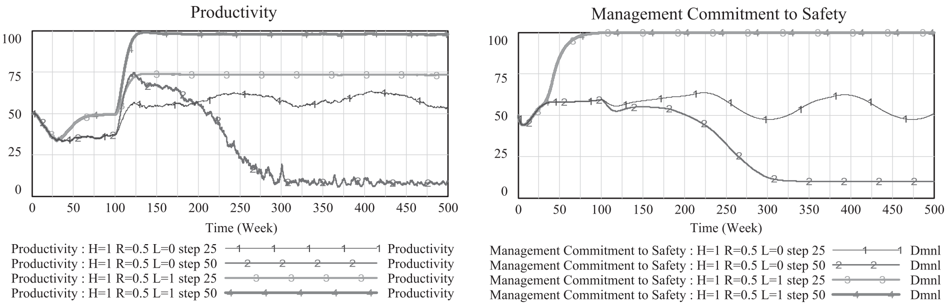

Productivity with and without *Incident Learning*. With *Incident Learning* ($L = 1$), the initial behavior is the same as the Base Case. However, once the *Lessons Learned* and *Safety Improvements* take effect, *Productivity* moves closer to the goal set by management. Thus, the motivation for organizations to implement learning systems is based on performance improvement to both the safety and business systems.

Incident learning reinforces *Management Commitment to Safety* so that system response is robust to changes in the *Productivity Goal*. This is shown in Figure 13 by the response to step changes in the *Productivity Goal* from 50 to 75 (step 25) and from 50 to 100 (step 50). The system with the larger step change and no incident learning ($L = 0$, step 50), eventually collapses as *Management Commitment to Safety* gives way to *Productivity Pressure*.

Implications for safety management

The model of the incident learning system discussed in the previous section could be made much more complicated without adding a lot of value. Its strengths and weaknesses could also be critiqued, as could any of the assumptions on which the model is based. However, the true value of this model is to highlight the incident learning system as a continuous improvement process, to promote discussion of barriers to organizational learning and to suggest ways in which the learning process can be strengthened and improved. We offer the following suggestions for strengthening the incident learning system, which may go some way towards addressing the barriers to organizational learning identified by authors such as Sagan (1993) and Rijpma (1997):

1. Management should not be discouraged by those who dismiss the creation of a safety culture as a myth or fantasy. The simulations show that it could take several years to build a safety culture, as represented by variables such as *Management Commitment to Safety* and *Personal Commitment to Safety*, and that this culture can wax and wane as productivity pressures come and go. We suggest that safety culture can be reinforced by implementing 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. Implementing such a system would go a long way towards demonstrating that management can "walk the talk" in terms of safety commitment.
2. Organizations should put their focus on determining the causal structure of an incident in the context of the overall business system, rather than on finding a single "root cause." Obviously the extent of such an analysis will depend on the severity of the incident. Incident investigation teams should be trained and equipped with the tools of systems thinking as set forth by Senge (1990) and Sterman (2000). 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.
3. Organizations should implement a reward system that encourages reporting of incidents and implementation of corrective actions. These are two important 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 unachievable safety targets such as "zero spills" or "no 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 or her colleagues were expecting to get a bonus for one million person-hours without a recordable injury. Eliminating the blame game is difficult, but it can be done by following the 14 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).

4. The importance of an incident learning system for strengthening risk communications cannot be overemphasized. Managers and supervisors should discuss lessons learned 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 findings from incidents reported in the previous week. Conversely, communicating knowledge gained 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 do not specifically mention learning from incidents, the reporting and problem-solving processes required for incident learning 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. (For example, see Jones *et al.*, 1999, for information on the European Commission's Major Accident Reporting System; Selby, 2003, for a description of the safety alert database and information exchange for safety in the U.K. offshore drilling industry; and Haunschild and Sullivan, 2002, who analyzed incident data shared with the National Transportation Safety Board by the U.S. commercial airline industry).
5. The external loop for shared learning, shown in Figure 2, should be supported by a benchmarking process that analyses best practices in other organizations and adapts these practices to improve the business system. Sharing knowledge from an incident externally encourages other organizations not only to learn from the incident itself but to share what they have learned. Other organizations may contact the organization at which the incident occurred to say, "We have experienced that particular problem and we are willing to share our solution with you."
6. 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, the overall level of incident severity, and success in completing corrective actions. The possibility of people fabricating incidents to "look good" is very unlikely compared to the possibility of people not reporting real incidents in order to avoid blame.

7. 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 explored by Sawicka and Gonzalez (2003) 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."
8. Management should use climate surveys and other feedback tools to measure management and personal commitment to safety and incident learning. The survey and feedback information can be useful in designing and implementing policies and procedures to encourage a proactive learning response to incident reporting. It could also be used by the CEO and Board of Directors to monitor the "safety pulse" of the organization. Support for such a communication system can be found in Grabowski and Roberts (1997), who note that "the necessity for good communication cannot be overemphasized" and that "a strong organizational culture—and implementation of norms that reinforce that culture—is an important risk measure." Climate surveys are an important tool for assessing organizational culture.
9. 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 of information in databases is not enough. There should also be a "management of change" process for modifying facilities, upgrading operating or maintenance procedures, engineering standards etc. based on learning from incidents. For more

information on management of change see, for example, the guidelines published by the Center for Chemical Process Safety (1989).

10. Finally, many organizations will find opportunities to integrate quality, safety, environment, health, security, 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, perhaps by building on economic models developed for quality improvement.

Conclusions

We suggest than an incident learning system can help to 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 frequency and 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 in part on the analysis and simulation of the Westray mine disaster (Cooke, 2003), which demonstrated that the fatal explosion was not a spontaneous event, but a consequence of a chain of events or precursor incidents that could have been detected by an effective incident learning system, thereby breaking the chain. We have shown that "disaster" can result from productivity pressures in a simple system model, and that disaster can be averted by learning from the precursor incidents. Given the importance of disaster prevention to the continuing survival of a socio-technical organization, we would suggest that an incident learning system should be just as central to an organization's mission as its production or service delivery system.

As pointed out in Sterman (2002), all models are “wrong” and we agree with this sentiment in that our incident learning model is not a perfect reflection of reality. The usefulness of the model is in the insights applicable to the real systems being studied. In this regard, we believe our model is useful.

The model shows that organizational losses from incidents can be reduced dramatically by focusing on the learning cycle that reduces unsafe conditions and the severity of incidents. By maintaining a “stock” of lessons learned about past incidents, future accidents and disasters can be averted. The value of this approach is no more evident than in NASA’s most recent shuttle accident. Indeed, the Columbia Accident Investigation Board (2003) found that the most likely explanation for the Columbia disaster was a large chunk of foam breaking away from the solid rocket booster during launch and causing severe damage to the thermal protection tiles on the wing of the shuttle. Incidents involving foam breaking away from the solid rocket boosters had happened many times in the past, but the Investigation Board found that NASA had normalized the risk and were operating under the perception that a disaster was unlikely. Our model shows that high levels of unsafe conditions can persist, especially when an organization is focused on a business system goal (like a shuttle launch). Our model also shows that learning from all incidents, regardless of whether they lead to a significant loss, can lead to higher safety levels and better achieve the business outcomes.

In terms of future research there is ample opportunity to explore the value of incident learning in real environments. There has been little or no evidence provided in the literature of the quality of the incident learning systems, if any, being in operation at the other organizations experiencing the disasters cited in the Introduction. 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. Of course, hindsight is 20/20, and unfortunately, the absence of something does not demonstrate that its presence would make a difference. While our model shows the potential value of incident learning, there is little or no empirical evidence in the literature showing whether or not an incident learning system makes a difference. In particular, industry studies of normal accident theory like that of Sagan (1993) or Wolf (2001) should be extended to specifically explore the effectiveness of the organizations’ incident learning systems.

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 disasters that would happen 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. While it is probably true that an effective incident learning system is more likely to be found at a petrochemical plant than at a farm or construction site, further research is needed to determine which industries have implemented systems for incident learning. Furthermore, more research is needed to test the hypothesis that the operating performance of individual companies and plants within those industries will be related to the organizational effectiveness in learning from incidents.

The authors are currently engaged in the implementation of an incident learning system at a health care organization. Studying this implementation will be the start of researching the following propositions which we expect to hold true:

1. Given the same degree of industry and organizational risk, 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.

Finally, we have discussed many barriers to the implementation of an effective incident learning system in an organization, and suggested ways in which they can be overcome. There are also many external barriers, not the 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. For example, the Louisiana State University Law Center website[2] states that there "is often a non-attorney 'risk manager' who sends only 'important' incident reports to the attorney. The problem is that the incident reports not sent to the attorney will be discoverable." As the results of our research suggest, this legal climate needs to change for long-term systemic improvement and accident prevention to occur. We would hope that the models and ideas discussed in this paper would help motivate organizations to remove barriers to learning from incidents, since doing so will improve not only safety performance but also general business performance.

Notes

1. The model is built in the Vensim modeling language; see www.vensim.com for more details. Readers who are interested in details of the model should contact the authors to obtain a copy.
2. See http://biotech.law.lsu.edu/Books/aspen/Aspen-LEGAL.html and http://biotech.law.lsu.edu/Books/aspen/Aspen-INCIDENT.html for a discussion of legal privilege and incident reports in the context of medical care delivery. [Accessed 12 September 2006].

References

Anand V, Manz CC, Glick WH. 1998. An organizational memory approach to information management. *Academy of Management Review* **23**: 796–809.

Argote L. 1999. *Organizational Learning: Creating, Retaining, and Transferring Knowledge*. Kluwer Academic: Boston, MA.

Bird FE, Germain GL. 1986. *Practical Loss Control Leadership*. Institute Publishing: Loganville, GA.

Boyd C. 2001. HRM in the airline industry: strategy and outcomes. *Personnel Review* **30**: 438–453.

Carroll JS. 1995. Incident reviews in high-hazard industries: sensemaking and learning under ambiguity and accountability. *Industrial and Environmental Crisis Quarterly* **9**: 175–197.

Carroll JS. 1998. Organizational learning activities in high hazard industries: the logics underlying self-analysis. *Journal of Management Studies* **35**: 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 for Chemical Engineers: New York.

Columbia Accident Investigation Board. 2003. *Report on the Columbia Space Shuttle Accident*. U.S. Government Printing Office: Washington, DC.

Cooke DL. 2003. A system dynamics analysis of the Westray mine disaster. *System Dynamics Review* **19**: 139–166.

Cullen DJ, Bates DW, Small SD, Cooper JB, Nemeskal AR, Leape LL. 1995. The incident reporting system does not detect adverse drug events: a problem for quality improvement. *Joint Commission Journal of Quality Improvement* **21**(10): 541–548.

de Holan PM, Phillips N, Lawrence TB. 2004. Managing organizational forgetting. *MIT Sloan Management Review* **45**: 45–51.

Deming WE. 1989. *Out of the Crisis* (9th edn). Massachusetts Institute of Technology Center for Advanced Engineering Study: Cambridge, MA.

Edmondson AC. 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*. United States Environmental Protection Agency: Washington, DC.

Gonzalez JJ, Sawicka A. 2003. The role of learning and risk perception in compliance. In: JJ Gonzalez (ed.) *From modeling to managing security: A system dynamics approach*. Norwegian Academic Press: Kristiansand, Norway.

Grabowski M, Roberts KH. 1997. Risk mitigation in large-scale systems: lessons from high reliability organizations. *California Management Review* **39**(4): 152–162.

Haunschild PR, Sullivan BN. 2002. Learning from complexity: effects of prior accidents and incidents on airlines' learning. *Administrative Science Quarterly* **47**: 609–643.

Jones S, Kirchsteiger C, Bjerke W. 1999. The importance of near miss reporting to further improve safety performance. *Journal of Loss Prevention in the Process Industries* **12**: 59–67.

Kim DH. 1994 Managing organizational learning cycles. In *Reflections on Creating Learning Organizations*. Pegasus Communications: Waltham, MA; 43–52.

Kletz T. 1993. *Lessons from Disaster*. Gulf Publishing: Houston, TX.

La Porte TR, Consolini PM. 1991. Working in practice but not in theory: theoretical challenges of "high reliability organizations". *Journal of Public Administration Research and Theory* **1**: 19–47.

Langewiesche W. 1998. The Lessons of ValuJet 592. *The Atlantic Monthly* **281**(3): 81–98, March 1998.

Lieberman MB, Demeester L. 1999. Inventory reduction and productivity growth: linkages in the Japanese automotive industry. *Management Science* **45**(4): 466–485.

National Safety Council.1995. *Accident Investigation*. NSC: Itasca, IL.

Perrow C. 1984. *Normal Accidents*. Basic Books: New York.

Perrow C. 1999. Organizing to reduce the vulnerabilities of complexity. *Journal of Contingencies and Crisis Management* **7**: 150–155.

Phimister JR, Kleindorfer PR, Kunreuther H. 2003. Near-miss incident management in the chemical process industry. *Risk Analysis* **23**: 445–459.

Ragan PT, Carder B. 1994. Systems theory and safety. *Professional Safety* June: 22–27.

Reason J. 1997. *Managing the Risks of Organizational Accidents*. Ashgate Publishing: Aldershot, U.K.

Repenning NP, Sterman JD. 2001. Nobody ever gets credit for fixing problems that never happened: creating and sustaining process improvement. *California Management Review* **43**: 64–88.

Richard KP. 1996. *Report of the Westray Mine Public Inquiry*. Province of Nova Scotia, Canada.

Rijpma JA. 1997. Complexity, tight-coupling and reliability: connecting normal accidents theory and high reliability theory. *Journal of Contingencies and Crisis Management* **5**: 15–23.

Roberts KH, Bea R. 2001. Must accidents happen? Lessons from high-reliability organizations. *Academy of Management Executive* **15**: 70–79.

Rudolph JW, Repenning NP. 2002. Disaster dynamics: understanding the role of quantity in organizational collapse. *Administrative Science Quarterly* **42**: 1–30.

Sagan SD. 1993. *The Limits of Safety: Organizations, Accidents, and Nuclear Weapons*. Princeton University Press: Princeton, NJ.

Sawicka A, Gonzalez JJ. 2003. Choice under risk in IT-environments according to cumulative prospect theory. In: JJ Gonzalez (ed.) *From modeling to managing security: A system dynamics approach*. Norwegian Academic Press: Kristiansand, Norway.

Selby C. 2003. Steeling a march on offshore safety. *Chemical Engineer* June: 34–35.

Senge PM. 1990. *The Fifth Discipline: The Art and Practice of the Learning Organization*. Doubleday: New York.

Sitkin S. 1992. Learning through failure: the strategy of small losses. In *Research in Organizational Behavior*. JAI Press: Greenwich, CT: 231–266.

Sterman JD. 2000. *Business Dynamics: Systems Thinking and Modeling for a Complex World*. Irwin/McGraw-Hill: Boston, MA.

Sterman JD. 2002. All models are wrong: reflections on becoming a systems scientist. *System Dynamics Review* **18**: 501–531.

Sterman JD, Wittenberg J. 1999. Path dependence, competition, and succession in the dynamics of scientific revolution. *Organization Science* **10**: 322–341.

Report of the Presidential Commission on the Space Shuttle *Challenger* Incident. 1986. U.S. Government Printing Office: Washington DC.

Tucker AL. 2004. The impact of operational failures on hospital nurses and their patients. *Journal of Operations Management* **22**: 151–169.

Turner BA. 1978. *Man-Made Disasters*. Wykeham Sciences Press: London.

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

Weick KE, Sutcliffe KM. 2001. *Managing the Unexpected: Assured High Performance in an Age of Complexity.* Jossey-Bass: San Francisco.

Wolf FG. 2001. Operationalizing and testing normal accident theory in petrochemical plants and refineries. *Production and Operations Management* **10**: 292–305.